SOCIAL FEEDBACK: SOCIAL LEARNING FROM INTERACTION HISTORY TO SUPPORT INFORMATION SEEKING ON THE WEB

A Thesis Submitted to the College of
Graduate Studies and Research
In Partial Fulfillment of the Requirements
For the Degree of Doctor of Philosophy
In the Department of Computer Science
University of Saskatchewan
Saskatoon, Saskatchewan, Canada

By

SCOTT BATEMAN

Permission to Use

PERMISSIONS

- Figure 2.7 – Reprinted from [54], with permission from Elsevier.

- Figure 2.8 – Reprinted from [70], with permission from ACM.

- Figure 2.11 and Figure 2.12 – Reprinted from [167] with permission from the IEEE, © 2007 IEEE.

- Figure 2.10 – Reprinted from [156], with permission from ACM.

- Figure 2.16 – Reprinted from [136] with kind permission from Springer Science and Business Media.

- Figure 2.17 – Reprinted from [106] with kind permission from Springer Science and Business Media.

ABSTRACT

Information seeking on the Web has become a central part of many daily activities. Even though information seeking is extremely common, there are many times when these tasks are unsuccessful, because the information found is less than ideal or the task could have been completed more efficiently. In unsuccessful information-seeking tasks, there are often other people who have knowledge or experience that could help improve task success. However, information seekers do not typically look for help from others, because tasks can often be completed alone (even if inefficiently). One of the problems is that web tools provide people with few opportunities to learn from one another's experiences in ways that would allow them to improve their success.

This dissertation presents the idea of social feedback. Social feedback is based on the theory of social learning, which describes how people learn from observing others. In social feedback, observational learning is enabled through the mechanism of interaction history – the traces of activity people create as they interact with the Web. Social feedback systems collect and display interaction history to allow information seekers to learn how to complete their tasks more successfully by observing how other people have behaved in similar situations.

The dissertation outlines the design of two social-feedback systems, and describes two studies that demonstrate the real world applicability and feasibility of the idea. The first system supports *global learning*, by allowing people to learn new search skills and techniques that improve information seeking success in many different tasks. The second system supports *local learning*, in which people learn how to accomplish specific tasks more effectively and more efficiently. Two further studies are conducted to explore potential real-world challenges to the successful deployment of social feedback systems, such as the privacy concerns associated with the collection and sharing of interaction history. These studies show that social feedback systems can be deployed successfully for supporting real world information seeking tasks.

Overall, this research shows that social feedback is a valuable new idea for the social use of information systems, an idea that allows people to learn from one another's experiences and improve their success in many common real-world tasks.

ACKNOWLEDGEMENTS

I would like to thank my two supervisors, Carl Gutwin and Gordon McCalla, for their unconditional support, and for sacrificing so much of their time in order to guide me for the last seven years. A graduate student is extremely fortunate to be able to have a talented, kind, and well-respected supervisor, but few are lucky enough to have two. I look forward to passing on what I have learnt from observing my two role models. I cannot thank both of you enough for all that you have given me.

I would also like to thank the members of my committee for their interest, guidance, and encouragement. Thank you to Regan Mandryk, Ian McQuillan, Julita Vasileva, and Alison Muri (cognate), and to my external examiner David McDonald.

The work in this dissertation has benefited from many great collaborations. I would like to thank Ryen White and Jaime Teevan who provided me with twelve great weeks in Redmond, and who directly contributed to my dissertation work by guiding my work on the Search Dashboard, which eventually grew into the idea of social feedback. Again, I would like to thank Regan Mandryk for providing me with fantastic guidance on several projects and acting as a third supervisor many times. Thanks to my collaborators during throughout my PhD: Christopher Brooks, David Flatla, Lennart Nacke, Tadeusz Stach, Andre Doucette, Andy Cockburn, Robert Xiao, Aaron Genest, David McDine, Nathaniel Osgood, Michael Muller, Jill Freyne, and Miguel Nacenta. I would also like to thank the members of the Interaction (HCI) and ARIES Labs for providing an amazing environment in which to work and learn.

I could not have made it through the last seven years without the amazing encouragement and unconditional support of my family. First, thank you to Christine for being a fantastic mother and partner, and making me feel confident that I could spend many days away without leaving our family wanting. Thank you to Mara and Mateo for giving me motivation and focus, and for reminding me that I will always have a lot to learn. Thank you to Barb and Tom for making the last several years so much easier and more enjoyable; I owe you so much. Thank you to my brother Adam for inspiring me and instilling me with a curiosity about our world.

Finally, thank you to my parents, Bob and Carolyn – my original teachers. Thank you for showing me that there is no challenge in life that cannot be overcome.

Dedication

For all of my family.

TABLE OF CONTENTS

page

PERMISSIONS ii

ABSTRACT iii

LIST OF FIGURES xi

LIST OF TABLES xvi

LIST OF VIDEO FIGURES xvii

PUBLICATIONS OF THE AUTHOR xviii

  Publications with Content from this Dissertation xviii

  Publications Related to this Dissertation xviii

  Other Publications (selected – during the dissertation) xix

CHAPTER 1 INTRODUCTION 1

  1.1 Research Problem 2

    1.1.1 Important Concepts in the Problem 3

  1.2 Motivation 4

  1.3 Solution: Social Feedback 5

    1.3.1 A Conceptual Framework of Social Feedback 6
    1.3.2 The Creation and Evaluation of a Social Feedback System for Global Learning 6
    1.3.3 The Creation and Evaluation of a Social Feedback System for Local Learning 6
    1.3.4 An Examination of the Real World Challenges to Social Feedback Systems 6

  1.4 Contributions 8

  1.5 Overview of the Dissertation 8

CHAPTER 2 BACKGROUND 10

  2.1 Information and Information Seeking 10

    2.1.1 Mental Models 10
    2.1.2 Information Needs and Information Seeking 13

  2.2 Web-Based Behaviour 17

    2.2.1 Web-Based Information-Seeking Tasks 18
    2.2.2 The Principle of Least Effort in Information Seeking 19

2.2.3 Characterizing and Improving Search Performance 19

2.3 Tools and Techniques for Supporting Information Seekers on the Web 20

2.3.1 Original Concepts of Hypertext..................20
2.3.2 Browser-Based Tools..................21
2.3.3 Information Visualization for Supporting Information Seeking 21
2.3.4 Feedback and Persuasion 22

2.4 Interaction History and its Use in Information Seeking 24

2.4.1 Interaction History and Implicit Feedback for Information Seeking 24
2.4.2 Interaction History for Implicit Feedback..................25
2.4.3 User Modelling and Personalization..................27
2.4.4 Social Interaction History 28

2.5 Relevant Concepts from Computer Supported Cooperative Work 29

2.5.1 Coupling..................31
2.5.2 Awareness to Support Collaboration..................32
2.5.3 Overview of Informal Collaboration 35

2.6 Collaboration and Social Support in Information Seeking..................38

2.6.1 Social Information Seeking and Social Search 39
2.6.2 Social Navigation 42
2.6.3 Collaborative Information Seeking..................53
2.6.4 Privacy Concerns in Awareness Systems on the Web..................58
2.6.5 Summary of Collaboration and Social Support in Information Seeking 60

CHAPTER 3 SOCIAL FEEDBACK: SOCIAL LEARNING FROM INTERACTION HISTORY 63

3.1 Social Learning Theory 64

3.1.1 The Requirements of Modelling (Learning from Observation) 64
3.1.2 The Identity of a Model 65
3.1.3 The Basic Forms of a Model..................66
3.1.4 The Requirements for Supporting Self-Regulation..................66
3.1.5 Other Theories Related to Social Learning..................67
3.1.6 Applying Social Learning to Information Seeking..................68

3.2 Interaction History to Describe Behaviour 68

3.2.1 Capturing and using interaction history to describe behaviour..................69
3.2.2 Challenges to Using Interaction History on the Web 70

3.3 Social Feedback: Designing for Social Learning from Interaction History 72

3.3.1 Design Questions for Social Feedback Systems 73
3.3.2 Contextualizing the Social Feedback Design Questions to Web-Based Information Seeking 83

3.4 Summary of Social Feedback 87

CHAPTER 4 GLOBAL SOCIAL FEEDBACK: THE SEARCH DASHBOARD 88

4.1 The Search Dashboard: Social Feedback for Global Learning 89

4.1.1 What tasks should be supported? 90
4.1.2 What information do people need in these tasks? 91
4.1.3 How can the information be displayed?..95

4.2 Study...98

4.2.1 What data was gathered?...99
4.2.2 Conditions 99
4.2.3 Participants...99
4.2.4 Procedure..100
4.2.5 Data Analysis 101

4.3 Results...101

4.3.1 The Search Dashboard is Engaging 102
4.3.2 Techniques & Tendencies More Insightful Than Topics 102
4.3.3 Self-Assessment of Search Skills Changed..103
4.3.4 Comparison Leads to Increased Engagement and Insights.......................................104
4.3.5 Behaviour Change Observed for Tendencies & Techniques 105

4.4 Summary of Results 106

CHAPTER 5 LOCAL SOCIAL FEEDBACK: WEBWEAR 108

5.1 WebWear: Designing Social Feedback for Local Learning 109

5.1.1 What tasks should be supported? 109
5.1.2 What information do people need in these tasks? 110
5.1.3 How can the information be displayed?...110
5.1.4 WebWear...111

5.2 Empirical User Study of WebWear..115

5.2.1 Interface conditions 116
5.2.2 Tasks and scenarios 116
5.2.3 Participants and Apparatus..118
5.2.4 Design and Procedure 118
5.2.5 Data Analysis 119

5.3 Results...119

5.3.1 Were participants more effective with WebWear?..119
5.3.2 Were participants more efficient when using WebWear? 120
5.3.3 Perceived Usefulness of WebWear and BroadNav 122

5.4 Summary of Results 122

CHAPTER 6 REAL WORLD CHALLENGES FOR SOCIAL FEEDBACK 124

6.1 Field Trial of the WebWear System 125
6.1.1 Participants and Procedure 126
6.1.2 Results 127
6.1.3 Summary of Results for the Field Trial 133
6.2 Diary Study 134
6.2.1 Participants 136
6.2.2 Procedure 136
6.2.3 Data 138
6.2.4 Data Analysis 139
6.2.5 Results 140
6.2.6 Summary of Results for the Diary Study 147
6.3 Summary of Results 148
CHAPTER 7 DISCUSSION 150
7.1 Summary of Findings 150
7.2 Global Learning and the Search Dashboard 151
7.2.1 Search Dashboard: Explanations for Main Results 151
7.2.2 Search Dashboard: Generalizing the Results, and Deployment Issues 155
7.3 Local Learning and WebWear 156
7.3.1 WebWear: Explanations for Main Results 157
7.3.2 WebWear: Generalizing the Results and Deployment Issues 160
7.4 Real World Challenges for Social Feedback 161
7.4.2 Real World Challenges: Generalizing the Results 166
7.5 Overall Discussion 167
7.5.1 The Importance of Appropriate Models: Tightly-Knit Groups and Role Models 167
7.5.2 What are the differences between the approaches to supporting global and local learning? 168
7.5.3 The Importance of Loosely-Coupled Information Seeking 170
7.5.4 Information Seeking is Essential 171
7.6 Scope and Limitations 171
7.7 Reflection on Broader Motivations 172
CHAPTER 8 CONCLUSION 174
8.1 Contributions 175
8.2 Future Research 176
8.2.1 Social Feedback System Features 176

8.2.2 Questions for Evaluation 179
8.2.3 The Application of Social Feedback Outside of Information Seeking on the Web ... 181

REFERENCES 184

APPENDIX A: STUDY MATERIALS 202

University Ethical Approval for the Dissertation Research 202

WebWear – Ethical Consent Form 203

WebWear – Textual Excerpt Consent Form 204

WebWear – Study Tasks 205

WebWear – Study Handouts (Used to explain the task scenarios and study systems) 206

WebWear – Post Task Questionnaire 211

WebWear – Post System Questionnaire 215

WebWear – Post Experiment Questionnaire 217

Diary Study – Informed Consent Form (webpage for participant records, was also available in the study system) 221

Diary Study – Survey 225

Appendix B: Implementation Details 229

The Search Dashboard 229

WebWear 229

SaskWatch Diary System 229

LIST OF FIGURES

Figure page

Figure 2.1. Spence's framework of navigation. Rectangles represent mental artifacts, and ovals are actions taken by the navigator. Recreated from [139]..12

Figure 2.2. Pirolli and Card identified two loops that take place during information seeking. A foraging loop, where information is extracted and stored provisionally, and sensemaking loop where information is extracted and used. Recreated from [54]....14

Figure 2.3. The author’s search activity displayed in Google Trends. 23

Figure 2.4. The distinction drawn between interaction history (a), which is limited to human-object interaction, and Indratmo and Vassileva’s social interaction history (b) the direct user-user communication that takes place around an object. Recreated from [75]..29

Figure 2.5. The CSCW matrix, positioning groupware technologies in a time versus place grid. Public domain image from [164]...30

Figure 2.6. Kraut et al.’s organization of workplace interactions. Informal interactions are all interactions that are not scheduled. Including planned and unplanned interactions. Unplanned interactions can be either opportunistic or spontaneous.........................36

Figure 2.7. Evans and Chi’s canonical social information seeking model (reproduced from [54] with permission). 41

Figure 2.8. Hill and colleagues original scrollbar designs showing *wear* patterns of reading and writing. Reproduced from [70] with permission. 44

Figure 2.9. Social Navigation in different online systems. (A) Amazon.com (i) uses interaction history to provide recommendations, (ii) displays aggregations of user created ratings, and (iii) allows users to correct inferred recommendations with explicit feedback. (B) Twitter users often use the surface for direct social navigation, by sharing links they find interesting..45

Figure 2.10. The Footprints System displaying its tools: the browser (on the right) is displaying the Media Lab website; the map tool shows all pages and how they are connected (top left); the path tool shows different paths starting from a page with information describing relative use (lower left). Reproduced with permission,from [156]. 49

Figure 2.11. Alternative designs for scented widgets. From [167], used with permission. © 2007 IEEE 49

Figure 2.12. Different visual mapping possibilities for scented widgets. From [167], used with permission. © 2007 IEEE...50

Figure 2.13. A global list of accesses in Diebergger's CoWeb system. The visual cues used were small icons, which used color to describe the user activity on the linked page. "New" icons represented the age of the page (red = less than a day old, yellow = less than 3 days old, or blue = less than a week old). Footprint icons represent the relative amount of visitation to a page (red = lots of visitation, yellow = moderate visitation, blue = little visitation). From [46]. 52

Figure 2.14. Indicators of interaction history from two of the author's friends on Facebook. Friend's names have been blurred. 52

Figure 2.15 A search on Google for CHI 2012 on Google, shows a social navigation indicator. Because a friend, Michael, is also one of the author's contacts on Google+, Google augmented the search result snippet with an indicator that says Michael shared a workshop, an implicit recommendation to follow the link. 53

Figure 2.16. Fishlar-DT allows collocated searchers to find and organize videos in a collection. Reproduced with permission, from [136]. 55

Figure 2.17. The search interface for the S^3 System. Search dialog is at the top where users can issue queries. The lower left-hand frame displays a web browser, and the right-hand sidebar displays the summary of the search session. Reproduced with permission, from [106]..56

Figure 4.1. The Search Dashboard, displaying the Tendencies section (referred to as Performance in the study). 90

Figure 4.2. The Domains data in the Topics section...96

Figure 4.3. Alternative presentations of numeric data: without comparison data (left), and with comparison data (right). 97

Figure 4.4. An example of the Traffic direct answer. When a user hovers over an answer type an example is shown. 98

Figure 4.5. Mean ratings and ±SEM of insights provided by the Dashboard for each of the three Data Types. All questions were sig. diff. (*p*<.001). Sig. diff. pairs are indicated by lines. 102

Figure 4.6. Mean and ±SEM for self-assessment ratings before (pre) and after (post) using the Search Dashboard. Sig. diff. were found for all ratings (*=*p*<.01, **=*p*<.001). 103

Figure 4.7. Mean and ± SEM for ratings of insights between Dashboard variants. Sig. diff. were found for all aspects (*=*p*<.05, **=*p*<.001). 104

Figure 4.8. Mean and ±SEM usage statistics for both Dashboard variants. For all aspects, usage was significantly higher in the compare condition than the no compare condition (*p*<.01). 104

Figure 5.1. WebWear being displayed on a search results page for “Data Mining”. 1) Glyph presenting visitation to the current page (purple) and website (green). 2) Activity glyphs on a particular link. 3) Details of activity that appears from hovering over a glyph. 4) Overview, providing the location of glyphs on the page. 112

Figure 5.2. (left) The visual glyphs displayed beside a link indicating the number of unique visitors to the specific webpage that the link leads to. (right) The glyphs displayed beside a link indicating the timeframe for the most recent visit to the website the link leads to. 112

Figure 5.3. The details popup for activity on the wikipedia.org website. Shown are details of all contacts’ visits to the Wikipedia.org website. 113

Figure 5.4. The WebWear Controls Menu, which provides access to filtering the traces presented by user (currently the WebWear will present traces for John and Mary, but not Lynn). WebWear users can view and selectively delete any history collected by the system. Users can also selectively change whether the system is collecting or displaying traces. 114

Figure 5.5. Group-based query completion displayed on Bing; a recency-based based on the queries of contacts...115

Figure 5.6. Glyphs as displayed in the BroadNav condition of the study. Mousing over the glyphs in the BroadNav system did not provide any additional information.....................116

Figure 5.7. (Left) Task success rate. (Right) Subjective certainty about task success. 120

Figure 5.8. (Left) Average task completion time. (Right) Participant ratings of the system being helpful and distracting..121

Figure 5.9. (Left) Participant agreement with having completed the task efficiently. (Right) Ratings of needing to communicate to be successful in the task............................122

Figure 6.1. A request to complete a survey made by the diary study system (indicated by the red arrow)...137

Figure 6.2. A survey from the diary system embedded in the current webpage being visited by a participant...138

Figure 6.3. The mean agreement with being concerned if interaction history was shared with colleagues for the page currently being visited (by participant). The higher the value the more concerned a participant is on average, the lower the value, the less concerned...141

Figure 6.4. The proportion of work tasks to non-work tasks by participant. 142

Figure 6.5. The mean agreement with being concerned if interaction history was shared with colleagues for the page currently being visited (by participant), split by whether it was a work or non-work task...143

Figure 6.6 The mean agreement of being concerned with interaction history by task type........144

Figure 6.7. The number of occurrences of each task type reported in the diary study in work and non-work situations...145

Figure 6.8. The total number of information-seeking tasks that involved some form of collaboration, split by work and non-work situations...146

Figure 6.9 The mean ratings of agreement with being concerned if colleagues could view interaction history, split by non-collaborative and collaborative tasks....................146

Figure 6.10. Participant ratings of agreement with specific concerns for sharing interaction history with colleagues..147

LIST OF TABLES

Table page

Table 2.1. The summary of goal categorization for information-seeking tasks, provided by Kellar et al. [81] 18

Table 2.2. A categorization of user behaviour by minimum scope that the behaviour acts upon that can be used for implicit feedback. Originally proposed by Oard and Kim [114], and extended by Kelly and Teevan [82] (additions are highlighted to clarify provenance). Table reproduced from [82] 26

Table 2.3. Elements of workspace awareness that relate to the past. Recreated from [67]. 33

Table 2.4. Specific questions relating to 'who' based on an artifact-based, person-based, or workspace-based view of change awareness. Recreated from [145] 34

Table 2.5. Goecks' characterization of different sources of community data used in social navigation. Recreated from [63]. 46

Table 3.1. The high-level design questions and important component questions that should be considered in designing social feedback systems. 73

Table 3.2 The specific details of interaction history that can support the important concepts of social learning theory, and a description of how they can relate with example scenarios. 77

Table 4.1. The 11 metrics for the 12 data elements presented pre and post visiting the Dashboard. Significant differences were found *within subject* and *between groups* for the same variables (bolded). For these variables comparison led to significantly higher usage. 106

LIST OF VIDEO FIGURES

Video page

The Search Dashboard Video Figure, available through the University of Saskatchewan Library website or at:
http://delivery.acm.org/10.1145/2210000/2208311/paperfile236-3.mp4 or
http://scottbateman.ca/permanent/258-search-dashboard-video-figure.mp4 89

The WebWear Video Figure, available at (address to be updated, work under review):
http://scottbateman.ca/permanent/webwear-video-figure.m4v 111

PUBLICATIONS OF THE AUTHOR

Publications with Content from this Dissertation

Bateman, S., Gutwin, C., and McCalla, G. WebWear: Social Navigation for Loosely-Coupled Information Seeking in Tightly-Knit Groups. In Submission to the *ACM 2013 Conference on Computer Supported Cooperative Work (CSCW '13)*. 10 pages. In submission. (In Chapters 2, 3, 5, and 7).

Bateman, S., Teevan, J, and White, R. The Search Dashboard: How Reflection And Comparison Impact Search Behavior. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '12),* 2012, 1785-1794. (In Chapters 2, 4, and 7).

Bateman, S., Gutwin, C., McCalla, G., White, R. Collective Information Seeking: Supporting Search with Low-Cost Activity Sharing and Collaboration. In the *Proceedings of the Workshop on Human-Computer Interaction and Information Retrieval (HCIR '11)*. 4 pages. (In Chapters 2 and 3).

Bateman, S., Gutwin, C., Osgood, N., and McCalla, G. Interactive usability instrumentation. In *Proceedings of the Symposium On Engineering Interactive Computing Systems (EICS '09).* 45-54. (Chapter 3).

Bateman, S. Using Group Interaction History in the Wild. In the *Extended Proceedings of the of the 2010 ACM Conference on Computer Supported Cooperative Work – Doctoral Colloquium (CSCW 2010)*, 523 - 524. (In Chapter 2 and 3).

Publications Related to this Dissertation

Bateman, S., Muller, M.J., and Freyne, J. Personalized retrieval in social bookmarking. In *Proceedings of the International Conference On Supporting Group Work (GROUP '09)*, 91-94.

Other Publications (selected – during the dissertation)

Bateman, S., Mandryk, R.L., Gutwin, C., and Xiao, R. Analysis and Comparison of Target Assistance Techniques for Relative Ray-Cast Pointing. In submission to the *International Journal of Human-Computer Studies (IJHCS).* Submitted (March 2012).

Flatla, D., Gutwin, C., Nacke, L., Bateman, S., and Mandryk, R. Calibration Games: Making Calibration Tasks Enjoyable by Adding Motivating Game Elements. *Proceedings of ACM Conference on User Interface Software Technology (UIST '11)*, Santa Barbara, California. October 16-19, 2011. 403-412.

Bateman, S., Stach, T., Mandryk, R.L., Gutwin, C. Using Target Assistance for Subtly Balancing Competitive Play. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '11)*, 2355-2364.

Bateman, S., Doucette, A., Xiao, R., Gutwin, C., Mandryk, R.L., and Cockburn, A. Effects of View, Input Device, and Track Width on Video Game Driving. In the *Proceedings of Graphics Interface (GI '11)*, 207-14.

Bateman, S., Mandryk, R.L., Gutwin, C., Genest, A., McDine, D., and Brooks, C. Useful Junk? The Effects of Visual Embellishments on the Comprehension and Memorability of Charts. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '10)*, 2573-2582.

CHAPTER 1
INTRODUCTION

> *... the environment people are embedded in is, among other things, a reservoir of resources for learning, problem solving, and reasoning. Culture is a process that accumulates partial solutions to frequently encountered problems. Without this residue of previous activity, we would all have to find solutions from scratch. We could not build on the success of others.*
>
> *- Hollan, J., et al., 2000 ([71], p. 178).*

People look for information on the World Wide Web billions of time per day. Even though information seeking on the Web has become a central part of our daily lives, there are many times when searches fail and information needs go unmet. One of the problems is that web tools provide few opportunities to learn from one another's experiences. For example: We struggle to book flights at the best price and schedule – which airline did my colleagues book, again? We try to find a webpage that our friend says has the answer to our question – where did he say I could find it? We get frustrated trying to find the keywords that will get us the website we know exists – why can my officemate always seem to find things so easily!? We get overwhelmed by the number of options available when shopping for a new laptop – which ones would an expert consider? We try to direct our parents to a family video on Facebook through email – how can I tell them how I got there?

These scenarios all share a common characteristic; like so much that we do in our day-to-day lives, they are embedded in a social context, where we can use the knowledge we get from other people to guide our actions. When we need help or would like guidance, we seek out others to provide it, *if* they are available to give it. To further illustrate how such tasks are embedded in a social context, I will describe a specific example in more detail.

> *An academic named John is travelling to Vancouver for the first time, and he wants to find a reasonably priced hotel in a safe neighbourhood. He knows that several of his colleagues have travelled to Vancouver and had good*

> *experiences with their accommodations. However, John has left booking his hotel to the last minute; it's late in the evening and his colleagues are not around to advise him. John would really like to get the hotel booked because he is travelling in the morning, so he decides to book the first hotel he finds that seems good enough.*

This example and the examples above are all *loosely-coupled information-seeking task*s. In loosely-coupled information seeking people work alone to accomplish their tasks, but when the need for help arises, they use or look for information from other people to learn more about how they could proceed. However, the people they would like to guide them are not always available. In the above example, John's contacts are not available, and he has decided the overhead of emailing, or otherwise tracking down and bothering his colleagues is too high. This demonstrates an important characteristic of loosely-coupled information-seeking tasks: they are often *discretionary*. This means that looking for guidance is often at the discretion of the seeker, so it does not always need to occur, because the tasks can be completed individually (even if the result is not optimal), or can be postponed or abandoned. If a person seeking information online perceives the costs of looking for guidance as being too high, then they will simply opt not to look for it.

When help or guidance is not provided in loosely-coupled information-seeking tasks there are missed opportunities, where the experiences of others could have been useful. These missed opportunities can lead work to be less efficient (taking more time and effort) and less effective (finding incorrect or less than ideal information). In some cases, this causes annoyance and inconvenience (e.g., "I can't find the article Jim was telling me about, I will have to ask him about it tomorrow"); in other cases it can lead to more severe problems and task failures (e.g., "I can't find the online form Sally told me to complete before tonight's deadline, and her flight just left").

1.1 Research Problem

Despite the importance of information seeking in daily lives, we lack basic knowledge of the social contexts in which information seeking takes place. This lack of knowledge means that current information seeking tools do not support common loosely-coupled information seeking situations sufficiently, and this leads many information seekers to fail or have less success in completing their tasks.

1.1.1 Important Concepts in the Problem

An *information need* is a gap in a person's knowledge [44]. People engage in *information seeking* to fill the gap [44]. On the World Wide Web, information seeking can involve several activities, including performing a keyword query on a search engine, reviewing search engine result lists, reading the content of webpages, considering which links to follow, and creating a record of what has been found [104]. A person who is working through an information-seeking task is referred to as an *information seeker* (or simply a *seeker).*

Interaction history is the traces of activities that people create as they conduct their work with a system's interface. On the Web, people create interaction history as they use search engines, click links, and read webpages.

Coupling describes the amount of work that people can do individually before they require communication with or actions from another person [67]. Coupling is described as being on a spectrum between *loose* coupling – where collaborators work independently – to *tight* -coupling – where work is highly dependent on others. I define *loosely-coupled information seeking* (*LCIS)*, then, as: any information-seeking task where the seeker is working primarily individually, but could learn how to better accomplish some aspect of the task from some other person (the *provider*). The provider is willing to share the information needed by the seeker, but transferring the information to the seeker may or may not be practical or possible.

Support and *guidance* are information provided by some tool or other person in an attempt to help resolve a seeker's information need with the most success. It should also be noted that the definition means that an LCIS task can be part of collaborative work or completed individually. In LCIS the potential for guidance exists, but the knowledge or experience of some other person or people (the *guides*) must be transferred to the seeker somehow.

Learning can be defined broadly as "… acquiring new or modifying existing knowledge, behaviours, skills, values, or preferences [165]." I will consider an information seeker to have learned from someone else, if the seeker makes use of presented information and changes their behaviour to reflect it. In this sense, I will describe two forms of learning: *global learning* and *local learning.* Global learning refers to learning new skills or acquiring new behaviour that can be applied in many different information seeking situations. Local learning refers to learning that occurs within a specific context or situation, and can lead to the specific task instances being accomplished more effectively or efficiently.

The *effectiveness* of an information-seeking task is the degree to which the information found in completing the task is accurate, appropriate for the task, and meets the information seeker's goal. *Efficiency* is how much time and how much effort the seeker needs to complete the task. Finally, being *successful* in an information-seeking task is a combination of having completed the task both effectively and efficiently.

1.2 Motivation

Along with using email, information seeking is the most common task people do online, with 91% of online US adults using search engines to find information on the Web, and 54% using them every day [125]. A recent estimate puts the collective time spent by people searching and making sense of online information at over 70 billion hours per year in the United States alone [32]. Even though information seeking is a central online activity, there are few traces of that activity: once people have completed their tasks and filled a gap in their knowledge, their work most often disappears, and benefits no one else [57].

This disappearance would not be a problem if all information seeking was successful – but difficulty finding information on the Web is common: for example, 71% of search engine users report that they do not always find what they are looking for [125]. A recent survey found that 5% of searches completely fail, and that when failures occur they take twice as long as successful searches [54]. This means that roughly 6.4 billion hours per year are wasted on failed tasks, and this does not consider the time lost when tasks could have been completed more efficiently.

Because information seeking is such a common and frequent activity, any improvement in the information seeking process can have a profound impact on productivity. However, the basic tools used to accomplish search tasks have evolved slowly [104]. The main approach for improving information seeking has been through improving search engines to provide the most accurate and relevant results for a given query. While search engines are an important aspect of information seeking, they represent only one part of the typical process, which also includes reviewing webpages and selecting links to follow. Further, while search engines have substantially improved since they started appearing on the Web, search failures and difficulties are still a common occurrence.

1.3 Solution: Social Feedback

The problem is that there is a lack of knowledge about the social contexts in which people engage in information-seeking tasks. As a result, current approaches for supporting seekers do not adequately address the needs of loosely-coupled information seeking. I address the problem by designing and building new support techniques based on the new idea of *social feedback.* Social feedback allows people to learn from the experiences of others by viewing representations of past information-seeking activities. Social feedback systems allow people to learn in two ways: *global learning*, in which people learn new search skills and techniques; or *local learning*, in which people learn how to accomplish specific tasks more effectively and more efficiently.

In this dissertation, I show that information seekers can learn how to improve their information seeking skills and accomplish information-seeking task more successfully by using tools that provide social feedback – representations of other people's past information seeking behaviour. The solution of social feedback has four main steps:

1. Create a conceptual framework that describes how social feedback systems can support information seekers in learning how to complete their information-seeking tasks more successfully.
2. Create and evaluate a system based on social feedback to provide opportunities for global learning.
3. Create and evaluate a system based on social feedback to provide opportunities for local learning.
4. Examine the two potential challenges to widely deploying social feedback systems in the real world. *1) The Overlap Challenge:* Is there sufficient overlap between information seekers to provide opportunities for social feedback? *2) The Privacy Challenge:* Are people willing to share their past information-seeking activities (called *interaction history*) to enable social feedback in at least some situations?

Below, these steps are described in greater detail.

1.3.1 A Conceptual Framework of Social Feedback

In order to focus the design space for social feedback systems and make the most important requirements clear, a conceptual framework must be created that provides:

1. A description of the important elements of social learning theory.
2. A description of the main approach and concerns in collecting and applying interaction history to web-based information-seeking tasks.
3. A list of design questions that consolidate the elements of social learning theory with the information transfer mechanism of interaction history.

1.3.2 The Creation and Evaluation of a Social Feedback System for Global Learning

To show that social feedback can enable global learning, I designed, built, and evaluated the Search Dashboard, a system to help information seekers learn search skills and techniques from expert searchers. Based on the conceptual framework of social feedback, the system allows individual search engine users to compare descriptions of their search engine usage with search experts. Through a comparative study of the Search Dashboard with an alternative state-of-the-art system, an evaluation showed that the Search Dashboard allowed searchers to learn and adopt new search behaviour that was in line with that of the search experts.

1.3.3 The Creation and Evaluation of a Social Feedback System for Local Learning

To show that social feedback can enable local learning, I designed, built, and evaluated WebWear, a system that helps seekers to complete typical LCIS tasks more successfully. Based on the conceptual framework of social feedback, WebWear allows seekers to encounter and make use of the behaviour traces of other people during their information-seeking tasks. Through a comparative study with two alternative systems (a typical web browser and a broad social navigation system), I show that WebWear allows seekers to complete LCIS tasks more efficiently and more effectively.

1.3.4 An Examination of the Real World Challenges to Social Feedback Systems

A long-term and broad deployment of social feedback systems in the real world is not feasible within the scope of this dissertation. Therefore, I address the two main challenges through two separate studies. The separate studies allow examine both of the two challenges in different ways.

First, the *overlap challenge* might occur in social feedback systems because information seekers may visit the same webpage too infrequently, which means social feedback would not ever be available when it was needed. The overlap challenge might occur due to the breadth and depth of the Web, because an information seeker could potentially spend most of their time visiting unique webpages that no one else has visited. This is of particular concern to local learning systems, such as WebWear, that provide support based on encountering traces during an information-seeking task.

Second, the privacy challenge might occur in social feedback systems because people's web activity encompasses many personal and potentially private details. Social feedback systems rely on people willingly sharing their interaction history for others to learn from. Yet many seekers may be unwilling to share their interaction history because of concerns of what others may think of their Web activity. However, in certain situations people are very willing to share their activities, and actively do so (e.g., social networking sites such as Twitter or foursquare). A better understanding is needed about the contextual details of information-seeking tasks and when people might be willing to share their interaction history.

To provide further insight into the two challenges that social feedback systems face, I examined the challenges in two separate ways: a field trial of WebWear, and a diary study. The two studies allowed both of the challenges to be studied in different ways.

First, I conducted a week-long field trial of the WebWear system within a real workgroup. The field trial showed that the system design effectively provided many opportunities to encounter the traces of colleagues. Further, the field trial also showed that traces actually provided support for several LCIS tasks. Further, while participants did not report any bad experiences that arose out of privacy concerns or having their interaction history shared with colleagues, several participants expressed some concerns about their privacy.

To provide an initial understanding, I conducted a two-week diary study that allowed participants to periodically describe the contextual details of their current information-seeking tasks and their privacy concerns. The diary study shows that people are much more willing to share their interaction history when they are actively engaged in work-related information seeking, and that people are most frequently have few concerns about sharing their interaction history with their colleagues. The concerns that were expressed occurred mostly in non-working tasks, and were largely reported by the same few individuals.

Together the two studies provide further evidence that social feedback is an appropriate solution for supporting real world information-seeking tasks.

1.4 Contributions

The main contribution of this dissertation is the idea of social feedback, which provides new knowledge about how loosely-coupled information-seeking tasks can be supported. Further, it provides five minor contributions that provide guidelines and examples for the application of the social feedback idea, and strong evidence that systems based on social feedback provide workable solutions for supporting loosely-coupled information seekers. The five minor contributions include:

1. A conceptual framework for the design of social feedback system that allows information seekers to learn how to improve their success in information-seeking tasks.
2. A social feedback system for global learning and empirical evidence that social feedback systems can enable the acquisition of new behaviour in real information seeking practices.
3. A social feedback system for local learning and empirical evidence that information seekers can learn how to accomplish collaborative LCIS tasks with increased success.
4. An understanding that the overlap challenge would not be a major concern for the long-term deployment of social feedback systems in the real world.
5. An understanding that the privacy challenge can be contended with through the use of social feedback systems in the appropriate contexts, and that the workgroup context provides an appropriate context for the use of identifiable interaction history.

1.5 Overview of the Dissertation

Chapter 2 introduces the problem domain of the dissertation – information seeking – and focusses on theories of how people address their information needs. I give important background for the solution of social feedback, by describing work in computer-supported cooperative work and the use of interaction history. I also describe two current approaches that allow people to collaborate in information seeking, and explain why they do not provide adequate support for loosely-coupled information seeking.

In Chapter 3, I describe the idea of social feedback in depth. First, I describe social learning theory, which is the basic theory of how people can learn through the observation of others. Next, I describe how the behaviour of others can be instantiated through interaction history. Finally, I describe how they can be combined to create systems that provide social feedback.

In Chapters 4 and 5, I describe the design, development, and evaluation of two social feedback systems. Chapter 4 focuses on the Search Dashboard, a social feedback system for the global learning of search engine skills. The Search Dashboard allows people to reflect on their own behaviour and compare it with expert searchers. In Chapter 5, I describe a system for local learning, the WebWear system, which allows people to complete common types of collaborative tasks more efficiently and more effectively, through encountering and making use of the traces of their colleagues' past information-seeking history.

In Chapter 6, through a field trial of the WebWear system and a separate diary study, I examine two of the most important challenges to the widespread deployment of social feedback systems in the real world: the privacy challenge and the overlap challenge.

Chapter 7 discusses the findings presented in the three preceding chapters. Chapter 8 concludes the dissertation by summarizing the findings and the contributions, and providing directions for future research.

CHAPTER 2
BACKGROUND

This chapter provides a basic understanding in four general areas: i) how people engage and behave in information seeking on the Web, ii) how information seekers are supported using tools, and how the tools are improved through using interaction history, iii) the relevant concepts about how people collaborate and work together in information seeking, and iv) the proposed solutions for supporting different forms of collaboration in information seeking.

2.1 Information and Information Seeking

Marchionini defined *information* as “anything that can change a person’s knowledge ([95], p. 5).” Belkin highlighted that information defined in this way suggests that sources can, therefore, include: objects, reflecting on memories, and the thoughts and ideas of other people [15]. So, the information in a book, or any other object, can transfer knowledge to a person’s cognitive system, updating their knowledge. Knowledge can also be changed into physical representations – such as writing on paper or sounds from speaking – that relay abstractions of knowledge – such as words or numbers ([95], p. 5). Buckland defined different forms of information: information-as-process (or the act of communicating), information-as-knowledge (a change, up or down, in the level of uncertainty), and information-as-thing (the physical objects that may convey knowledge) [27]. As people encounter and acquire new information they have mental processes for its selection, comparison, and storage [95].

2.1.1 Mental Models

Marchionini points out that there exist many theories to help explain how human mental activity works; e.g., theory of short and long-term memory, semantic networks, frames, and mental models ([95], p.12). Marchionini [95,94], and others [44,112,121,129,139,150] have focused on the concept of mental models, as it seems to appropriately identify how information is used and how new information fits into a gap – a missing part of the model.

Mental models are built using information ([15] as cited by [95]), and are ever-changing representations of the real world that are created and stored in the mind [112]. People create mental models to predict the results of actions they are considering. For example, a mental model of a book allows the reader to create a strategy of how one should begin and proceed in reading, and to consider how much work it will take to read a chapter. Mental models are also created for different domains (topics). So, as new information is encountered it can be integrated into an existing model (the process of understanding) [95].

2.1.1.1 Acting in the World

The basis of mental models is the view that humans require information to plan and carryout actions – whether mental or physical [112]. To do this effectively people need to have "plausible mental models (or understandings)" of the world around them ([95], p.28). People actively pursue and acquire new information, so they can best interact within different contexts (i.e., environments, situations, people, constraints, etc.) [95].

Several approaches for describing how people act in the world have been proposed, for example Furnas' View-Traversal [62] and Norman's Action Cycle [113]. In his action cycle [113, p.51], Norman describes four basic steps: perception of the world (which includes building a model), evaluation (interpretation of the internal model), goals (formulating a strategy), and execution.

2.1.1.2 (Mental) Modelling

The process of *mental modelling* constructs new information into an internal mental model [139]. Carmel et al. also described it as the process whereby registered content is integrated into existing models ([31] as cited by [139]). For example, an electric guitarist might model the resulting sound characteristics that are created by varying the input levels of two sound effects (e.g. reverb and delay). The guitarist might then know what types of sound results given different levels of each sound effect. Many times, and as is the case for the example, the act of modelling occurs at the same time as exploring the space of possibilities (also referred to as *browsing*)[139]. New internal models of information often do not fit into other existing mental models. Tversky referred to the situation of several misaligned models as *mental collages.* The gaps between internal models may lead to errors in judgment and memory [150].

Spence described navigation as "the creation and interpretation of an internal mental model ([139], p.920)"[1]. Spence's proposed a framework of navigation (see Figure 2.1) is primarily inspired by how the Web is used, but is also put forward as a general framework for

---

[1] Note: navigation is often used as metaphor for movement through or to websites, and not necessarily associated with mental models as suggested by Spence [139].

describing mental modelling in information systems. It is comprised of four basic parts: browsing, modelling, interpreting, and formulating a browsing strategy. Therefore, it describes both modelling and the actions undertaken to update the model. In the navigation cycle, seekers interpret their mental model in combination with what can be seen in a system interface (e.g., a webpage). From this seekers formulate a browsing strategy – how they should proceed to collect new information. Seekers then browse[2] the interface to gather information about what they see. This is followed by modelling what has been seen and incorporating it into the mental model [139]. This is very similar to Norman's action cycle [113] and Spence indicates that the acts of building a model and gathering information to update the model can occur at the same time [139].

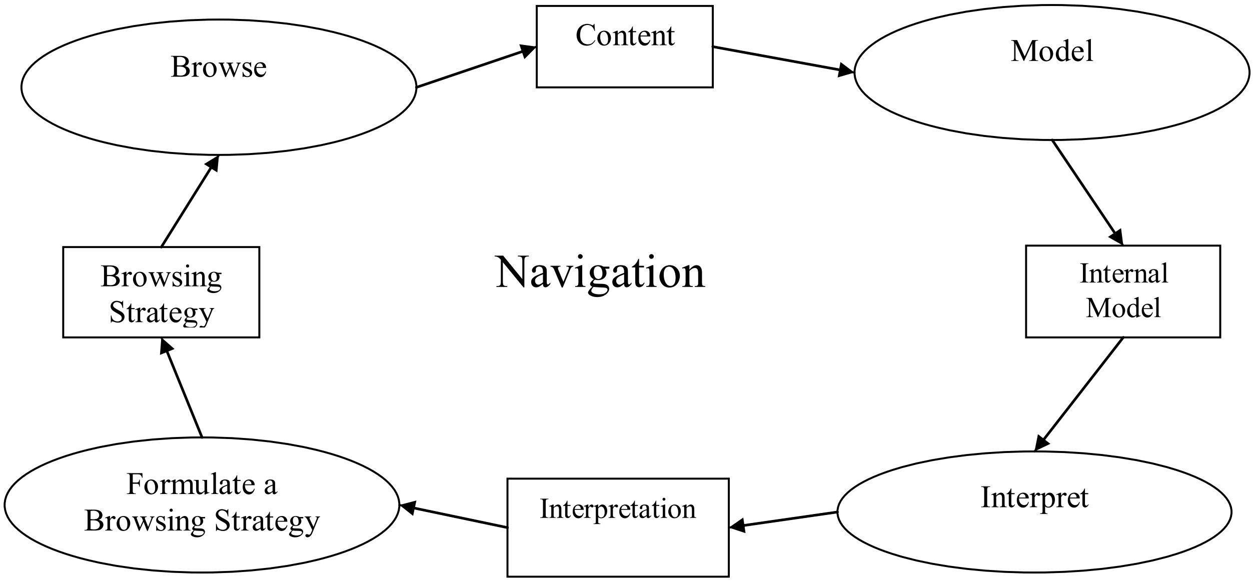

Figure 2.1. Spence's framework of navigation. Rectangles represent mental artifacts, and ovals are actions taken by the navigator. Recreated from [139]

2.1.1.3 Characteristics of Mental Models

Norman further characterized mental models, from his many observations, in a few general statements [112]:

1. Mental models are not complete representations.
2. People have an extremely limited ability to work their way through their models.

---

[2] Note: browsing is also put forward as a specific strategy used in information seeking as described below.

3. Mental models are always changing; details are forgotten, especially over time.
4. Mental models have weak boundaries; models of similar systems, devices and domains overlap and are confused with one another.
5. “Mental models are ‘unscientific’”: people keep unfounded beliefs even when they understand their problem, since it saves physical and/or mental effort.
6. People prefer physical action to mental manipulations of their model. For example, a simple action – like rebooting a computer or cell phone to fix it, rather than trying to diagnose an error – that applies to many objects is preferred, as it reduces the chance for confusion.

2.1.2 Information Needs and Information Seeking

When a mental model is incomplete or deficient in some way the holder of the model is said to have an *information need*; referred to as a “gap” by Dervin [44] or an “anomalous state of knowledge” by Belkin [15]. *Information seeking,* then, is “the search for information… a process in which humans purposefully engage in order to change their state of knowledge ([95], p.5)”, in order to fill the gap [44] or to clarify the anomalous state [9].

Marchionini points out that information seeking has been given much attention in the literature, particularly in information science and communications. He focuses on two general “Human-centered models of information seeking ([95], p.5)”, in the research of Belkin [9] and Dervin [44][3].

Belkin led research into how information seekers deal with anomalous states of knowledge. This theory posits that when information seekers are confronted with a problem, the problem itself and how they should proceed to solve it are not clear. So, the information seeker must go through a series of clarification steps, so that the problem can be communicated and a search action may be made [15]. This implies that search systems should provide some support for an iterative search process that has many steps. However, the model is limited in that it focuses on searches that are open-ended and ill-defined (an *exploratory search task* as described by [159]), but does not focus on other cases such as looking up the answer to a specific question (called *fact-finding*) or long-term research by a domain expert [95].

---

[3] Dervin’s work is treated more at length in Sensemaking.

Pirolli and Card identified two loops that take place during information seeking (as shown in Figure 2.2). The first loop is a foraging loop (described in the next Section) where information is identified, filtered, and extracted into a provisional schema. The second loop, is a sensemaking loop where a mental model is iterated over and updated based on information from the provisional model [123].

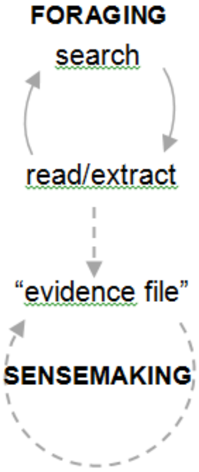

Figure 2.2. Pirolli and Card identified two loops that take place during information seeking. A foraging loop, where information is extracted and stored provisionally, and sensemaking loop where information is extracted and used. Recreated from [54].

2.1.2.1 Information Foraging

Information foraging theory provides a set of models to describe and predict information seeking behaviour based on optimal foraging theory – how living organisms behave in their search for food. The theory is based on the notion that humans are “informavores” ([42] as cited by [122]) – similar to animals in the wild, we seek out and consume information. Pirolli and Card use several models in their overarching theory:

- *the patch model*: deals with how time should be dedicated and optimized in foraging;
- *information scent models*: address how valuable information can be identified given some cues in the environment;
- *diet model*: deals with which information items should be foraged for and consumed [122].

An important underlying assumption of the theory is that people will adapt their strategies or change environmental factors such that the rate of gain of valuable information is maximized [121,122].

The Patch Model

The patch model relates to where information can be found in an environment; typically information (like food) is found in clusters. This can relate to many real world situations such as a stack of documents, a bookshelf, or an email inbox. It focuses on how time should be allocated in the pursuit of particular information items and information clusters. Foragers acquire knowledge rapidly when it is encountered and spend time moving between patches of information. A forager must make estimates about the cost of searching for a new patch and whether to stay or leave when new information starts to be exhausted [122]. Pirolli and Card provide a definition for the profitability of an information source "as the value of information gained per unit cost of processing the source [121,122]." The cost of processing includes the amount of time spent, opportunities lost, and any other resources used [129]. It is important to note that foragers may modify their environment to increase the likelihood of finding valuable information, for example, by organizing a pile of documents into subject-based stacks, or by sorting files in a file explorer by some feature (e.g., date of last access) [122].

The Information Scent Model

The information scent model is of particular interest to this dissertation, as it relates to how people make decisions to use information that they encounter in an information environment. It also describes how information is used at intermediary locations between patches to decide on which directions should be pursued. This concept was originally described as *residue* by Furnas [62], and later adapted by Pirolli [122]. Card and colleagues described scent as the "(imperfect) perception of the value, cost, or access path of information sources obtained from proximal cues, such as WWW links ([30] as cited by [159])." For example, a forager conducting research in a digital library might use cues such as titles, authors, publication dates and citations to determine what documents might contain the best information. However, different environments offer different amounts of scent. If an environment has a strong scent then the cues around information items are sufficient to allow the forager make an accurate assessment of the value of an information item. At the other extreme, if there is no scent, selecting information is completely random [122]. Several works have focused on improving information scent (e.g. [116,156,167]), with the understanding that “better proximal cues lower the cost structure of information foraging and improves information access ([167], p. 129).”

The work of Olston and Chi [116] provided an algorithm for information scent. Based on the keywords used in search, their system – TrailScent – visually augments hyperlinks during browsing, increasing the visual saliency of links encountered with stronger scent. They demonstrated two ways that visual saliency could be increased: by increasing the font-size of links with higher scent, or by using icons that conveyed information about the scent of a link. During a series of tasks using the system, they found that TrailScent significantly reduced the amount of time to completion over browsing and searching alone.

The Diet Model

The diet model focuses on what type of information should be consumed and which information can be ignored. Typically, real world environments offer more than one type of food, just as information environments offer more than one information type. Information “prey” may be selected to maximize the rate of gain of information. So, when considering what information to pursue, people take into account both the return and cost of the information type in an environment. For example, people often completely ignore junk mail, as it contains no valuable information. Or, in order to get an overview of a research field, a forager may only choose literature surveys of known authors or with a large number of citations [122].

2.1.2.2 Sensemaking

Dervin’s model of sensemaking or how people “need to make sense of the world” has been particularly influential ([95], p.29). As Klein, et al., described it “[sensemaking] is a motivated, continuous effort to understand connections … in order to anticipate their trajectories and act effectively [84].” The model developed around the “cognitive gap” that people experience as they try to understand the information they encounter. Sensemaking can also be described as the process of creating situational awareness and understanding during complex tasks, with the goal of decision making [159]. However, sensemaking can equally be about unmaking sense (i.e. learning where there is a gap in knowledge), as it is about making sense (i.e. filling in the gaps in a mental model) [43].

The sensemaking process can be described in several ways, but it typically involves an iterative process of gap-defining and gap-bridging. It is also a two-way process of fitting new information into a mental model and fitting a model around new information; neither of which comes first; new information leads to new or updated models; and, new models result in newly selected and connected information [44,43].

Marchionini describes it as a three step "situation-gap-use" process ([95], p.29):

1. *the situation*: the context for the information need;
2. *the gap*: a question representing a gap in knowledge of what the information seeker needs and what is known to address the current situation;
3. *the use*: the information found, used as an answer to address the gap.

White and Roth summarize the literature identifying five typical steps in the sensemaking cycle ([159], p. 31):

1. the recognition that a gap in knowledge exists;
2. a model of the information needed to address the gap;
3. the search for the needed information;
4. the analysis and reconciliation of the information found to provide new understanding;
5. the creation of new knowledge or follow-up actions based on the fulfilled information need.

2.2 Web-Based Behaviour

Web search often connotes the issuing of keywords to a search engine (such as Google, Yahoo, or Bing), and receiving a series of results [104]. However, recent literature treats web search as synonymous with information seeking, but specifically on the Web. Web search includes "browsing to specific URLs, making sense of found content, iteratively revising a query and disseminating results ([104], p.1)". In this section web-based behaviour refers to the analysis of seeker behaviour on the Web, both generally and specifically in information-seeking tasks.

2.2.1 Web-Based Information-Seeking Tasks

Kellar et al. created a categorization of task-based needs in seeking behaviour and found four high-levels categories [81]:

- *Fact-finding* is the case where the seeker is looking for a specific piece of information like the weather, a recipe, or a piece of software;
- *Information Gathering* involves collecting information from multiple sources over a period of time, and unlike fact-finding there may be no set finish;
- *Browsing* involves reviewing information and web sites with no specific goal in mind, although, a broad goal such as “what’s new” may be set;
- *Transactions* are some online actions such as sending an email or banking.

Finally, Kellar et al. describe all tasks that do not fit into the previous four categories as *other*. Kellar et al. also noted that these categories aligned closely with other previously characterizations of web-based information seeking in reviewing the work of Choo et al. [35], Morrison et al. [108], and Sellen et al. [133], and created the following table (see Table 2.1) to align the categorizations from previous research.

Broder [23] also provided a widely cited taxonomy that focuses on seeker behaviour in performing search engine queries, or the “need behind the query” taxonomy. Broder’s taxonomy contains three high-level needs: informational, navigational, and transactional. In navigational queries the intent is to reach a specific website. Informational queries, broadly, aim to acquire information, which is assumed by the searcher to be available on at least one website. In transactional queries seekers are looking to engage in a particular web-based activity; e.g., looking for a website to play a game or to view online videos.

Table 2.1. The summary of goal categorization for information-seeking tasks, provided by Kellar et al. [81]

Choo et al. (2000)	**Morrison et al. (2001)**	**Sellen et al. (2002)**
Informal Search	Find	Finding
Formal Search	Collect	Information Gathering
Undirected Viewing	Explore	Browsing
Conditioned Viewing	Monitoring	N/A
N/A	N/A	Transactions/Communications/ Housekeeping

2.2.2 The Principle of Least Effort in Information Seeking

Zipf's *principle of least effort* describes how people naturally choose solutions that they find the easiest to acquire as long as the solution is minimally acceptable [171]. In terms of information seeking, Zipf's principle has been used to describe how seekers tend to use the most convenient methods and source of information possible, and then stop seeking information when minimally satisfactory information has been found [33]. This tendency has be exemplified in particular information seeking scenarios where people will choose a less cognitively demanding strategy for getting to target information, rather than choosing the most efficient strategy [146].

2.2.3 Characterizing and Improving Search Performance

Moraveji et al. [101] identified several factors that affect an individual's information seeking performance, including knowledge of search engine features [72, 102] and the resources being sought [160], topical expertise [161], level of general literacy [87], and differences between task types [9].

There has been much work in characterizing the differences between search experts and novices, and in identifying characteristics that predict search performance. Several studies have used different seeker characteristics to determine who is an expert, including: having more than 50 hours experience on the Web [90], browsing the Web more than 5 hours/week [83], and performing searches as part of a job for at least 3 years [72]. Behavioural differences have also been between novices and experts, where experts tend to take more [22] or less time [130] to complete search tasks, use more query terms [72,158], use advanced operators [8], and have higher [83,90] or comparable performance levels [22].

White and Morris [158] used a simple approach of identifying advanced search engine users by their use of four search operators (any of quotation marks, '+', '–', or 'site:'). By having external judges rate the relevance of Web pages seekers visited for a given query, they showed that, on average, the identified advanced users visited pages judged to be more relevant than non-advanced users.

Recent work has also examined how search performance can be improved by increasing search engine knowledge. A controlled study showed that seekers who were taught advanced search engine functionality that would greatly improve their performance on a search task, were able to successfully apply the knowledge, both at the time of learning about the functionality and a week later [102]. Bing recently added a rewards program[4] that provides periodic tasks, each of which introduces new search engine functionality. When tasks are completed, seekers earn points redeemable for merchandise. A Google A Day[5] allows seekers to practice their search skills, by providing daily questions. Seekers race against a timer to search for an answer, and if they become stuck can receive a hint query that will lead to the correct answer.

2.3 Tools and Techniques for Supporting Information Seekers on the Web

In this section work that aims to support information seeking is briefly described. This is divided into three parts: the original visions of the Web, work that has developed tools to improve browser use based on observed information seeking behaviour; and, work that seeks to abstractly represent information to improve information seeking in general.

2.3.1 Original Concepts of Hypertext

Original concepts of hypertext systems, which would eventually lead to the World Wide Web, incorporated the idea of augmenting paths based on their usage and explicit annotations. Venevar Bush famously wrote about his concept of the Memex system in 1945 [28]. In his vision he described how the indexing of documents would be created by individuals explicitly creating trails through documents, annotating them to make clear linkages, and sharing them for others to follow. Less well known is the work of Paul Otlet, who in the 1930's started developing a physical classification system in an elaborate library, called the Mondaneum [170]. In addition to a top-down classification system, he created a bottom-up system that collected the history of people who accessed particular pieces of data. This created implied relationships between individual resources. In this way a growing record would augment each data source providing a secondary, evolving access method based on the accesses of previous people. Otlet referred to this as the 'social space' of the document.

---

[4] http://bing.com/rewards
[5] http://agoogleaday.com

2.3.2 Browser-Based Tools

These conceptualizations of hypertext systems (especially those of Bush) have been echoed in web browsing tools. Some tools have been created that allow seekers to explore pre-defined paths (e.g. [109]), that automatically organize history into paths, or other graphical structures (e.g. [37,103,124]). Tools that have been designed to improve the use of web browsers generally include improved mechanisms to facilitate navigation and refinding previously visited content. For example, an improved back button was proposed to allow seekers to jump to important hub pages that have been previously visited [37,99]. Others provide more detailed expansions of classic browser history and bookmarking tools: Kaasten and Greenberg integrate history and bookmarking tools [76]; Takano and Winograd dynamically built a tool for creating bookmarks based on an analysis of visitation and page link structure [144]; MacKay, et al. created a tool to bookmark content within a page instead of the entire page itself [92]; and Morris et al. created a history tool that is organized by navigation after search queries [103]. Teevan et al., created DiffIE, which is a tool that highlights content changes since the seeker last saw the page [148].

2.3.3 Information Visualization for Supporting Information Seeking

Information visualization is the process of transforming data into graphical representations to convey insights [140]. Card describes “Information visualization [as] a set of technologies that use visual computing to amplify human cognition with abstract information [29].” Research in information visualization seeks to find new ways to take advantage of the high bandwidth of human vision. People have the ability to visually scan a complex display, and recognize subtle changes in color, size, shape, texture and movement [135]. Shneiderman highlights the need for interaction for exploring for information and insight with his information seeking mantra: “Overview first, zoom and filter, then details-on-demand ([135], p. 363).” These also form the first four of his seven tasks for visual information seeking:

- *Overview*: allow for a view of the entire collection under examination.
- *Zoom*: change the view to focus in on items of interest.
- *Filter*: hide items that are not of interest.

- *Details-on-demand*: select an item or a grouping of items to receive more details about them.
- *Relate*: make clear the relationships between items.
- *History*: a history of actions that have been performed by the seeker so actions can be undone, replayed, or refined.
- *Extract*: allow features of the data and groupings of items to be extracted for use elsewhere.

2.3.4 Feedback and Persuasion

The area of personal informatics aims to design systems that help people learn about and understand their own behaviour, with the goal of providing new insights, increasing self-control, and promoting the acquisition and maintenance of desirable behaviour [91]. Similarly, persuasive technologies have been described as systems that seek to change behaviour or attitudes, without the use of coercion [59]. Work in education and in persuading users to adopt new behaviours have both highlighted the importance of feedback on personal behaviour for reflection and learning [20,91]. Additionally, observing other skilled practitioners can improve learning [12,101], and knowing what others do can lead to positive choices [149]. Personal informatics and persuasive systems have been created for reflecting on past behaviour and promoting behaviour change in several domains including physical activity [39], environmental impact [61], and webpage visitation [151].

In the domain of information seeking there has been little research into what, and how data should be displayed for feedback. In one study of personal and shared web activity, participants found views containing the data of others to be more useful [151]. The Google and Bing search engines both provide “Search History” functionality that allows seekers to review past queries (google.com/history; bing.com/profile /history). Such systems have been shown to improve performance in re-finding and resuming search tasks [103], but it is unclear whether they would be effective tools for feedback or persuasion. Google’s search history includes “Trends”, which displays a seeker’s most-frequent queries, most-visited sites, most-often clicked search results, and the total number of searches executed over different time frames (see Figure 2.3). While search engines have recognized some value in presenting summaries of behaviour, there is little information about what data should be presented, and what effect it might have on seeker attitudes and behaviour.

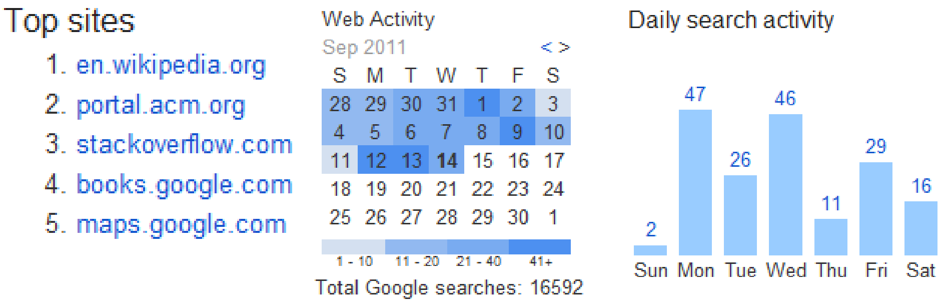

Top sites

1. en.wikipedia.org
2. portal.acm.org
3. stackoverflow.com
4. books.google.com
5. maps.google.com

Web Activity

Sep 2011 < >

S	M	T	W	T	F	S
28	29	30	31	1	2	3
4	5	6	7	8	9	10
11	12	13	**14**	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	1

Total Google searches: 16592

Figure 2.3. The author's search activity displayed in Google Trends.

2.3.4.1 Persuasive Technology

Persuasive technology is defined as "any interactive computing system designed to change people's attitude and behaviours ([60], p.1)". The concept of persuasive technology was developed by B.J. Fogg, and he cites several examples of persuasion on the Web: Amazon.com persuades users to buy more merchandise through recommendations (e.g., "users who bought that item, also bought this item"); and, Iwon.com tries to persuade users to make the site their default search engine by rewarding them with prizes to do so.

Fogg says changes in attitude and behaviour can occur at two levels as a result of successful persuasion: *macro* and *micro*. He describes *macrosuasion* as persuasion used in products that have a sole focus of attempting to persuade (e.g., a website that attempts to persuade people to give up smoking). *Microsuasion,* in comparison, is the use of "smaller persuasive elements" to achieve a different overall goal. He cites the praise that some systems give as an example to promote a target behaviour (e.g., Quicken gives praise to the user for completing mundane but important tasks, such as balancing a cheque book).

Persuasive technology is broadly defined, but focuses specifically on the concerns for designing technology that will lead to a desired behaviour change, rather than on how behaviour can be most effectively presented to learn from others. While Fogg recognizes that social influence can be a powerful persuasive factor (citing several theories, including social learning theory), he most often describes the interventions to achieve persuasion that come from the system designer. Further, he makes clear that persuasive technology "…investigates how people are motivated or persuaded when acting *with* computing products rather than *through* them ([60],

p. 16).” He distinguishes persuasive technology from work in computer-mediated communication or computer-supported cooperative work (described in Section 2.5) where “the computer facilitates communication; it does not persuade ([60], p. 16).” This *is* the focus of this dissertation: facilitating learning through viewing representations of the behaviour of other seekers.

2.4 Interaction History and its Use in Information Seeking

In this section, I describe interaction history in more detail. In particular, I outline how interaction history is used by information seeking tools, such as search engines and recommender systems, where seeker preferences are inferred and can be used as feedback for algorithms to improve accuracy.

2.4.1 Interaction History and Implicit Feedback for Information Seeking

Interaction history has been given many different names and is represented in many different disciplines. For example:

- interaction history in general HCI and CSCW [70,156,157];
- implicit relevance feedback [82,159] in information seeking, retrieval, and web search;
- traces [79] or monitor data [80] in user modelling; and,
- implicit indicators of interest in recommender systems [36].

These works share the fact that they aim to provide systems with information that can be used for improved functionality. Marchionini describes interaction history in terms of “moves”:

> *Moves are finely grained actions manifested as discrete behavioural actions such as walking to a shelf, picking up a book, pressing a key, clicking a mouse, or touching an item from a menu. Moves are manifestations of tactics, and although they are conceptually uninteresting individually, taken in context, they offer observable evidence for interface assessment and mapping the intellectual at higher levels of action ([94], p.74).*

2.4.2 Interaction History for Implicit Feedback

Approaches for employing seeker interactions with systems as implicit feedback have been used for retrieval, filtering, and recommending many different information items, including: web links, webpages, academic articles, email messages, movies, books, TV shows, jobs and stocks [82]. Implicit feedback for improving information retrieval dates to work carried out in the 1970's [82,95]. Implicit relevance feedback can be used for query expansion – adding additional search terms automatically to help narrow an information seeker's query – for short-term modelling, which represents a seeker's immediate information need. However, it can also be used to build a profile (or model) to capture a seeker's longer-term interests and preferences [82]. Traditionally, such relevance feedback methods have required seekers to explicitly interact with the system to directly provide feedback for the system. Providing keywords representing interests, selecting or otherwise marking documents as exemplars of interests, or by answering questions, are all examples of explicit relevance feedback [82]. However, because the cost, in terms of seeker effort, is high, and because seekers are not always immediately aware of any benefit, it can be difficult to adequately collect and maintain explicit feedback data [82,162]. Because of these reasons, explicit feedback is not often used in practice [147].

Implicit feedback, through interaction history, is appealing because it reduces the cognitive load for seekers [82,162]. White, et al., found that seekers preferred giving a search system implicit relevance feedback – from actions such as scrolling, reading, and saving – to explicitly marking search results as relevant. This was particularly true for novice searchers. Further, the preference for implicit feedback being used became stronger as the complexity of the search task increased [162].

Kelly and Teevan reviewed important work that has investigated the use of implicit measures for feedback. Table 2.2 shows a categorization of user actions by the minimum scope they act upon (in terms of an information item or document) [82].

Kelly and Teevan note that placing user behaviour in specific categories in Table 2.2 is not always clear-cut because the scope and intent of interactions can be ambiguous. They use the example of creating and saving a new document. If the "save" action is considered alone then it may seem that this is a "Retain" action (see the Retain behaviour and Object scope cell in Table 2.2), while it is likely better categorized as a "Create" behaviour. The table is intended to only provide a sample of possible useful behaviours rather than an exhaustive list [82]. To exemplify the work in the area of implicit feedback to improve information access, a selection of three works is described. However, work in this area has been quite extensive (see [82] for a review). These descriptions are intended to demonstrate how different types of interaction history have been assessed for implicit feedback only. The first two works focus on how implicit interaction history relates directly to individual assessments of interest or relevance. However, implicit feedback measures can be used as votes for the overall quality of an information item, or they can be used in combination with collaborative filtering algorithms (described below) to offer personal recommendations based on people who have similar interests.

Table 2.2. A categorization of user behaviour by minimum scope that the behaviour acts upon that can be used for implicit feedback. Originally proposed by Oard and Kim [114], and extended by Kelly and Teevan [82] (additions are highlighted to clarify provenance). Table reproduced from [82].

			Minimum Scope	
		Segment	**Object**	**Class**
Behaviour Category	**Examine**	View Listen Scroll Find Query	Select	Browse
	Retain	Print	Bookmark Save Delete Purchase Email	Subscribe
	Reference	Copy-and-paste Quote	Forward Reply Link Cite	
	Annotate	Mark up	Rate Publish	Organize
	Create	Type Edit	Author	

Arguably the most well-known type of implicit feedback in information retrieval is Web link analysis. Because webpages provide a high number of links and are authored by people, the links on webpages can be used as feedback with regards to the importance of a page. For example, the work of Kleinberg [85] identified *hubs* and *authorities.* Authorities provide an authoritative source of information on a topic, while a hub provides links to many authorities.

Good authorities have many incoming links from hubs, while good hubs have many outgoing links to many authorities. Larry Page, et al., famously described the PageRank algorithm, which adapts the social network analysis concept of eigenvector centrality to webpage linking. As in relationships for people, important pages can be identified by being connected to by many other pages, but not all connections are equal. Links from other highly-linked-to pages are worth more [117].

2.4.2.1 Group-Based Implicit Feedback

Some work has looked at the use of implicit feedback from a group or community of individuals to improve search. The motivation is to allow a wider variety of history to be incorporated [54]. For example, the work of Smyth [137] has looked at user-created search communities. A search community is created by a single user, and other friends and colleagues can be invited to the community. Community interaction history is aggregated for all community users. These are used for reordering search results, but also to compute the similarity to other existing communities. When sufficiently similar communities are found, the interaction history is shared between groups, vastly increasing the coverage of the interaction history.

2.4.2.2 Personalization in Search Engines

Teevan, et al., developed an algorithm for personalizing search based on automatically creating a model of interests and activities from a seeker's past queries, browsed webpages, personal documents, and email [147]. They showed that re-ranking search results for a seeker based on their personalized profile significantly improved user relevance ratings for search results over a typically ordered result list.

2.4.3 User Modelling and Personalization

The area of *user modelling* has focused on creating detailed descriptions of users based on their passive or active behaviour using interactive systems; particularly in the educational domain [79,80,86,96,127]. The goal of creating a user model is to allow for computational reasoning to occur that can provide personalized support for a particular user. In research exploring adaptive hypertext systems, personalized support has taken the form of adaptive information presentation (e.g., hiding text that an e-learner has already been tested on) or adaptive navigation support (e.g., adaptively annotating links based on preference) [25].

Similarly, personalization has been seen in information retrieval and recommender systems, which filter larger lists and/or reorder lists based on a user's profile (model). For example, a user model describing a user's interests in movies is often based on their previous ratings of other films. This would allow a recommender system to give a list of movie recommendations based on this model (collaborative filtering and recommender systems are described in more detail below).

2.4.3.1 Open User Modelling

Open user and open learner models allow users to inspect a system's beliefs about their knowledge or preferences. Openness usually takes the form of a visual representation of the model (or part of a model) that a user can inspect to understand the system's current beliefs about the user and how it arrived at those beliefs. The user may also be able to enter into a dialogue with the system to negotiate beliefs. The predicted benefits of such an approach is that it will allow the system to create a more accurate user model, but it also allows the user to reflect on the system's beliefs, thereby providing a better understanding of their own beliefs [47,153].

2.4.3.2 The Ecological Approach

McCalla's "ecological approach" describes a method whereby models of users and their interactions are attached to the information they interact with [96]. The goal of the ecological approach is to find patterns within the interactions of users in e-learning environments to support them in their learning goals, thus providing a new way to describe each digital learning object in a way that is appropriate to each individual learner. Each learner and each learning object (information item) is uniquely described by a history of interactions. In the ecological approach, the rich history of interactions are attached to objects and learners. Computational agents act as intermediaries and work on behalf of the learning objects and learners by interpreting low-level interaction data, and negotiating the exchange of data and interpretations.

2.4.4 Social Interaction History

Indratmo and Vassileva proposed a framework of social interaction history [75]. They described social interaction history as:

> *[the] records of communication among users in an information space. While interaction history concerns user-object interaction, social interaction history focuses on user-user interaction regarding an information item ([75], p.540)*

So, while a reader dog-earing a page would be a form of interaction history, the provided definition means that social interaction history would be people discussing the dog-eared page. It should be noted this definition differs from the "personal versus social" dimension of Wexelblat's framework of interaction history, which suggests the social aspect of interaction history is simply the aggregation of interactions with an information item [157]. Indratmo and Vassileva's definition of social interaction history is demonstrated in Figure 2.4.

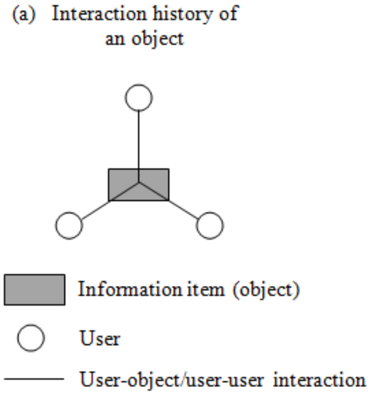
(a) Interaction history of an object
(b) Social interaction history of an object
Information item (object)
User
User-object/user-user interaction

Figure 2.4. The distinction drawn between interaction history (a), which is limited to human-object interaction, and Indratmo and Vassileva's social interaction history (b) the direct user-user communication that takes place around an object.
Recreated from [75].

In a similar manner to the previously described work on interaction history, Indratmo and Vassileva found that social interaction history correlated strongly with the user ratings. This work differs from the motivation of this dissertation, which instead seeks to examine how the aggregated interaction history (i.e., not including direct communication) on an information item can guide information seekers.

2.5 Relevant Concepts from Computer Supported Cooperative Work

Computer Supported Cooperative Work (CSCW) has been described as "a generic term which combines an understanding of the way people work in groups with the enabling technologies of computer networking and associated hardware, software, services, and techniques ([168], p. 6).” Groupware describes technology that is designed for collaborative work to be carried out with groups. Groups are typically teams that are working together on shared tasks under time constraints [168]. Baecker and colleagues [10] describe groupware in terms of the well-known CSCW matrix (see Figure 2.5), which illustrates groupware along time and space axes. Collaboration with groupware can occur in different *places*: collocated (e.g., in the same room) or remotely (e.g., rooms are separate and can be separated by small or great distances). Work can take place at different *times*: synchronously (i.e., at the same time) or asynchronously (i.e., at different times).

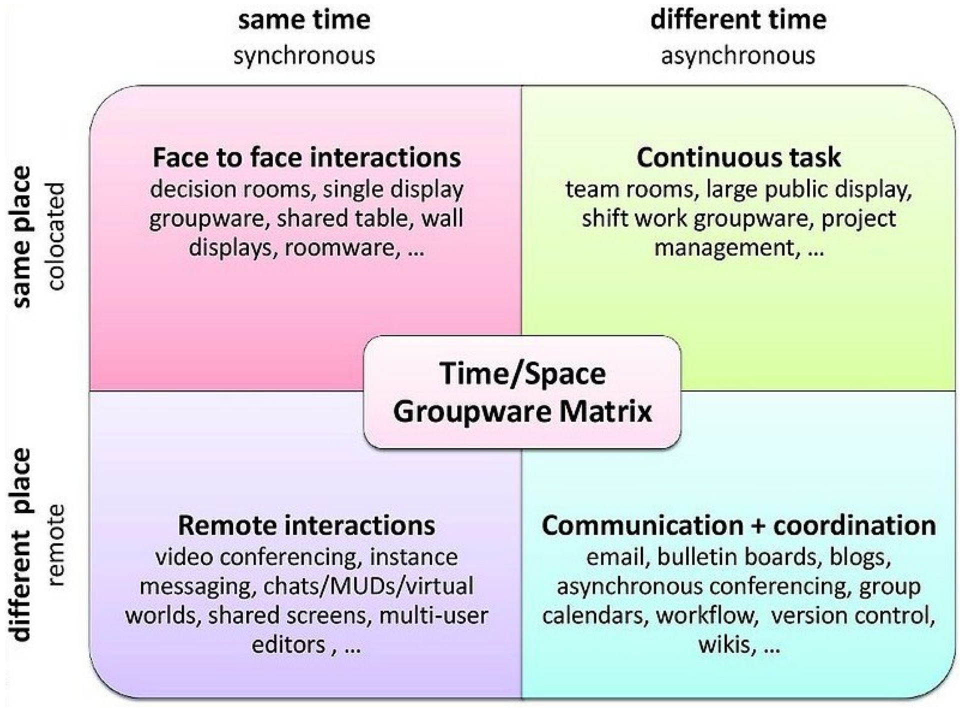
same time
synchronous
different time
asynchronous
same place
colocated
different place
remote
Face to face interactions
decision rooms, single display groupware, shared table, wall displays, roomware, ...
Continuous task
team rooms, large public display, shift work groupware, project management, ...
Time/Space Groupware Matrix
Remote interactions
video conferencing, instance messaging, chats/MUDs/virtual worlds, shared screens, multi-user editors , ...
Communication + coordination
email, bulletin boards, blogs, asynchronous conferencing, group calendars, workflow, version control, wikis, ...

Figure 2.5. The CSCW matrix, positioning groupware technologies in a time versus place grid. Public domain image from [164].

In this section, I review some of the key terms from the field of CSCW that underlie collaborative information seeking, to provide a basic understanding of how people work together. I first focus on the concept of coupling, which describes the degree of independence that collaborators have in their work towards a shared goal. I then describe work in informal collaboration, which describes work that is both discretionary (it does not necessarily need to happen) and where the collaboration is unplanned. These concepts are of concern to this dissertation because in most information-seeking tasks that involve other people, seekers are not planning or relying on advice from others to complete a task; they are instead loosely coupled and working independently. However, even if seekers are working completely independently, an appropriate awareness of what others have done can give guidance in completing many tasks.

2.5.1 Coupling

Coupling describes the amount of work that people can do individually before they require communication or actions from another person [67]. Coupling is usually described as being on a spectrum between *loose* coupling – where collaborators can work independently – to *tight* – where work between individuals is highly dependent on others. In both synchronous and asynchronous collaboration, situations can arise where collaborators move between tightly coupled and loosely coupled interactions. In this dissertation, I focus on loosely coupled work where individuals conduct tasks and use tools that allow them to work individually (such as using a typical web browser).

Pinelle and Gutwin [120] carried out interviews and field studies of mobile healthcare providers, who collaborated with one another in a loosely coupled fashion. They summarize four relevant results from their observations:

- When workers work autonomously and communication requires a large amount of effort, workers must carefully consider the amount of effort required before sharing information. Explicit (direct) communication is typically for urgent issues only. They prefer methods with the lowest overhead and effort.

- When artifacts and work locations are shared, "traces" of work are left behind. Workers prefer to learn about one another this way, because the evidence from work is stored implicitly (as a result of performing other tasks) and it is easy for others to discover and retrieve.

- When workers need to explicitly communicate with one another, they prefer using asynchronous channels, as it provides them with flexibility in dealing with the variability in their schedule and location.

- Workers whose jobs require mobility and working alone have difficulty with synchronous communication. These constraints limit the amount of information that can be delivered back and forth.

2.5.2 Awareness to Support Collaboration

Dourish and Bellotti described awareness as “an understanding of the activities of others, which provides a context for your own activity ([52], p. 107).” This means awareness information can help a seeker refine their own tasks. For example, two collaborators who are scouring a website for hidden information can be aware of each other’s activities, thus avoiding duplicating effort. Dourish and Bly more specifically described what is involved in achieving awareness and its longer term benefits in working with collaborators: “Awareness involves knowing who is ‘around,’ what activities are occurring, who is talking with whom; it provides a view of one another in the daily work environments. Awareness may lead to informal interactions, spontaneous connections, and the development of shared cultures—all important aspects of maintaining working relationships which are denied to groups distributed across multiple sites ([49], p.541).”

Gutwin and Greenberg provide four characteristics of awareness:

1. Awareness is knowledge about the state of an environment bounded in time and space.
2. Environments change over time, so awareness is knowledge that must be maintained and kept up to date.
3. People interact with and explore the environment, and the maintenance of awareness is accomplished through this interaction.
4. Awareness is a secondary goal in the task – that is, the overall goal is not simply to maintain awareness but to complete some task in the environment ([67], p. 416).

They provide a framework for *workspace awareness*, which is an “up-to-the-moment” knowledge of other people’s interactions with the shared workspace. Here a workspace is considered a bounded area where people can manipulate objects that form the focus of collaboration [67].

Gutwin and Greenberg's framework is focused on supporting awareness for synchronous groupware; however, they also consider elements of workspace awareness that are related to the past (the history of the workspace), which is the focus of this review. The elements in Table 2.3 provide considerations for groupware designers who aim to support workspace awareness. The table provides specific questions relating to *how* (How did this happen?), *when* (When did this happen?), *who* (Who was here? And when?), *where* (Where has that person been?), and what (What did that person do?).

Table 2.3. Elements of workspace awareness that relate to the past. Recreated from [67].

Category	Element	Specific questions
How	Action history Artifact history	How did that operation happen? How did this artifact come to be in this state?
When	Event history	When did that event happen?
Who (past)	Presence history	Who was here, and when?
Where (past)	Location history	Where has a person been?
What (past)	Action history	What has a person been doing?

Tam and Greenberg [145] proposed an extended framework of Gutwin and Greenberg's workspace awareness, which focused on asynchronous changes to shared documents; they refer to this as change awareness. They define *change awareness* as the ability of people to track the changes made asynchronously by others in working with shared documents or workspaces. Such asynchronous changes would happen at different times, where only a single collaborator is working on the document at a time. Tam and Greenberg point out that providing visuals of all history may be overly complex; changes can happen in many places, can be overlapping and can be made by many individuals. So, people may want to query the interface to see only certain changes made by certain individuals at certain times. They describe three views of change that viewer may want:

1. an artifact-based view,
2. a person-based view, and
3. a workspace-based view.

They then re-examine and extend Gutwin and Greenberg's questions relating to how, when, who, where, and what (see Table 2.3). For example, Table 2.4 shows Tam and

Greenberg's extended questions relating to 'who' for change awareness. Tam and Greenberg's work relates directly to this review, which considers interaction history that arises in information seeking on the Web. However, information seekers can be considered consumers rather than authors or editors of the artifacts (webpages and search results) they work with. Therefore, questions in Tam and Greenberg's framework such as 'Who has changed this artifact?' have less direct application. A framework does not exist to guide the awareness provision for collaborative information seeking on the Web.

Table 2.4. Specific questions relating to 'who' based on an artifact-based, person-based, or workspace-based view of change awareness. Recreated from [145].

Information Elements	**Artifact based view**	**Person based view**	**Workspace view**
Presence history	Who has looked at this artifact?	Who has this person interacted with?	Who has been in the workspace?
Identity	Who has changed this artifact?	Who made changes with this person?	Who has looked at the workspace?
Readership history			Who has made changes to the workspace?
Presence history			
Authorship history			

2.5.2.1 How Awareness Information is Used

Awareness can reduce the amount of effort, decrease the number of errors, and increase the efficiency in collaborative activities. Gutwin and Greenberg highlight five types of activity that workspace awareness improves [67].

1. *The management of coupling.* In many face-to-face collaborations people move between tight and loose coupling, and awareness of others in the group can support these transitions. Even if work is loosely coupled, collaborators benefit from knowing what others are doing and have done.
2. *Simplification of communication.* Having an awareness of a workspace and the artifacts in it allow users to simplify their communication, because they have a shared representation to reference; this makes communication more efficient.
3. *Coordination of actions.* Collaborative tasks require that actions performed in carrying out the task are properly coordinated. Having awareness about what others are doing and

have done allows collaborators to make sure they align their actions correctly. For example, two seekers (John and Mary) are conducting a literature survey and aim to divide the task so that they can maximize their coverage. Because John can see that Mary has covered particular articles he can skip them to avoid duplicating effort.

4. *Anticipation.* In anticipation people take actions based on predicting or expecting what others will do or need in the future. Anticipation can aid in collaboration in several ways: people can avoid conflicts (such as duplication of effort); they can provide resources for a collaborator before they are needed; or they can get ready for a required action while others are completing a requisite task.
5. *Assistance.* Having an awareness of others and when they encounter difficulty allows users to appropriately switch to a helping role, and provide direction or help in a timely manner.

2.5.3 Overview of Informal Collaboration

Gutwin et al. defined informal collaboration as the “unplanned and opportunistic interactions that occur in an impromptu fashion in the workplace ([68], p. 1413).” Unplanned communication tends toward the spontaneous and arises out of chance encounters in the environment when opportunity presents itself [68,89,163].

Billotti and Bly found, in observations of a large international design firm, that remote collaborators were severely disadvantaged in their communications [16]. Collaborators who were not local had no means of learning what their colleagues were doing and no opportunities for chance encounters to take place. Because face-to-face conversation is a lightweight entry point into work-related discussions, there was no technology available that could provide similar opportunities [16]. Spontaneous interactions facilitate many quick exchanges, which provide useful information. The resulting awareness of what others are doing creates shared knowledge and facilitates further interaction to occur. Therefore, informal interaction is key to providing a successful collaborative work environment [16,163].

From this work, colleagues in a workplace can be considered as having an ongoing active collaboration that is loosely-coupled. The colleagues are all working towards a shared objective (e.g., the successful operation of the company), and from time to time they enter into a tighter forms of collaboration when appropriate.

2.5.3.1 Characteristics of Informal Collaboration

Gutwin et al. [68], summarized the work of Kraut et al. [89], Whittaker et al. [163], and Bellotti and Bly [16], highlighting several characteristics of informal collaboration described below. In the summarized work, communication (i.e., face-to-face conversation) is used as an exemplar of lightweight collaboration.

Informal interactions are common in collocated groups

Kraut et al. categorized the most common communication form, face-to-face meetings, from observations made in a R&D firm [89]. Conversations were categorized as one of the following:

- *scheduled*: a conversation that was previously scheduled or planned (formal);
- *intended*: an initiator set out specifically to visit another party (informal);
- *opportunistic*: an initiator planned to have the conversation at some time, and took advantage of a chance encounter (informal);
- *spontaneous*: the initiator did not plan to talk to others (informal).

Kraut and colleagues describe *informal interactions* as being those that are not scheduled (i.e., one of intended, opportunistic, or spontaneous). *Unplanned* interactions can be either opportunistic or spontaneous (see Figure 2.6). It was found in this categorization that 88% of conversation were informal (not scheduled), 53% were unplanned, 21% were opportunistic (they were started by the initiator seeing another person), the other 32% were completely spontaneous [89]. This type of collaboration was echoed in a study of information seeking where 62.9% collaborative, collocated searches in the workplace were unplanned [5].

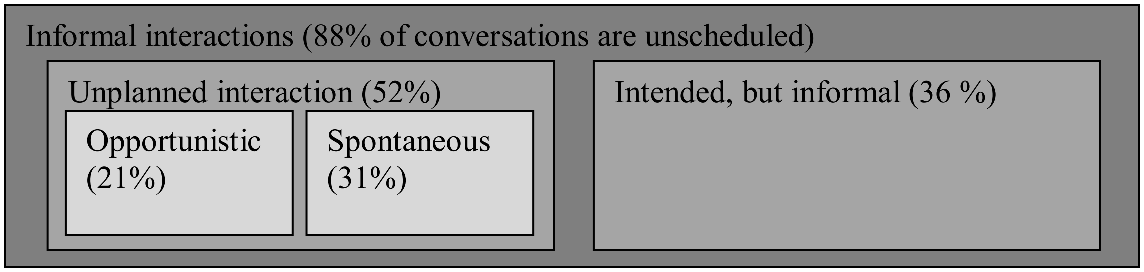
Informal interactions (88% of conversations are unscheduled)
Unplanned interaction (52%)
Opportunistic (21%)
Spontaneous (31%)
Intended, but informal (36 %)

Figure 2.6. Kraut et al.'s organization of workplace interactions. Informal interactions are all interactions that are not scheduled. Including planned and unplanned interactions. Unplanned interactions can be either opportunistic or spontaneous.

Informal interaction is grounded in awareness of the work environment

People maintain awareness through the people, objects, and activities that their colleagues are undertaking in the environment [*68*]. Kraut et al. ([89], p. 10), describe the process of social browsing:

> *…people walk down the halls, peering into open offices and public spaces as they go to the printer, copy machine, bathroom, or other ultimate destination… This process of browsing the social environment while on other business provides people with a substantial amount of information about the world in which they live.*

Bellotti and Bly highlighted that mobility was a key to informal collaboration, and observed that people often walked around to gain information. They described a "walkabout" as simply moving around to see what others are doing and working on, but in overhearing conversations and viewing the environment valuable information is gained. There is a lack of such an equivalent tactic for workers at a distance because "technology only affords explicit communication rather than the kinds of implicit communication available through co-presence and mutual awareness ([16], p. 214)."

Informal collaboration can be triggered by people, objects, actions, or interaction

Cues for opportunistic collaboration can include seeing a person, seeing an object in the work environment, seeing someone doing a particular action, or hearing others in conversation [*68*]. The most common cue is seeing someone else, but there is no clear distinction in the form of the cue and the quality of the communication [89]. Whittaker et al., found that people used documents in over half of the informal conversations observed, and that they were used both as cues and props for the participants to focus on once communication had started [*163*]. Billotti and Bly observed that being close enough to overhear other conversations led people to engage in informal communications that they would not have otherwise had [*16*].

Informal collaboration is often discretionary

Informal communication allows participants to make choices about whether they need to engage in them, and whether they ever need to occur at all; often such conversations can be easily postponed. In these situations the work being undertaken is loosely coupled. The communication may never need to happen at all, which allows individuals to decide if the cost of entering into communication is too high [68].

Informal interactions are easy to initiate

The only requirement for engaging someone in informal collaboration is that they are available and there is a topic to initiate communication [89]. Being collocated with others allow judgments about availability to be simple, and the environment itself provides many topics to initiate communication and collaboration; people can simply approach one another, point at something, and ask a question about it [*68*]. In many situations, if more effort is required people simply will not engage in communication [*16*].

2.5.3.2 Basic Requirements for Informal Collaboration

Gutwin and colleagues [68] described basic requirements for informal collaboration to be successfully supported in groupware systems.

- *Individual work*: people must be able to work on their own tasks, until an opportunity for collaboration presents itself;
- *Awareness of the community*: informal collaboration is based upon having an up-to-date and rich awareness of other people and the environments they are working in;
- *Lightweight initiation*: due to informal communication being discretionary, engaging in collaboration must be low-cost or people will choose simply not participate;
- *Interaction with task artifacts*: many tasks involve objects that need to be referred to and manipulated during communications.

2.6 Collaboration and Social Support in Information Seeking

In this section, I review work that is directly related to the work carried out in this dissertation in Chapters 4 and 5. First, I describe the main work in characterizing how and why people engage others during the information-seeking tasks. Second, I describe work in social navigation systems that aim to provide support in information-seeking tasks through using the collected actions of others, but where the use of the information is typically unplanned. Third, I describe work in collaborative information seeking where information seekers are actively engaged and working together towards a shared goal.

2.6.1 Social Information Seeking and Social Search

Many researchers have described information seeking as task that is completed individually; however, this view is steadily changing to recognize the role that other people play throughout an information seeker's search process [34,54]. Social search, social information seeking, and other social software systems (such as social networking sites) have the common feature of fast and simple sharing of information with others (who may not necessarily be known) [100]. Evans and Chi cast the net widely in their characterization of social search pointing out that the term has been used to describe information seeking that:

- makes use of social or expertise networks;
- is carried out in a shared workspace; or
- contains information that is inferred from observed behaviour (such as through data mining) [54].

They provide the following definition to frame their work in social search: "'Social search' is an umbrella term used to describe search acts that make use of social interactions with others. These interactions may be explicit or implicit, co-located or remote, synchronous or asynchronous ([54], p.657)." This definition includes situations in social search where there is direct collaboration between searchers.

Research into social search encompasses the use of feedback that can be either explicitly or implicitly created by users. Implicitly created data can come from inferred indicators of interest or quality, and are based on system observed user behaviour. Explicitly created data comes from user actions that directly inform the system, such as votes or ratings. Both implicitly and explicitly created data are used in what Chi calls *social feedback systems*[6], which use this data for ranking information items [34].

Chi distinguishes social feedback systems from *social answering systems*, which allow people to explicitly post questions for people to answer. Answerers can be within the asker's social network, extended social network, or just have some interest or expertise to share [34].

---

[6] Chi's use of the term "social feedback" is not in direct conflict with the use of it in this dissertation. Although, I define social feedback systems as providing information to people, not systems, and the source of the feedback information as coming exclusively from implicit data.

Social answering systems are sufficiently different from the general information seeking process that they are beyond the scope of this review.

2.6.1.1 A Study and Model of Social Search

Evans and Chi have recently provided a model of "social information search" [54]. Their model is based on a survey of 150 Amazon -Mechanical Turk workers. Their survey focused on communications with other people before, during and after their most recent information seeking incident. The model (see Figure 2.7) extends current information seeking models by characterizing key points of social interaction. I focus on describing the points of information exchange between the seeker and others throughout the information seeking process.

Communication with Others Before Search is Common

Before the search, over 30% of respondents were motivated to conduct their information-seeking task from an external source; e.g., one respondent conducted research based on a customer request. The others, approximately 69% of respondents, were self-motivated in performing their tasks.

After the establishment of an information need, and regardless of whether externally or internally motivated, the influence of other people continued to play a role. In roughly 43% of all tasks, seekers communicated with others. This helped the seeker in setting guidelines for the task and provided criteria for refining their information need.

Communication with Others During Search is Rare

Evans and Chi subdivided respondent search queries using Broder's taxonomy of the "need behind the query" ([23], previously described in Section 2.2.1). They found that during the search engine use, only two searchers (2 of 89) reported explicit communication with others. While another 5 searchers reported using implicit communication channels, through searching the discussion boards of Wikipedia.

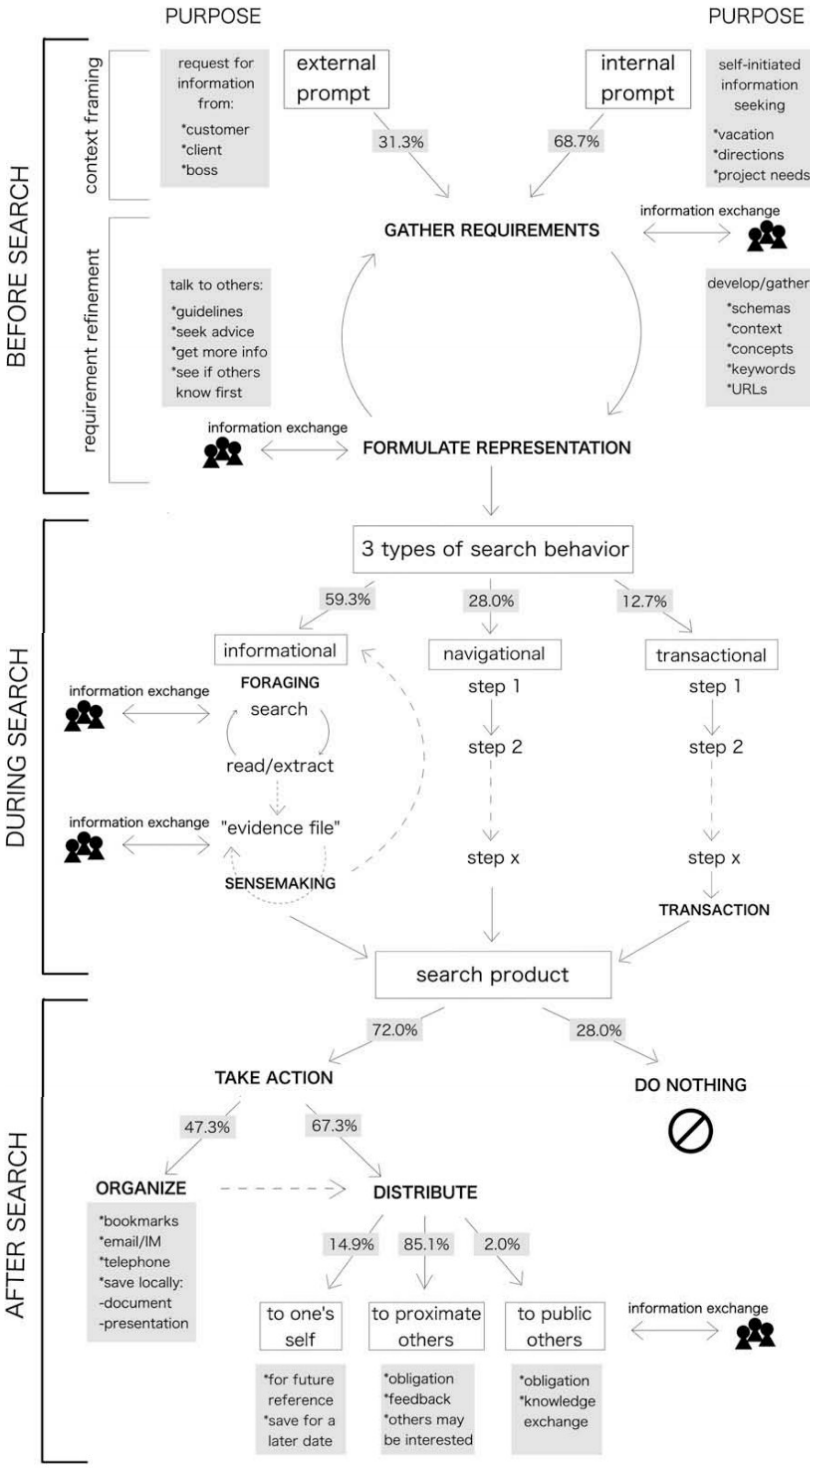

Figure 2.7. Evans and Chi’s canonical social information seeking model (reproduced from [54] with permission).

Communication with Others After Search is Common

After the search was completed a large portion of seekers (28%) took no action with their search results. 47.3% of all seekers who were surveyed (71/150) reported creating some organizational artifact, including printing the results, bookmarking the page, or creating a presentation, and 67.3% of all users (101/150) reported distributing the results somehow.

Seekers shared their results often with others, but most of the time this was with users who were proximate to them, most often face-to-face or over the phone (86/101). These people tended to be close friends or colleagues. Only two seekers shared their results with the general public, by updating websites to include their search results.

While sharing results of information-seeking tasks was widespread among seekers, the reasons for sharing information were quite diverse. Some seekers reported feeling obliged to share what they found or because they knew others who might be interested. Other seekers shared the information because they wanted some advice or feedback on what they had found.

Communications with Others During Difficult Searches is More Common

Evans and Chi also reported on a second study that surveyed an additional 150 seekers, and placed the focus on the details around failed searches. When searches were particularly difficult, seekers reported greater social involvement with others during foraging and sensemaking. When seekers had difficulties during the search process they looked for help in refining the query or getting feedback about the direction they were pursuing. In these cases, seekers turned to others for help during the search 28% of the time, compared to only 1.3% of the time in the original study.

2.6.2 Social Navigation

Social navigation is a concept coined by Dourish and Chalmers [50] that characterizes an approach for navigation in information spaces. Social navigation is concerned with guiding people through a larger information space based on the interpreted actions, expressed thoughts or explicit recommendations of others. Svensson and colleagues refer to social navigation as being “driven by the actions from one or more advice providers ([142], p.342)”.

Dourish and Chalmers also made the distinction between *spatial, semantic* and *social* navigation. Spatial navigation refers specifically to movement through real or virtual space, such as in a virtual world. “[Semantic navigation offers] the ability to explore and choose perspectives of view based on knowledge of the semantically-structured information ([50] , p.2).” “[In] social navigation, movement from one item to another is provoked as an artefact of the activity of another or a group of others ([50], p.1).” Traditional interfaces contain a lot of semantically structured information in which control is limited to a few editors, whereas social navigation information can be structured by the actions of many people [50].

Svennson and Hook [143] elaborate on two further points in characterizing social navigation. First, an advice provider for navigation can be a person or artificial agent. Second, an important distinction should be made between *direct* and *indirect* social navigation. Direct social navigation results from the direct communication or advice (e.g. through an emailed link) from another person or group. Indirect social navigation requires the monitoring and analysis of behaviour of a group of people upon which recommendations are inferred.

Hill, et al. [70] originated the idea of “[capturing] on computational objects (e.g. documents, menus, spreadsheets, images, email) the events that comprise their use ([70], p.3).” Their goal was to graphically depict these histories to give the viewer a graphical representation of *wear* such that the history of their use could be better understood. They referred to this as *computational wear*, alluding to the metaphor of the wear that objects accrue in the physical world. They developed *read wear* and *edit wear* tools that visually depicted the reading and editing history of all users working on a document (see Figure 2.8). Their goal was to make apparent from glancing at graphical traces the answers to questions such as: What parts of the document are most stable? Who wrote that part of the document? Did my colleague read that part? What is the relative age of that paragraph? [70]. These questions relate to the concept of awareness (as previously discussed).

Dourish has said that the “navigation” in social navigation is best understood as an information-seeking task [51], which is related to making decisions; seekers make many decisions during the information seeking process [63]. Goecks says there are three general activities in making decisions using social navigation systems: (1) becoming aware of choices or paths; (2) seeking information about the paths or choices; (3) considering the information that is available and making an informed decision (these often include community data) [64].

Wexelblat and Maes have also suggested that making decisions is central to using social navigation systems [156]. Goecks concurs, saying that decision making is the principal activity that users engage in when using social navigation systems [63]. However, Wexelblat and Maes suggest that because social navigation systems typically aggregate data over all users, awareness questions are limited to those that answer *what* (e.g., "What link has been taken?" or "What link is the most popular?") [156]. Recall, that awareness related questions might include: what has been done here, who has done it, why did they do it, and how did they do it.

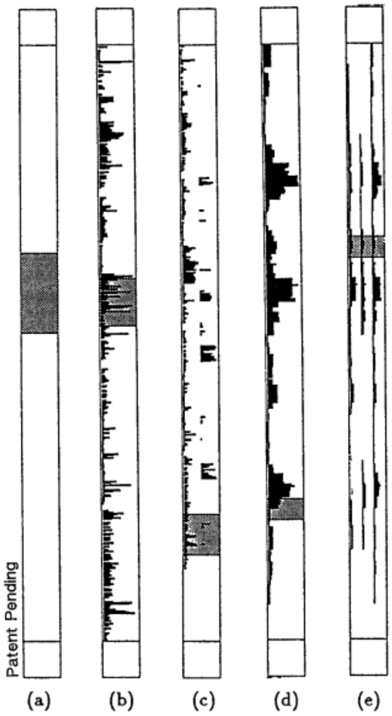

Figure 2.8. Hill and colleagues original scrollbar designs showing *wear* patterns of reading and writing. Reproduced from [70] with permission.

Social navigation has focused on guidance in information selection – through supporting decision making – and the notion has been applied widely in application areas such as recommender systems [88], collaborative groupware systems [50], online recipe exploration [73], reading newsgroup messages [126], browsing the Web [156], viewing online university lectures [13,98] or learning materials [13,26], and in virtual worlds [88]. In this diverse set of application types, implicit and explicit markers of past use are employed to help the current seeker make decisions. Further, many popular web-based systems use social navigation as a highly visible part of their systems and often as the primary or sole means of selecting or ranking information items (see Figure 2.9).

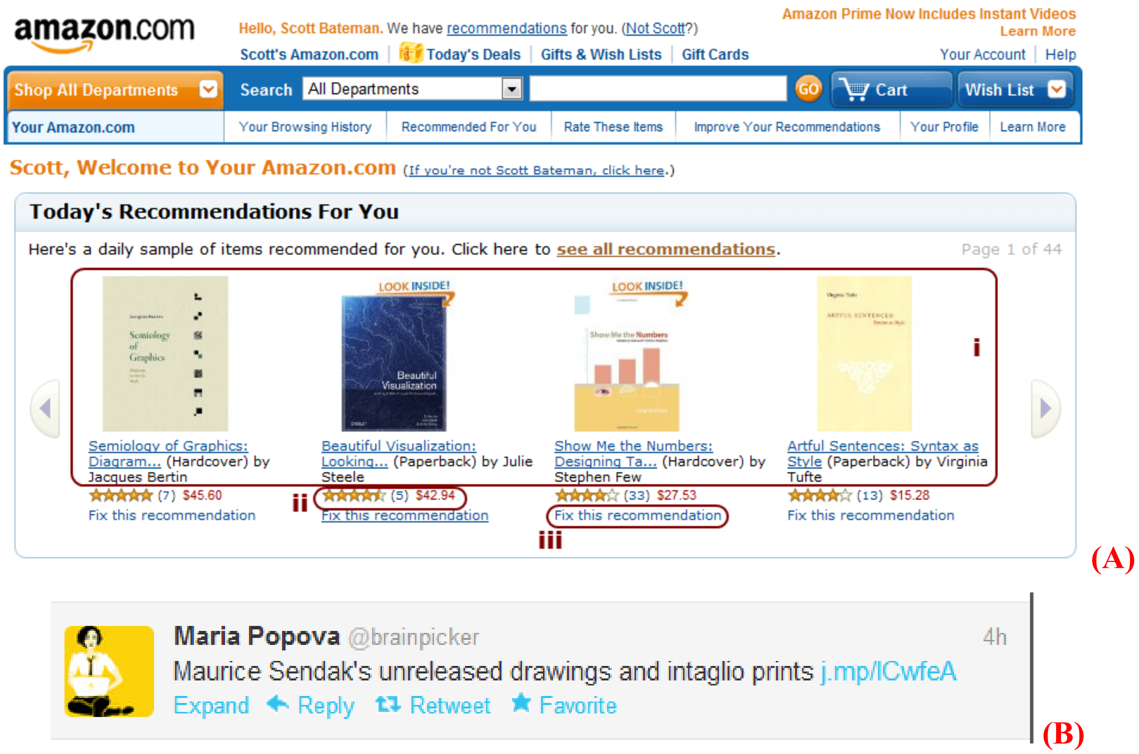

Hello, Scott Bateman. We have recommendations for you. (Not Scott?)

Scott, Welcome to Your Amazon.com (If you're not Scott Bateman, click here.)

Today's Recommendations For You

Here's a daily sample of items recommended for you. Click here to see all recommendations.

Semiology of Graphics: Diagram... (Hardcover) by Jacques Bertin
★★★★★ (7) $45.60
Fix this recommendation

Beautiful Visualization: Looking... (Paperback) by Julie Steele
★★★★★ (5) $42.94
Fix this recommendation

Show Me the Numbers: Designing Ta... (Hardcover) by Stephen Few
★★★★☆ (33) $27.53
Fix this recommendation

Artful Sentences: Syntax as Style (Paperback) by Virginia Tufte
★★★★☆ (13) $15.28
Fix this recommendation

Maria Popova @brainpicker
Maurice Sendak's unreleased drawings and intaglio prints j.mp/ICwfeA

Figure 2.9. Social Navigation in different online systems. (A) Amazon.com (i) uses interaction history to provide recommendations, (ii) displays aggregations of user created ratings, and (iii) allows users to correct inferred recommendations with explicit feedback. (B) Twitter users often use the surface for direct social navigation, by sharing links they find interesting.

2.6.2.1 Components of Social Navigation Systems

Goecks highlights that social navigation systems can be characterized as having three basic components: input composed of community data; an aggregation algorithm; and output [63]. He defines a *community* to be the set of users who contribute data to the system, and *community data* as the aggregation of data that is used in information item presentation. Output can take many forms such as list ranking, item filtering and highlighting items to draw user attention [63].

Input: user data

Goecks characterizes typical types of community data in terms of how they are collected (implicit or explicitly), the severity of burden on the user in contributing data (low, medium or high), the difficulty in aggregating the data (easy, moderate or hard), and how expressive the resulting output is (see Table 2.5). Please note that how data is collected, implicitly (from collection of interaction data) or explicitly (e.g., user input such as ratings) should not be confused with the form of social navigation, direct (advice received from some form of communication) or indirect (inferred advice from implicit data).

Table 2.5. Goecks' characterization of different sources of community data used in social navigation. Recreated from [63].

	Date Collection	User Burden	Aggregation	Expressiveness
Activity Data	Implicit	Low	Easy	Low
Ratings	Explicit	Moderate	Easy	Moderate
Free Text	Explicit	Moderate	Moderate	High
Tagging	Explicit	Moderate	Moderate	Moderate-High

Goecks considers activity data (or interaction history) as the only truly implicitly collected data. Implicitly collected data is characterised by the fact that it can require no extra effort beyond typical system use. This means that the amount of data collected can be higher than through explicit methods, but the trade-off is that the insights yielded can be more limited. Other classically used data sources include ratings, free text (e.g., reviews), or tagging (keywords used to describe and/or categorize information items). Consider writing a free-text description on an experience with a hotel. This activity would be considerably more difficult and time-consuming for a user than simply visiting the page or selecting the appropriate star on five-point rating scale. Free-text, though, possibly provides the richest source of information. In free text all questions relating to who, what, where, when and how can be answered, if the user takes time to provide this information. However, it is difficult to accurately extract this information from multiple reviews and aggregate them in a meaningful and consistent way [63].

Aggregation algorithms: creating community data

Most social navigation systems do not display or employ input data on a per-user basis. Aggregation algorithms are used to select and weight inputs appropriately to create community data. Goecks highlights that most aggregation algorithms perform the following steps [*63*]:

1. identify a group of users who use the system;
2. weight selected users based on some heuristic;
3. average the weights of selected users.

Algorithms vary either in the way users are selected or by weighting users differently. At one end of the spectrum, along these two varying dimensions, an all-inclusive equal-weighting approach is an extremely common and simple approach, which include all user data with equal weights. At the other end of the spectrum are collaborative filtering systems, which often select a set of users most like the current user, and then weight data according to the calculated level of similarity [63].

Output.

There has been relatively little research designing and comparing the effectiveness of different social navigation user interfaces [*63*]. Systems and studies that have investigated the use of providing visual cues regarding the activity of others are described in Sections *2.6.2.3* and *2.6.2.4*.

2.6.2.2 Challenges: Herding, Cold Start, Interest Drift[7]

Research into recommender systems has suggested that the presentation of social navigation information does have an influence on users [40,63]. However, this may not always be a desirable thing, as some work has suggested that this leads to *herding* behaviour. Herding occurs when users are unduly influenced by social navigation information and end up making a poor decision based on this information [40,63]. Herding can lead to decisions that the users will regret later on, or provide little value, or be incorrect [63,141]. For example, Svennson and colleagues call this the *snowball effect*; "… watch out for the snowball effect where the social trails lead more and more users down a path they do not perceive [as] valuable in the long run [142]."

An inherent problem with social navigation systems is that opinions and interests change over time, which is called *interest drift* [141]. The challenge for social navigation systems, then, is not to provide guidance down stale paths; what is of interest today may not be in a month's time. This can be partially accounted for by giving higher weights to actions that are more recent in a system or by making clear the timeframe of presented information.

The cold start problem is a well-known issue with systems that rely on user-generated data. When the first users come to a social navigation system, there is no information available

---

[7] The challenges section is after Farzan ([55], p. 24-25).

to provide guidance. Because of this, initial users can have a lot of influence, and if they make poor choices it can set up a situation for a snowball effect to occur. The most common method to deal with this is to combine semantic navigation along with social navigation, giving users an additional mechanism upon which to base decisions [131].

2.6.2.3 Broad Social Navigation Tools

In this section, I describe work in creating implicit trace-based social navigation tools. In general, *broad social navigation tools* have focussed on graphically representing the most visited webpages and links, from collected interaction history, and typically aggregate and present the behaviour of all users together. For example, the work of Wexelblat and Maes created tools for navigating the MIT Media Lab website, which were based on common real-world navigation tools: maps, paths and signposts (see Figure 2.10) [156]. The tools were displayed in separate windows around a webpage, and allowed users to navigate webpage links based on explicit markers (signposts) created by other users, or by viewing how other users moved between pages (using paths and maps). In a study, they found that users could find destinations pages equally well with or without the tools. However, with the social navigation tools they needed significantly fewer steps to reach destination pages.

Willet et al., created “scented widgets” to improve exploratory data analysis [167]. The scented widgets extended the typical functionality of well-known widgets (e.g., radio button, sliders, checkboxes, or list boxes) with tightly integrated visualizations of community activity (see Figure 2.11 and Figure 2.12). They conducted an evaluation where users were instructed to explore a visualization system in two conditions: a version that contained typical non-scented widgets; and a version that contained scented widgets from pre-seeded social navigation information. Participants were instructed to find as many unique pieces of evidence, in support of several hypotheses, as possible. They found that using social navigation cues in their widgets led users to explore twice as many unique views in the system. Users preferred the scented widgets to typical widgets and found that they did not overly clutter the interface.

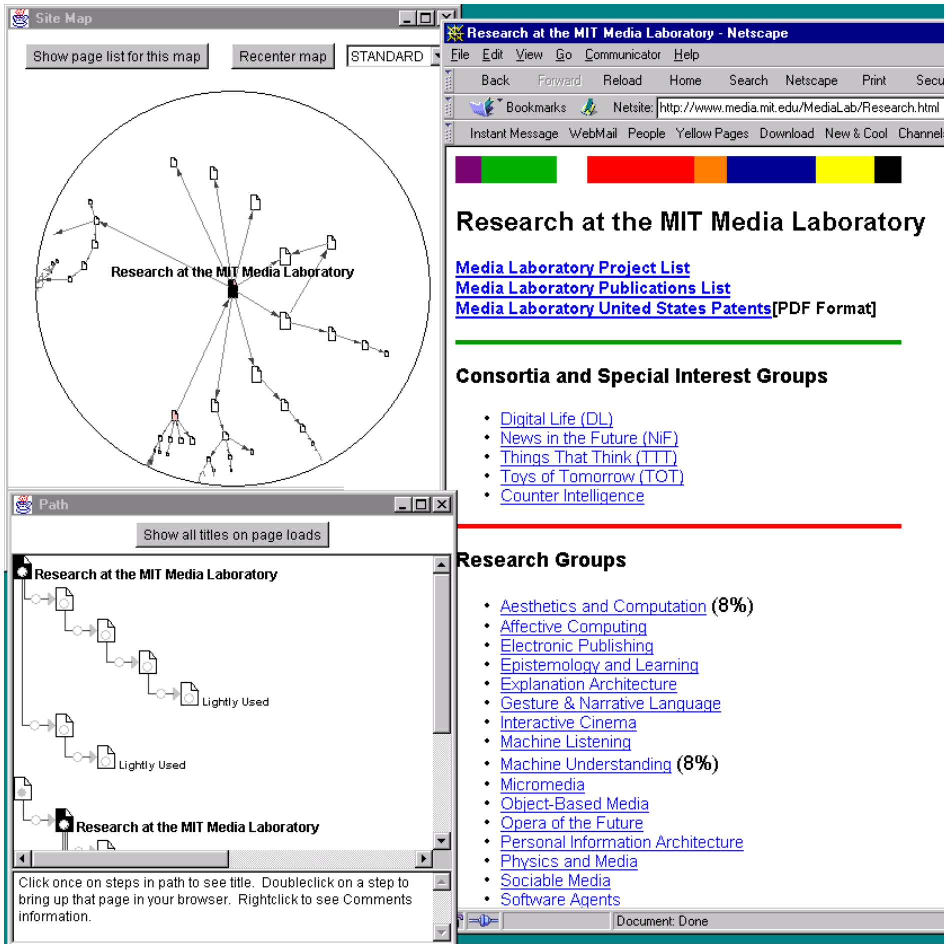
Site Map
Show page list for this map
Recenter map
STANDARD
Research at the MIT Media Laboratory
Path
Show all titles on page loads
Research at the MIT Media Laboratory
Lightly Used
Lightly Used
Research at the MIT Media Laboratory
Research at the MIT Media Laboratory - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Secu
Bookmarks Netsite: http://www.media.mit.edu/MediaLab/Research.html
Instant Message WebMail People Yellow Pages Download New & Cool Channel
Document: Done

Click once on steps in path to see title. Doubleclick on a step to bring up that page in your browser. Rightclick to see Comments information.

Research at the MIT Media Laboratory

Media Laboratory Project List
Media Laboratory Publications List
Media Laboratory United States Patents[PDF Format]

Consortia and Special Interest Groups

- Digital Life (DL)
- News in the Future (NiF)
- Things That Think (TTT)
- Toys of Tomorrow (TOT)
- Counter Intelligence

Research Groups

- Aesthetics and Computation (8%)
- Affective Computing
- Electronic Publishing
- Epistemology and Learning
- Explanation Architecture
- Gesture & Narrative Language
- Interactive Cinema
- Machine Listening
- Machine Understanding (8%)
- Micromedia
- Object-Based Media
- Opera of the Future
- Personal Information Architecture
- Physics and Media
- Sociable Media
- Software Agents

Figure 2.10. The Footprints System displaying its tools: the browser (on the right) is displaying the Media Lab website; the map tool shows all pages and how they are connected (top left); the path tool shows different paths starting from a page with information describing relative use (lower left). Reproduced with permission,from [156].

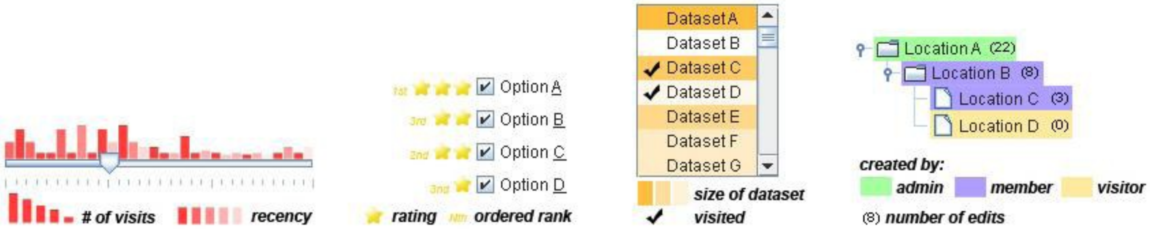

Figure 2.11. Alternative designs for scented widgets.
From [167], used with permission. © 2007 IEEE

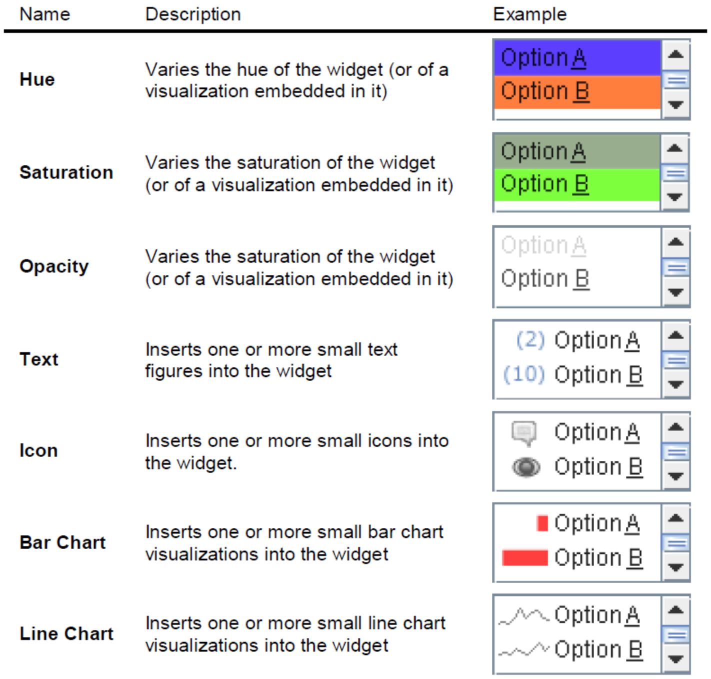

Name	Description	Example
Hue	Varies the hue of the widget (or of a visualization embedded in it)	Option A Option B
Saturation	Varies the saturation of the widget (or of a visualization embedded in it)	Option A Option B
Opacity	Varies the saturation of the widget (or of a visualization embedded in it)	Option A Option B
Text	Inserts one or more small text figures into the widget	(2) Option A (10) Option B
Icon	Inserts one or more small icons into the widget.	Option A Option B
Bar Chart	Inserts one or more small bar chart visualizations into the widget	Option A Option B
Line Chart	Inserts one or more small line chart visualizations into the widget	Option A Option B

Figure 2.12. Different visual mapping possibilities for scented widgets.
From [167], used with permission. © 2007 IEEE

Broad social navigation tools are often limited because they typically aggregate the behaviour of all users [63], limiting what can be inferred from the traces and what questions can be answered. For example, a system showing the navigation behaviour of a large group can indicate general visitation patterns and popularity (like in the Footprints System), but are less useful for specific questions relating to awareness (relating to *who* and *when*). This can lead them be subject to the challenges of social navigation systems (as described above), because traces lack important contextual details. For example, a visual trace may indicate that a particular news story in the most popular on a technology website: “Apple’s New iPad”. However, after following the link an information seeker sees that the story is about the release of the original iPad, something that they are not interested in because it is old news. In this case, the social navigation system was subject to interest drift. However, had the system provided an indication of how old the traces were on the iPad story link, then the seeker could have realized that the article was old news. Further, there have been no studies of broad social navigation tools on the open Web, where they can provide support throughout the information seeking process.

2.6.2.4 Narrow Social Navigation Tools

Narrow social navigation systems provide traces of user interactions, in much the same way as broad social navigation systems do. However, more details of the other seekers' interaction history are made visible, including who performed an action and when it occurred. These details can allow an information seeker to better determine if the presented information is of interest or not.

Diebergger and Lonnqvist created CoWeb [45] – an extension to a wiki system –that provided visual cues of what other users have been doing (edits and visits) displayed beside links (see Figure 2.13). The visual cues were small icons, which used color to describe the user activity on the linked page. "New" icons represented the age of the page (red = less than a day old, yellow = less than 3 days old, or blue = less than a week old). Footprint icons represent the relative amount of visitation to a page (red = lots of visitation, yellow = moderate visitation, blue = little visitation). They also provided a global list of what other users have been doing –users could visit a central page to view the activity of all users. They inferred from the results of a user study that the visualized interaction history cues led users to participate more (i.e., users visited and edited more pages than they otherwise would have), and they saw some importance in providing both global (see Figure 2.13) and contextual cues (cues that appear in the location of activity, e.g. annotations made on hyperlinks representing visits to the linked page).

Narrow social navigation systems are in active use on several social networking websites – such as Facebook (see Figure 2.14) or Academia.edu – and on search engines – such as Bing and Google (see Figure 2.15). These systems provide indirect social navigation from the traces of a seeker's social networking contacts. For example, Facebook displays contacts' reading behaviour (Figure 2.14), and Google displays what contacts have shared using Google+ when viewing search results (see Figure 2.15).

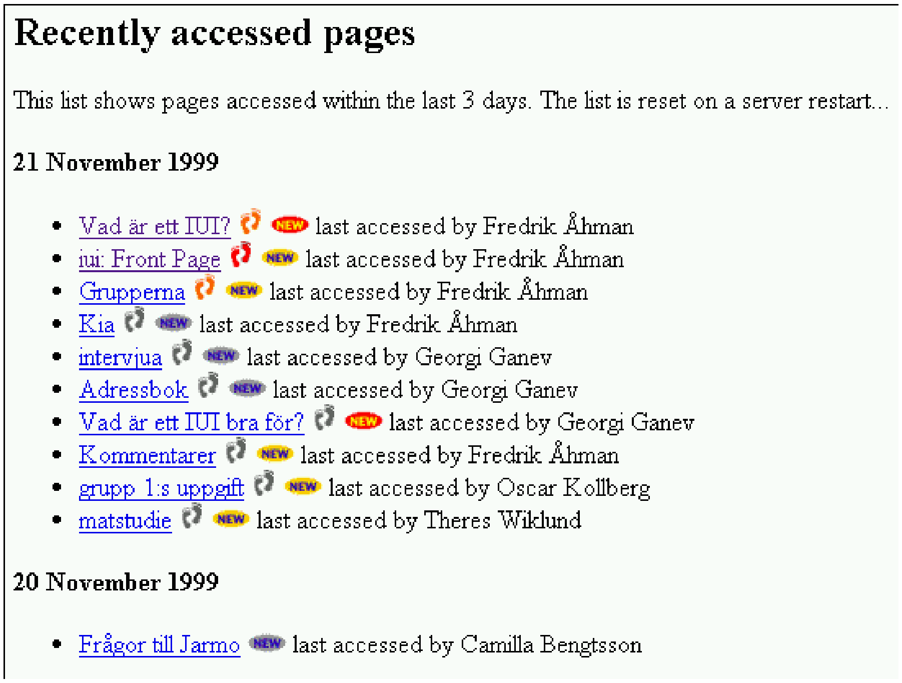

Recently accessed pages

This list shows pages accessed within the last 3 days. The list is reset on a server restart...

21 November 1999

- Vad är ett IUI? 👣 NEW last accessed by Fredrik Åhman
- iui: Front Page 👣 NEW last accessed by Fredrik Åhman
- Gruppema 👣 NEW last accessed by Fredrik Åhman
- Kia 👣 NEW last accessed by Fredrik Åhman
- intervjua 👣 NEW last accessed by Georgi Ganev
- Adressbok 👣 NEW last accessed by Georgi Ganev
- Vad är ett IUI bra för? 👣 NEW last accessed by Georgi Ganev
- Kommentarer 👣 NEW last accessed by Fredrik Åhman
- grupp 1:s uppgift 👣 NEW last accessed by Oscar Kollberg
- matstudie 👣 NEW last accessed by Theres Wiklund

20 November 1999

- Frågor till Jarmo NEW last accessed by Camilla Bengtsson

Figure 2.13. A global list of accesses in Dieberggerʼs CoWeb system. The visual cues used were small icons, which used color to describe the user activity on the linked page. “New” icons represented the age of the page (red = less than a day old, yellow = less than 3 days old, or blue = less than a week old). Footprint icons represent the relative amount of visitation to a page (red = lots of visitation, yellow = moderate visitation, blue = little visitation). From [46].

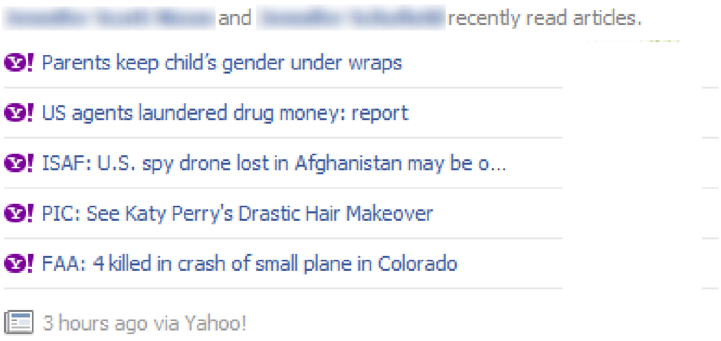

[redacted] and [redacted] recently read articles.

Y! Parents keep child's gender under wraps

Y! US agents laundered drug money: report

Y! ISAF: U.S. spy drone lost in Afghanistan may be o...

Y! PIC: See Katy Perry's Drastic Hair Makeover

Y! FAA: 4 killed in crash of small plane in Colorado

3 hours ago via Yahoo!

Figure 2.14. Indicators of interaction history from two of the authorʼs friends on Facebook. Friendʼs names have been blurred.

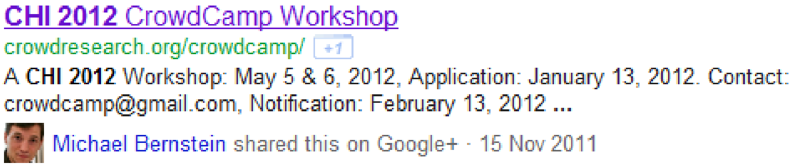

CHI 2012 CrowdCamp Workshop

crowdresearch.org/crowdcamp/ +1

A **CHI 2012** Workshop: May 5 & 6, 2012, Application: January 13, 2012. Contact: crowdcamp@gmail.com, Notification: February 13, 2012 ...

Michael Bernstein shared this on Google+ · 15 Nov 2011

Figure 2.15 A search on Google for CHI 2012 on Google, shows a social navigation indicator. Because a friend, Michael, is also one of the author's contacts on Google+, Google augmented the search result snippet with an indicator that says Michael shared a workshop, an implicit recommendation to follow the link.

The systems described above provide narrow social navigation support – by including details of *who* and *when.* These details could allow an information seeker to better understand if presented information is useful to their current task. However, in general, they are limited in scope, because they work only for certain types of behaviour and on certain websites. For example, Facebook only displays friends' article-reading history from particular sites (like *Yahoo! News* and the *New York Times*). This means the scope of the displayed activity does not extend to most information-seeking tasks, which can occur over many different websites and different activities. Whether such narrow systems that have a limited scope can provide support in typical information-seeking tasks is unclear.

2.6.3 Collaborative Information Seeking

People work together when seeking information [34,66,104]. Even if search system use is solitary, the overlying task can involve others [34,159]. However, most information seeking systems do not support collaboration at any stage of search [66,104]. This leads people to adopt workarounds [66], such as by sending URLs in emails, or by one searcher adopting the role of "driver" while others observe from the "backseat" [105]. Collaborative search has been defined as a situation where users have "explicitly shared information needs ([66], p.47)" and work together to fulfill those needs. Research has been carried out to address different conditions where collaboration information seeking may occur. This work has addressed several related themes that describe the collaborative setting [66,104]:

- *location*: collocated versus remote;
- *time*: synchronous versus asynchronous;
- *role*: symmetric versus asymmetric.

However, in addition to these characteristics are those that bridge two different classes of systems – systems that coordinate user activities (by inferring how users can help one another) and those that allow users to collaborate. Such characteristics include:

- *depth of mediation*: system mediated versus collaborator mediated;
- *group formation*: explicit versus implicit;
- *intent*: explicit versus implicit.

Further, in some of the systems that will be reviewed users can work towards a shared goal almost completely independently of one another, yet there is still a shared task. This implies another theme:

- *coupling*: loose (largely independent work) versus tight (work is dependent on others).

2.6.3.1 Systems and Studies of Collaborative Information Seeking

As described in the introduction, this section will introduce work along several themes that characterize collaborative information seeking: location, time, role, group formation, intent, and coupling. I further describe the overlap of the previously described social navigation systems where systems coordinate users but do not necessarily allow them to collaborate. Concepts that span collaboration and coordination are the depth of mediation, intent, group formation, and planning. Each characteristic is briefly described and, where possible, examples of systems that demonstrate the characteristic are provided. Systems can often fit into more than one category, but will only be presented once. Further, while there may be many systems that display a given characteristic or have been focused on collaborative web search (see [104] for a comprehensive review), I have selected recent systems and studies that clearly demonstrate the characteristic.

Location: Collocated versus Remote.

Collaborators can be physically proximal or separated by distance. Accordingly, there have been two classes of systems that have emerged. I will describe systems that are intended for collocated use here, as the systems that follow are all for remote collaboration. There is nothing that necessarily precludes a remote system from being used in a collocated situation; rather systems that work at a distance may be limited to some forms of interaction, and the need for additional communication channels is implied [*65*]. Unlike remote systems, most systems designed for collocated search also entail the requirement for collaborators to work at the same time (as described below in Time: Synchronous versus Asynchronous). The systems described below demonstrate interaction possibilities enabled by collocation. An example of a system designed for collocated search is Fischlar-DT (see Figure 2.16) allows users to search a video archive on a touch tabletop display. Users can type keywords on personal virtual keyboards. They can freely arrange video surrogates (keyframes) using direct touch [136].

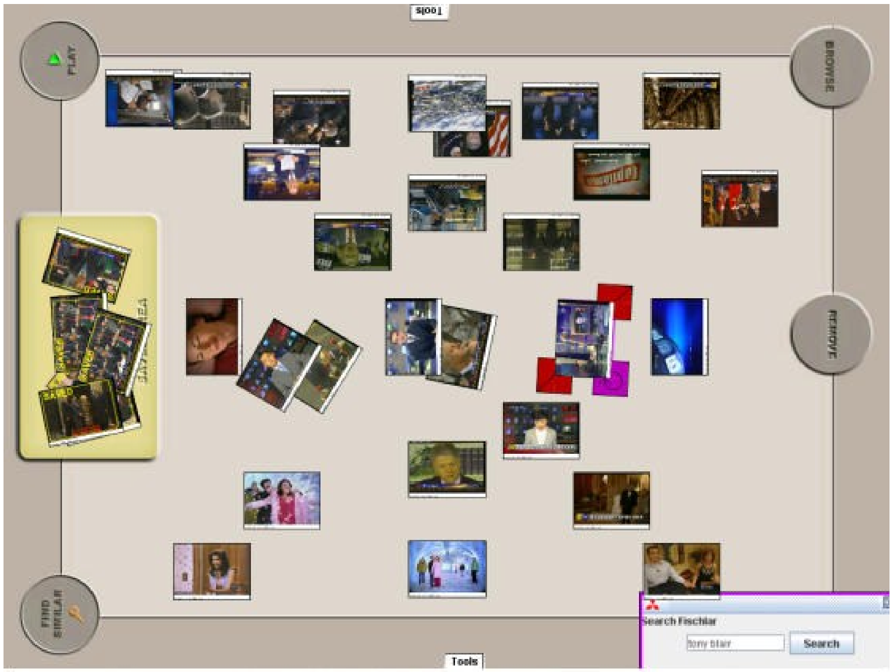

Figure 2.16. Fishlar-DT allows collocated searchers to find and organize videos in a collection. Reproduced with permission, from [136].

Time: Synchronous versus Asynchronous.

Collaborative activities can occur synchronously – where collaborators work at the same time to achieve some information seeking goal – or asynchronously – with collaborators completing parts of a shared task at different times. Systems that are designed to support asynchronous work do so by creating a persistent shared representation of search activities [*104*]. In asynchronous situations, previously performed searches are shared to influence later searches [66]. Recent work on synchronous search has focused on providing awareness of what other collaborators are doing [*104*], where the focus is on being able to provide an influence in real time [66]. The S^3 System (Storable, Shareable Search) was a system created to support asynchronous search [106]. As a user searches, search activities are stored, including: queries made, the result lists viewed, and the pages visited. These activities are categorized automatically (by query) in a web browser sidebar (see Figure 2.17).

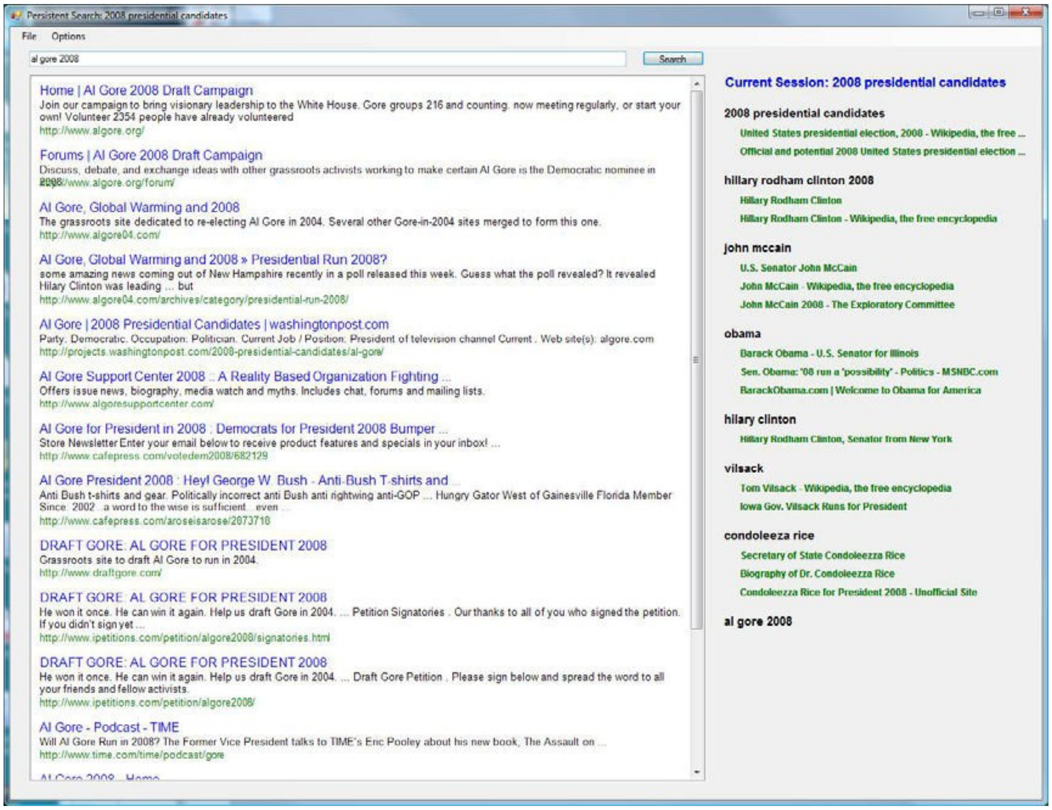

Home | Al Gore 2008 Draft Campaign
Join our campaign to bring visionary leadership to the White House. Gore groups 216 and counting. now meeting regularly, or start your own! Volunteer 2354 people have already volunteered
http://www.algore.org/

Forums | Al Gore 2008 Draft Campaign
Discuss, debate, and exchange ideas with other grassroots activists working to make certain Al Gore is the Democratic nominee in 2008
http://www.algore.org/forum/

Al Gore, Global Warming and 2008
The grassroots site dedicated to re-electing Al Gore in 2004. Several other Gore-in-2004 sites merged to form this one.
http://www.algore04.com/

Al Gore, Global Warming and 2008 » Presidential Run 2008?
some amazing news coming out of New Hampshire recently in a poll released this week. Guess what the poll revealed? It revealed Hilary Clinton was leading ... but
http://www.algore04.com/archives/category/presidential-run-2008/

Al Gore | 2008 Presidential Candidates | washingtonpost.com
Party. Democratic. Occupation: Politician. Current Job / Position: President of television channel Current . Web site(s): algore.com
http://projects.washingtonpost.com/2008-presidential-candidates/al-gore/

Al Gore Support Center 2008 :: A Reality Based Organization Fighting ...
Offers issue news, biography, media watch and myths. Includes chat, forums and mailing lists.
http://www.algoresupportcenter.com/

Al Gore for President in 2008 : Democrats for President 2008 Bumper ...
Store Newsletter Enter your email below to receive product features and specials in your inbox! ...
http://www.cafepress.com/votedem2008/682129

Al Gore President 2008 : Hey! George W. Bush - Anti-Bush T-shirts and ...
Anti Bush t-shirts and gear. Politically incorrect anti Bush anti rightwing anti-GOP ... Hungry Gator West of Gainesville Florida Member Since: 2002 .. a word to the wise is sufficient ... even ...
http://www.cafepress.com/aroseisarose/2873718

DRAFT GORE: AL GORE FOR PRESIDENT 2008
Grassroots site to draft Al Gore to run in 2004.
http://www.draftgore.com/

DRAFT GORE: AL GORE FOR PRESIDENT 2008
He won it once. He can win it again. Help us draft Gore in 2004. ... Petition Signatories . Our thanks to all of you who signed the petition. If you didn't sign yet ...
http://www.ipetitions.com/petition/algore2008/signatories.html

DRAFT GORE: AL GORE FOR PRESIDENT 2008
He won it once. He can win it again. Help us draft Gore in 2004. ... Draft Gore Petition . Please sign below and spread the word to all your friends and fellow activists.
http://www.ipetitions.com/petition/algore2008/

Al Gore - Podcast - TIME
Will Al Gore Run in 2008? The Former Vice President talks to TIME's Eric Pooley about his new book, The Assault on ...
http://www.time.com/time/podcast/gore

Current Session: 2008 presidential candidates

2008 presidential candidates
- United States presidential election, 2008 - Wikipedia, the free ...
- Official and potential 2008 United States presidential election ...

hillary rodham clinton 2008
- Hillary Rodham Clinton
- Hillary Rodham Clinton - Wikipedia, the free encyclopedia

john mccain
- U.S. Senator John McCain
- John McCain - Wikipedia, the free encyclopedia
- John McCain 2008 - The Exploratory Committee

obama
- Barack Obama - U.S. Senator for Illinois
- Sen. Obama: '08 run a 'possibility' - Politics - MSNBC.com
- BarackObama.com | Welcome to Obama for America

hilary clinton
- Hillary Rodham Clinton, Senator from New York

vilsack
- Tom Vilsack - Wikipedia, the free encyclopedia
- Iowa Gov. Vilsack Runs for President

condoleeza rice
- Secretary of State Condoleezza Rice
- Biography of Dr. Condoleezza Rice
- Condoleezza Rice for President 2008 - Unofficial Site

al gore 2008

Figure 2.17. The search interface for the S^3 System. Search dialog is at the top where users can issue queries. The lower left-hand frame displays a web browser, and the right-hand sidebar displays the summary of the search session. Reproduced with permission, from [106].

Role: Symmetric versus Asymmetric.

When collaborating, users can take on different roles in order to complete their shared task. These can be classified as symmetric or asymmetric roles. In symmetric collaboration searchers take on equivalent roles. Asymmetric collaboration refers to the case where searchers take on different roles. Asymmetric roles may be used by searchers to most efficiently complete a task and can be based on familiarity with search technology, a hierarchy in jobs, or some specific expertise [65,104]. Shoulder surfing is another example where users have asymmetric roles. The search technologies described so far can support either symmetric or asymmetric use, and it would be up to the group to decide what roles should be taken on. However, very few systems are designed specifically for supporting asymmetric roles. One system that does provide asymmetric roles for search is Cerchaimo [119]. Cerchiamo is a system for two users to search together for online video clips. One user acts as the prospector, whose job it is who creates new queries, trying to find new avenues of exploration. The successfully identified queries created by the prospector are collected for the other user – the miner. The miner explores the collected queries in depth to decide which videos are worth keeping. Each user has their own task specific interface, as well as shared interface that allows them to coordinate.

Users who work in symmetric roles can be considered to be working as peers [65]. However, work may be further divided by particular parts of a higher level information-seeking task. One approach is to use a *divide-and-conquer* strategy, where users coordinate what parts of the task they will complete. If no such coordination is provided or planned, they are said to be using a *brute force* strategy [105]; work may be repeated by multiple people.

Coordination versus Collaboration.

Golvochinsky and colleagues widen the scope of what might be considered collaboration by characterizing intent, group formation, and the depth of mediation [*65*,*66*]. However, as will be discussed, these characteristics do not necessarily describe situations where users are collaborating; users are not actively working together on a shared goal or information need. Golvochinsky et al. Admit: “In some sense this is not strictly collaboration, but rather a coordination of people’s activities ([65], p.1).”

Intent and Group Formation.

When intent is explicit in information seeking, a group of users work on a declared information need [65,66]. The explicitly identifiable groups may have been formed by self-identifying shared needs, through social relations, or through work requirements. Implicit information seeking systems instead infer information needs and the intent of searches based on user actions or opinions. Similarly, users are “grouped” based on their behaviour and their preferences. The previously described collaborative filtering systems and search engines which filter or order items in a larger collection based on an inferred preference are examples of both implicit intent and implicit group formation.

Depth of Mediation.

A different aspect of collaborative systems is where user coordination occurs. For example, Amazon makes recommendations based on inferring preference from user-created ratings and histories of items bought. This is an example of *system mediated* collaboration, where some system controls the flow of information between users. Typically, systems for collaboration, as the ones discussed here, allow users to coordinate and engage one another through the interface. These systems are referred to as being *interface mediated* [66].

Coupling: Loose versus Tight.

As previously described, coupling describes the amount of work that people can do individually before they require communication or actions from another person [67]. The awareness that collaborators have of one another facilitates moving from loose to tight coupling in synchronous scenarios. However, in asynchronous scenarios, work tends to be more loosely coupled because of the overhead and time requirements associated with communication at a distance. Awareness of the past can reduce or relieve the need to enter into tightly coupled collaboration, and allow collaborators to continue to work independently. For example, two students, Larry and Joe, are collaborating on a long essay by passing the Word file back and forth through email. Joe receives the latest updates from Larry, but isn't sure what Joe has done. Larry hadn't been using the track change features in Word, and he hasn't left any comments in the document. Joe decides he had better call Larry on the phone, so they can discuss what has been done to the document and what needs to be done next. During the discussion, Joe must ask Larry several times to slow down while he takes notes so that he can keep track of the details. In this example, had Larry been using the track change features, Joe might have been clear on what had been changed, relieving him of the need to enter into a tightly coupled collaboration with Larry over the phone. Shah and Marchionini studied different levels of awareness in collaborative information seeking and found that the appropriate level of awareness reduced the amount of communication that participants needed to successfully complete their tasks [134].

2.6.4 Privacy Concerns in Awareness Systems on the Web

The systems that have been described in this section (social navigation and collaborative information seeking systems) can be considered as awareness systems; they provide information that facilitates different forms of collaboration through being aware of other people's actions. Research into awareness systems has noted that too much awareness can be a bad thing because it raises privacy concerns [74]. In the physical world it is often immediately clear to people what is private and what is public, and people are able to adjust their behaviour accordingly to protect their own privacy and to respect the privacy of others [74]. However, in the digital world we often lack the cues that are needed to distinguish between public and private spaces. This means there is often a tradeoff between providing awareness information and protecting privacy. On the one hand, providing a too little awareness information preserves privacy, but does not facilitate collaboration. On the other hand, providing too much awareness information can compromise privacy, but improves collaboration [21,74].

Privacy was described as a control process by Altman [3], where people must manage the flow of information towards and away from one's self. This leads to two forms of privacy: confidentiality and solitude. Confidentiality describes the information that moves away from a person, and controls the personal information that other people can have access to. Solitude is control over the information moving towards a person, and determines how the information requires someone's attention [3].

For a person to feel that their privacy has been protected requires that they are comfortable with the information that a system has about them and their activities. In particular, people must feel they have control over the flow of information [21] and not feel like they are under surveillance [78] – that is, when information is used for purposes other than the original intention of sharing. The importance of user control has also been noted in different types of activity logging (e.g., logging of mobile phone usage [77]).

Managing and controlling privacy can be difficult because it is a dynamic process that is under continuous refinement, with the boundaries between public and private changing based on contextual factors [118]. However, there are still few good lightweight mechanisms that allow finely-grained control over the distribution of awareness information. Most solutions require too much user effort and only allow coarse-grained schemes for distributing information [21]. Because the cost of effort for users to maintain their privacy is too high, it can prevent people from sharing awareness information (in order to avoid privacy concerns). The tension between user effort and the need for fine-grained contextual control has led to the investigation of different factors that mediate privacy concerns. The intention has been to identify contextual factors that could allow awareness systems to control the distribution of awareness information in an appropriate and natural fashion [41, 115]. These factors have included relationship with the person receiving the information [41, 115], the type of information being distributed [24,41,115], and the physical location of where the information will be displayed [24].

There is some research that has investigated the use of web activity in awareness systems and what privacy concerns this might raise. It is known that most search engine users believe that having their search history tracked by search engines is an undesirable thing [125]. One survey of information sharing preferences within a large company has suggested that people may not want to broadcast the pages they are visiting to their entire company, and would prefer to share their web activity with smaller groups of close colleagues [24]. However, people are generally much more willing to share activity about the work-related websites they have visited than the non-work related websites. People treat the sharing of their work-related web activity similarly to other work related activities (e.g., sharing information about the work-related documents they have accessed, their availability, or their workplace calendar entries). Further, people have a high willingness to share their work-related website activity with trusted colleagues and managers [115]. Other research has shown that people are willing to make a large portion of their activity (42% on average) completely viewable to people who may incidentally see their screen, but people vary widely in the degree of webpage activity they are willing to share [69].

Overall, past research has shown that people are willing to share awareness information for a large portion of their web activity. Further, it seems people are more willing to share work-related web activity with work colleagues, but little is known about what other contextual factors may play a role in these situations.

2.6.5 Summary of Collaboration and Social Support in Information Seeking

2.6.5.1 Loosely-Coupled Collaboration

The majority of work interactions (81%) occur in an informal, unplanned, and opportunistic fashion [89]. Similar patterns have been observed in collocated information-seeking; e.g., in one study, 62% of collaborations occurred without pre-planning [4]. These kinds of collaborations often involve *loosely-coupled work* [89]. Coupling describes the amount of work that people can do individually before they require communication or actions from another person [68]. In loosely-coupled work, communication with collaborators must be low-cost, because the collaboration is often secondary to the underlying task; if the cost of communication is too high, people will simply not engage in it [4, 16,89,120].

It has been observed in loosely-coupled situations that traces of work left in the physical environment are a preferred channel of communication [16,120]. Traces in shared documents or work artifacts allow collaborators to communicate indirectly about the work they have undertaken [120], cost little to gather (since they are stored as a result of performing other primary tasks), and are easy for others to discover and retrieve [120]. Researchers have suggested that digital environments should, therefore, capture and display traces of work, since these are an important part of collaboration in physical environments [16].

2.6.5.2 Social Search and Collaborative Information Seeking

Social search encompasses any scenario where people work together to seek information on the Web [34,54,66,104]. Social interaction in web search is a common occurrence [34, 54,104]; for example, a recent survey found that 59% of search respondents shared the results of their most recent search [54]. Even when search is solitary, the underlying task can involve communication and collaboration with others [34,54]: before search – to help describe and focus what is sought; during search – to receive guidance and advice on avenues to pursue; and after search – to share what has been found. However, most existing search systems do not support any form of collaboration [66,95]. This leads people to adopt workarounds [66]; e.g., exchanging URLs by email [4].

Research in *collaborative information seeking* addresses situations where users have "explicitly-shared information needs" ([66], p.47), and actively work together to fulfill these needs. Work in this area has produced systems that support different forms of collaboration (see [104] for a review), including: collocated, synchronous distributed, and asynchronous distributed collaborations. This research has produced several successful systems (e.g, [106,107,134]); however, these systems typically focus on explicitly-shared information needs, meaning that they are oriented towards pre-planned work and more tightly-coupled interaction. As described above, however, most collaboration occurs in an unplanned way, and may be better supported by lower-cost solutions that fit into existing individual tools and practices.

2.6.5.3 Computational Wear and Social Navigation

Hill and colleagues [70] described the idea of *computational wear* – capturing "on computational objects the events that comprise their use" ([70], p.3). They developed read wear and edit wear tools that visualize traces of reading and editing history for people working on a shared document. Their goal was to allow people to easily answer awareness questions such as: Who wrote that part of the document? Or did my colleague edit that part? [70]

Computational wear shows that interaction traces provide implicit communication that can reduce the need for direct communication. Further, it provides a potential solution for loosely-coupled information seeking, because traces can be created at low cost and can be used in an informal fashion. However, there has been little work on how systems could be designed to enable appropriate sharing of traces in real-world web-based information seeking.

Social navigation describes how people can be guided through an information space based on the direct advice or interpreted actions of others [50]. The idea of indirect social navigation from interpreted behaviour (e.g., computational wear) led to Web systems such as Footprints [156] and CoWeb [45], or Scented Widgets [167] that show traces in a user interface. These systems provide visualizations of others' use of paths and links – they augment the information scent in the environment, providing information about social activity that allows users to make inferences about past user choices. Studies showed that these systems can reduce the navigation steps needed [156] and can increase the number of alternatives explored [167].

However, broad social navigation systems typically aggregate the behaviour of all users [63], limiting what can be inferred from the traces and what questions can be answered. For example, a system showing the navigation behaviour of a large group can indicate general visitation patterns and popularity [156], but is less useful for specific questions that are set into a local social context. While some other systems do offer narrow social navigation (including details of who and when), they typically only work for certain types of behaviour and on certain websites, which means their scope does extend to most information-seeking tasks.

CHAPTER 3

SOCIAL FEEDBACK: SOCIAL LEARNING FROM INTERACTION HISTORY

In many situations people conduct their information-seeking tasks individually. Asking for help from someone else in solo information seeking is discretionary (it does not always need to occur), and is based on a helper being ready and willing to help. These situations are best supported by loosely-coupled collaboration, where work occurs individually. In loosely-coupled collaborations people prefer using low-cost forms of communication, such as traces left in the environment; if communicating requires any more effort, then people will often choose not to engage in it, which means that work becomes less effective and efficient. Supporting loosely-coupled collaboration in information seeking activities has not been directly explored.

Enabling users to work together in information seeking has taken two general approaches in previous research: collaborative information seeking and social navigation. However, neither approach ideally lends itself to supporting some of the most common situations in which information seeking occurs, that is, loosely-coupled information-seeking scenarios.

On the one hand, collaborative information seeking allows users to work on "explicitly shared information needs ([66], p.47)" where a shared goal has been established beforehand. This means that collaborative information-seeking tools require explicit collaboration and direct communication before a task begins, and possibly during and after the task. However, information seekers most often work alone (even when a task is part of a collaboration), and when collaboration does arise, it is informal – meaning it is unplanned and at the discretion of the individual collaborators. Because of their more formal style of collaboration, traditional collaborative information seeking tools do not fit into these common situations.

On the other hand, social navigation systems are often either too broad or too limited in scope to support loosely-coupled information seeking. Typically, broad social navigation systems aggregate the information of all users, which means that critical details are hidden (including *who* performed an action and *when* the action occurred) (see Section 2.6.2.3). Broad social navigation tools only offer support in limited situations, because the missing details make it difficult for an information seeker to understand what support is being offered. While some other systems do offer narrow social navigation (including details of who and when) (see Section 2.6.2.4), they work only for certain types of behaviour and on certain websites. The scope of narrow social navigation tools does not extend to most information-seeking tasks, which can take place across many websites.

In this chapter I introduce the idea of *social feedback,* which supports loosely-coupled information seeking by allowing information seekers to learn from one another's efforts through shared traces of past activity. Social feedback is based on social learning theory, which describes how people learn from observing the behaviour of others. In social feedback, learning from observation is instantiated through the mechanism of *interaction history*, which are the traces of behaviour created as people interact with webpages during information-seeking tasks. Systems that provide social feedback allow people to conduct their information-seeking tasks as they normally would, but provide support for the essential parts of information seeking when needed. Social feedback can allow people to learn new skills and techniques to be used in their information-seeking practices, and can also allow people to learn how to get specific tasks accomplished more effectively and efficiently.

In the remainder of this chapter I describe social learning theory and briefly compare it to other related theories of learning. I then explain how interaction history is used for describing and creating representations of behaviour, and outline the important challenges of collecting interaction history on the Web. Finally, I describe how social learning can be enabled by using interaction history to create social feedback systems.

3.1 Social Learning Theory

The basic theory that explains how people learn through observation is called *social learning theory*, and was developed by Albert Bandura. Bandura [12] posits that a fundamental way that humans learn is from observing the behaviour of others. There are four main parts to the theory, as described below: the requirements for learning from others, the identity of a model, the basic forms of models, and the requirements for self-regulation.

3.1.1 The Requirements of Modelling (Learning from Observation)

Bandura refers to learning from observation as *modelling*, which involves four requirements.

1. ***Attention:*** people cannot learn from others unless they are paying attention while the learned behaviour is being demonstrated.

2. ***Retention:*** people must be able to effectively recall the components of the behaviour in order to effectively reproduce the behaviour.
3. ***Reproduction:*** people must be able to put the recalled components into action in order for the behaviour to be effectively reproduced. This implies that in order for behaviour to be both retained and reproduced, the observed behaviour must not be overly complex for the learner.
4. ***Motivation:*** the recall of the component actions and target behaviour will weaken unless the learner perceives some value in being able to reproduce the behaviour. Positive reinforcement can provide motivation, and can include:
   a. *Past reinforcement*: having had a positive experience for reproducing the behaviour.
   b. *Promised reinforcement*: the perception of receiving a positive experience for reproducing the behaviour.
   c. *Vicarious reinforcement*: seeing or remembering others having received a positive experience for producing the behaviour.

Together these four requirements are the core of social learning theory, and suggest that people can learn most effectively if they see and understand a demonstrated behaviour, and that the behaviour is not overly complex. However, a learner must also be motivated to learn the behaviour and they must see value in learning it.

3.1.2 The Identity of a Model

Bandura refers to the people we learn from through observation as *models*. The identity and characteristics of the model are essential to the modelling process, and influence both the viewer's attention to and motivation in learning a behaviour. Bandura identifies two ways in which a learner can more effectively learn based on the identity of the model, both of which lead the viewer to pay closer attention and have increased motivation to acquire the behaviour:

1. ***Having desirable or admirable characteristics*** – that is, if the model is attractive, appears to be competent, or seems prestigious to the viewer.

2. ***Social similarity or proximity*** – that is, if models can be self-identified with (i.e., the model seems to have characteristics of the viewer) or they are socially close (e.g., friends or parents).

Therefore, people are more likely to pay attention and value the demonstrated behaviour when they respect or admire a model, are closely connected to the model, or can see themselves in the model.

3.1.3 The Basic Forms of a Model

Traditionally, one might think of models as being an actual person. However, Bandura suggests that models for learning can take three basic forms:

1. ***A live model:*** a person who demonstrates or acts out behaviour.
2. ***A verbal description:*** descriptions and explanations of behaviour.
3. ***A symbolic model:*** real or created characters displaying the behaviour in a static form (e.g., in a book) or animated form (e.g., on film or a computer screen).

While Bandura does not directly take into account abstract or metaphorical visual representations, the notion that visual representations can be "read" and can convey precise meaning is widely recognized [166]. This suggests that as long as a behavioural representation can be understood and that it accurately conveys information, it should have the same effect as a verbal description that presents the same information.

3.1.4 The Requirements for Supporting Self-Regulation

Bandura also describes important steps for *self-regulation* – controlling one's own behaviour [11]. Once a behaviour is observed, the ability to adjust one's own behaviour may be improved – especially when a behaviour is complex or difficult to change – by following three steps:

1. ***Self-observation:*** first, we must have a realistic picture of our own behaviour.
2. ***Judgment:*** second, we need to be able to compare our own behaviour to a standard. A standard can be a personal goal or we can compare our behaviour (or compete) with others or ourselves. Standards must not be set too high, making them unrealistic, or set too low, making them meaningless.

3. ***Self-response:*** third, if we do well compared to our standard, we should reward ourselves. Rewards can range from extrinsic (collecting a material reward) to intrinsic (feelings of pride).

3.1.5 Other Theories Related to Social Learning

In describing persuasive technology (see Section 2.3.4.1), Fogg cites three theories that describe how social influence can persuade people to change their behaviour [60]. I summarize this review by briefly describing the relevant theories as they touch on how people can learn or be motivated by observing others. In all cases, these do not contradict what has been presented in social learning theory, but instead may work in concert with it.

Social Comparison Theory proposes that people naturally compare themselves with others in order to get a more realistic idea about their own attitudes and behaviour, and to provide a target for how behaviour might be changed [56]. This theory has implications for self-regulation of behaviour, where people need to have an accurate view of their own behaviour and need to judge themselves against a standard.

Reflection[8] has also been promoted as an essential part of learning [20,38]. Reflection is the process "… in which people recapture their experience, think about it, mull it over and evaluate it… working with experience [in this way] is important in learning [20, p.19]." This suggests that observing and reflecting on one's own behaviour can lead to important insights and learning.

Normative influence describes the pressure people perceive to change their own behaviour or beliefs, to match that of a particular group (such as family or colleagues) [60]. A normative influence may provide people with motivation to change, and can also have a role in the characteristics of an appropriate model.

Intrinsic motivation describes motivation that comes from enjoyment or interest in completing a task. Malone and Lepper describe three types of intrinsic motivation that come from other people and that can play a role in learning [93]. Many people are motivated by *competing* to outperform other people, while at the same time they can be motivated by

---

[8] Reflection was not considered by Fogg; I present it here for simplicity.

cooperating with others in their group. Finally, many people are motivated by *recognition*, which is having their work appreciated or recognized by other people.

3.1.6 Applying Social Learning to Information Seeking

In order for the lessons of social learning theory to be applied as support for information seeking on the Web, there needs to be a mechanism by which people can have access to the behaviour of others. Typically, social learning is studied and thought of in the case of live models. However, observing others as they work on the Web is not a common occurrence, and most often other people's guidance is not sought because we believe the potential costs to outweigh the benefits [19]. A mechanism to provide support would need to work within existing practices and meet the needs of loosely-coupled collaboration: allowing support to be available when needed, and allowing it to be provided and obtained with minimal effort.

3.2 Interaction History to Describe Behaviour

Interaction history (introduced in Section 2.4.1) consists of the traces of users' actions with digital objects such as webpages. Interaction history is most often collected without any explicit effort on the part of the user who creates it, because it results from actions that took place during some other task. Interaction history most often consists of time-stamped events that are gathered from the use of an interface or digital object. Interaction history can provide an account of the previous actions of a user, and allows events to be reviewed individually or to be replayed in sequence. Further, unlike social learning instantiated by a live model, interaction history (or "activity data" as shown in Table 2.5) can be aggregated over time periods and groups [63], which allow patterns of behaviour to be shown and compared.

Interaction history meets the basic requirements of loosely-coupled collaboration on the Web, by providing a low-cost method for the collection of users' experiences. Further, once interaction history is collected and stored, it can be presented at the discretion of a seeker that requires support. However, there are a few potential challenges that should be considered with any application of interaction history, including the loss of contextual information, the basic structure and dynamics of webpage content, and privacy concerns.

3.2.1 Capturing and using interaction history to describe behaviour[9]

The typical process of making use of interaction history to describe behaviour can be reduced to four main steps:

1. ***Instrumentation***: where logging capability is added to a system, such as a web application or a web browser;
2. ***System usage***: where the instrumented system is used in a real or simulated scenario and interaction history is collected;
3. ***Processing***: where the interaction history is optionally filtered, aggregated and transformed to create high-level descriptions of behaviour; and,
4. ***Presentation*****:** where the captured behaviour is displayed and observed to support some purpose.

First, a system should be *instrumented* (programming code added) to collect user interactions. Care should be taken during instrumentation, so that the collected interaction history enables behaviours to be accurately represented. In particular, the contextual details of collected interaction events must allow behaviour to be sufficiently described. For example, simply collecting an event that describes that "a link was clicked" carries very little information, and may not be useful. However, if instrumentation captures the appropriate contextual details of the event (e.g. "user U_1 clicked the link for page x while visiting page y at time t"), a richer picture of the original behaviour can be given. If the contextual details of interactions are not sufficiently logged, there is typically no way to recover the contextual information describing the original behaviour.

Second is *system usage*, where a system is used and interaction history is collected. System usage is typically done in realistic scenarios or in actual real-world use to provide accurate representations of user behaviour. Collected interaction history may be automatically transferred to a remote system or stored locally on the computer.

Third is *processing,* where the collected interaction history is optionally changed and combined to create higher-level representations of behaviour. Processing can involve three different activities: *filtering, transformation* and *aggregation*.

---

[9] This section builds on my work creating Interactive Usability Instrumentation [14] and the "Components of Social Navigation Systems" proposed by Goecks (see section 2.6.2.1).

In filtering, the interaction history is limited to a set of users, a set of particular interface actions, and possibly a particular timeframe. This allows only interface events that are necessary to creating a higher-level representation of behaviour to be considered.

In transformation, low-level interaction history events are combined to create higher-level descriptions of behaviour. For example, low-level logged interactions (such as "user U_1 loaded page x at time t") are often transformed into higher-level behaviour or action sequences. For example, a "webpage visit" could be created from two separate page load events. A webpage visit and its duration could be calculated as being the time between a user loading a page X at time t_1, and the user's next page visit, this time to page Y at time t_2. So, the resultant event might look like "user U_1 loaded page X at time t_1 and visited for $t_d = t_2 - t_1$".

In aggregation, the collected interaction history, that may have already been filtered and transformed, is combined to create a representation of higher level behaviour, over some set of users or interface actions, or some timeframe. When aggregating behaviour some function or algorithm is often used (such as taking the average value) to calculate a resulting value. For example, in the work that will be described in Chapter 4, I identified expert search engine users in interaction history logs and created a single "expert user" behaviour profile by taking the average behaviour of all of the identified search experts.

Fourth is *presentation*, where the behaviour is displayed for some purpose. The method of display should allow the key questions associated with the purpose to be answered as effectively and efficiently as possible. In the section on social feedback (see 3.3), I will describe how behaviour can be displayed to enable learning from social feedback. Other purposes include, the advice provided by social navigation systems (see Section 2.8.2.2) or usability evaluation (e.g., [14]).

3.2.2 Challenges to Using Interaction History on the Web

Below I discuss the two important challenges that could prevent the application of interaction history from web-based information seeking. These relate directly to *the privacy challenge* and *the overlap challenge* that will be examined in Chapter 6 and then discussed in Chapter 7.

First, there is the challenge associated with privacy. Because the Web now encompasses so much of our day-to-day activities, many people have serious privacy concerns about their interaction history being captured. People are now able to shop, bank, and form personal relationships online; and the activities that people wish to keep private and those that can be public are interwoven throughout someone's use of a web browser. In general, it is well understood that people like to stay in control of their logged behaviour [77]; however, interaction history is most often collected passively and is outside of a user's attention. While most systems allow collection to be turned on or off, people might see this as additional overhead and prefer not to have collection turned on in the first place. Further, most users are not aware of how to stop collection of their online activities, and many web users feel that the costs of interaction history collection outweigh the benefits [125]. However, in certain contexts people are willing to expend effort in order to share their otherwise personal activities (e.g., Facebook or foursquare). The situational characteristics in which people are more willing to share their activities are still not well understood.

Second, there is the challenge associated with the overlap of interaction history. The depth and breadth of the Web raises the issue of having sufficient interaction history in order for the behaviour that is described to be useful. Because there are so many individual pages on the Web, two users in a given group may never end up using the same page. This means that there may not be adequate interaction history collected to describe general user behaviour. Further, user behaviour can be idiosyncratic. Without sufficient overlap in usage, which can act as a filter for "noisy" behaviour, understanding interaction history may be difficult.

3.2.2.1 Further issues

I now describe two further issues that are not under direct study in this dissertation, but could pose potential challenges in the use of interaction history. Unlike the challenges above, these do not necessarily prevent the utility of a system that uses interaction history, but should be considered in its application. These additional issues will be discussed further in Chapter 7.

First, interaction history that is used to describe behaviour is subject to a *loss of context*. A user's interaction with an application is set within the specific details of a particular task and situation, which are often not represented by interaction history. For example: Why is the work being done? What is the goal of the work? Who else is working on this task? Answers to questions like these can be lost. In many situations, these details can provide important context to explain why a particular observed behaviour has occurred. Whether the loss of context would be a challenge for social feedback systems is unclear.

Second, there is the issue of *ephemerality.* The web content that people visit most frequently is ephemeral [1]. The content of a page itself could be, in whole or in part, dynamic. For example, news sites and weather forecasts are highly dynamic and often change from hour to hour, if not more frequently. This raises the question of "what does interaction history describe if the object of interaction has changed?" Whether someone would be confused or misled from viewing social feedback about a dynamic object is unclear.

3.3 Social Feedback: Designing for Social Learning from Interaction History

Social learning theory describes the fundamental requirements for learning from observation. It also describes who a learner might view as good model to learn from, and what form that observed behaviour can take. While interaction history has not directly been accounted for in this theory, I have argued that viewing a representation of behaviour (which can come from interaction history) can be little different from reading a verbal description of behaviour. Further, interaction history can be collected with little effort on the part of a provider, and it can be easily stored and presented when needed. Therefore, the use of interaction history fits well into loosely-coupled information seeking scenarios, because it provides a low-cost solution based on traces left in the environment, and can be used at the discretion of an information seeker.

Social feedback allows people to learn from observing the behaviour of others, behaviour which is instantiated using interaction history. Social feedback, then, can be defined as:

> *Support and guidance provided by learning from the behaviour of others, which is represented by interaction history.*

It should be noted that this definition of social feedback does not include information seeking. Social feedback can be generally applied because social learning theory and interaction history are concepts that are not specific to a particular application or topic. While this dissertation is focused on the use of social feedback for supporting loosely-coupled information-seeking tasks, social feedback can be applied more generally (several possibilities for its wider application outside of information seeking will be discussed in Section 8.2.3). Below I first present the basic design questions that must be addressed in designing and building social feedback systems for general applications. I then contextualize these design questions by considering their use in information seeking on the Web. I will address these questions again in

Chapters 4 and 5, where they are applied in focusing the design of two social feedback systems for supporting loosely-coupled information-seeking tasks. The four high-level questions (and their important component questions) that should be considered when designing social feedback systems are shown below (Table 3.1).

Table 3.1. The high-level design questions and important component questions that should be considered in designing social feedback systems.

The High-Level Design Questions for Building Social Feedback Systems

1. ***What are the tasks or activities to be supported?***
    - What is the basic behaviour used to complete these tasks?
    - Can task behaviour be captured and described through interaction history?
    - What type of learning best supports these tasks?
      – Global or Local Learning

2. ***What information do people need in the task?***
    - What are the important contextual details needed to describe behaviour?
    - Who will be the models for social learning?
      – Role Models or Tightly-Knit Contacts

3. ***How can the information be displayed?***
    - How can model behaviour fit into the task?
    - How can model behaviour be presented so that it is easy to learn?
    - Does there need to be support for self-observation or comparison?
    - Where will the motivation for learning come from?

4. ***How can potential challenges be contended with?***
    - How can the overlap challenge be contended with?
    - How can model privacy concerns be managed?

3.3.1 Design Questions for Social Feedback Systems

Below I describe the four high-level questions that should be considered when designing social feedback systems. In this section I refer to the *learner* as the person who is making use of a social feedback system to improve their task behaviour, and the *models* as the other people who are providing their interaction history to help the learner.

3.3.1.1 What are the tasks or activities to be supported?

First, a task should be identified where there is a real or perceived need for learning – that is, learners could potentially complete the task more successfully, by some measure of success, but lack the information that could improve their success. System designers should recall that learning does not just refer to acquiring new knowledge or changing behaviour, but can also include changes to skills, values, or preferences, and that learning in any of these ways could result in improvements to a learner's success. However, in social feedback the model must demonstrate their knowledge, skills, values or preferences through their behaviour as they interact with a system. This means that in order for social feedback to be useful, there needs to be enough information carried in the *behaviour* of a model that a learner can learn from their interactions with a system.

What is the basic behaviour used to complete these tasks?

Next, designers should consider what the basic behaviour is that people use to complete the task. This description of behaviour may not be as fine-grained as the level of individual interface interactions, but should instead constitute the basic steps carried out in completing the task. Specifically, the basic behaviour contains the points where a model makes decisions about how to act, and where the decisions that were made could affect task success. Identifying the basic behaviour is important because it determines the actions that can be observed by learning.

Can the task behaviour be captured and described through interaction history?

Once a candidate task has been identified and the basic behaviour used to complete the task is understood, designers must analyze the application and identify how the basic behaviour can be captured and described using interaction history. Designers may need to consider whether there are any challenges to the capture of the behaviour (such as those described in Section 3.2.2) that may prevent the capture of interactions that adequately describe behaviour (e.g., will models be willing to share their interaction history or can an application be instrumented appropriately).

What type of learning best supports the tasks: Global or Local Learning?

Designers must also decide what form of learning will best support the task. I identify two general types of learning: *global* and *local*[10]. *Global learning* refers to learning of new skills or techniques that can be applied in many situations. Examples of global learning are learning new mathematical skills and techniques or learning how to successfully assess the credibility of an information source (e.g., an article or a webpage).

Local learning refers to learning how to proceed within a specific situation, allowing a seeker to learn how to most effectively and efficiently accomplish a specific task based on previous actions. An example of local learning is having a friend advise you on what movie you should see or what novel you should read. It is important to note that through local learning, a learner may be able to generalize their experiences, allowing for global learning of skills or techniques that can be applied in many situations.

3.3.1.2 What information do people need in the task?

In order for learning to occur, a learner needs to see and understand appropriate details around a model's behaviour. A social feedback system must present appropriate contextual details of behaviour that allow the learner to infer its relevance, and to make judgments about its value.

What are the important contextual details needed to describe behaviour?

Contextual details are related to awareness (see Section 2.5.2) and allow a learner to answer questions relating to *what, who, when,* and *where* (examples can be seen in Table 3.2). Designers must ensure that the behaviour presented contain the appropriate contextual details to enable learning.

What. The basic requirement to provide support is a description of *what* the action was.

Who. The identity of *who* the model is, is critical in social learning because it allows the viewer to decide if the presented behaviour is of interest or not. The model can be an identifiable individual or a representation of behaviour that learners will see value in and be motivated to learn from. In the case of representations, these may created through some automated method,

---

[10] Note that these have some similarity to micro and macro persuasion as described in Section 2.3.4.1.

such as aggregating the behaviour of many individuals to represent the behaviour of an ideal model.

When and Where. In addition to the basic requirements of 'who' and 'what', further contextual information about interaction history can better allow the seeker to understand what the provider's behaviour means. This might include details of *where* and *when*.

Filtering. In a similar way that systems may filter interaction history (as described in 3.2.1), a learner may also choose to filter the traces of behaviour presented by a social feedback system. Through selectively filtering the presentations of behaviour, learners can more easily view traces that relate to a task they have. While filtering is not a contextual detail itself, it makes use of contextual details to limit the presentation of information to only that which is relevant to a learner.

In Table 3.2, I summarize the discussion on contextual details of interaction history by describing how they relate to the important concepts of social learning. This table illustrates how each of the learning concepts can make use of the details surrounding interaction history, including *what, who*, *when,* and *where*. It should be noted that these are not requirements for the learning concept, but rather possibilities. The connection between the social learning concept and the interaction history details are illustrated in the 'description' column.

Who will be the models for social learning: Role Models or Tightly-Knit Contacts?

The identity of the model is essential for social feedback (see 3.1.2) because it affects both the seeker's attention and motivation to learn. Further, being aware of *who* provided information allows a seeker to understand if the information is within the context of their current task. For example, if I need help with "non-parametric statistical tests", I would not ask my wife because I am aware that her experience in statistics has been limited to parametric tests. However, if I needed information about what muscle I have strained from sitting too long at the computer I would ask her, because I know that she is a physical therapist and has a lot of knowledge on the subject. In social feedback I suggest two general forms that models can take: *role models* and *tightly-knit groups*.

Role models include people whom the seeker feels have some competency or some prestige related to the topic or task they are trying to complete. Role models could be known to the seeker or may be unknown. If they are unknown, then the social feedback system should identify the model as being a potential role model and identify the qualities that make the model an appropriate role model. For example, expertise, experience, reputation, or past success could all be qualities that make someone an appropriate role model.

Tightly-knit contacts include people who the seeker is aware of or knows to have shared interests, goals, and tasks within a particular personal context, and from whom the seeker would appreciate support. For example, within the context of "children's activities in Saskatoon" I would appreciate support finding activities of interests for my children from the group of parents that I know personally. Similarly, within the context of "research in HCI" I would appreciate support from my extended social network of colleagues, who I know both personally and indirectly (through conducting my research). It is important to note that a tightly-knit contact may also be a role model.

Table 3.2 The specific details of interaction history that can support the important concepts of social learning theory, and a description of how they can relate with example scenarios.

Social Learning Concept	Interaction History Details	Description
Attention	• *What* (the behaviour of interest) • *Who* (identity of the model) • *When* (when an act occurred, and how long it lasted) • *Where* (where the action occurred)	People may selectively attend to information depending on who created it, what the action was, and when and where it occurred (e.g., I am looking for the hotel that my friend stayed at in Vancouver last month).
Motivation and Self-Reward	• *What* (the behaviour of interest) • *Who* (identity of the model) • *When* (when an act occurred, and how long it lasted) • *Where* (where the action occurred)	People could be motivated by who performed the actions, what the particular actions were, where it occurred, and how long they lasted (e.g., perhaps they would like to recreate or compete with the model's behaviour to execute it better or faster).
Retention and Reproduction	• *What* (the action: searching, reviewing a page, clicking a link) • *Who* (identity of the model) • *When* (when an act occurred, and how long it lasted) • *Where* (where the action occurred)	The details of displayed actions cannot be overly complex for the situation or it will be difficult to remember, recall and reproduce them when they are needed. An alternative presentation of behaviour may be able to facilitate learning. For example, telling a learner that "John stayed at Hotel A, then Hotel B, then Hotel A, then B, then C, then B, then A" may be too complex for the learner to follow. A better way to represent this past behaviour may be that John stayed at hotel A 3 times, hotel B twice, and hotel C once.
Self-Observation and Judgment	• *What* (the behaviour of interest) • *Who* (identity of the model) • *When* (when an act occurred, and how long it lasted) • *Where* (where the action occurred)	In certain situations, behaviour that is difficult to learn or maintain can be strengthened by supporting self-regulation. Allowing the seeker to compare relevant aspects of their own behaviour with that of a model, can enable them to get a more accurate understanding of their own performance and attitudes, and also allow them to use a model as target for change (e.g., I have been spending one hour a day on social networking websites, and the most productive member of my group only spends 15 minutes).

3.3.1.3 How can the information be displayed?

To determine how a model's behaviour can be displayed in a social feedback system, designers should consider four questions: How can the behaviour be presented so that it is easy to learn? How can the presentations of behaviour fit into the task? Does there need to be support for self-observation or comparison? And, where will the motivation for learning come from?

How can model behaviour be presented so that it is easiest to learn?

Designers must keep in mind that information describing behaviour should be made as simple to understand, recall, and reproduce as possible. To achieve this, only information relevant to making decisions in the process of completing the task should be displayed [167]. At the very minimum this should include a basic indication of what a particular behaviour or action was and who performed it. Other information may be required or be beneficial to help the leaner better understand the context or result of the behaviour, such as when the behaviour occurred and how long it lasted (see the previous discussion of "What information do people need in the task?").

How can model behaviour fit into the task?

Designers must also decide how model behaviour can fit into the typical process of conducting the task. I identify two general approaches for how behaviour can be presented: *externally* and *internally.*

An *external* presentation of behaviour means that the model's behaviour is viewed, explored, and learned separately from the typical actions associated with completing the task. An appropriate analogy for an external representation is that of a person using a map. A person first consults the map, orients themselves and makes a plan of how to move towards their destination. The map may need to be consulted several times, so that the person can verify that they are on course. The use of a map is done separately from where they will actually act; in the physical world. An external approach can be seen in Figure 2.10, where the Footprints system provided maps of peoples most commonly travelled paths through a website.

An *internal* presentation displays information directly where it can be used. An appropriate analogy for an internal representation is that of a worn path between two buildings on a campus. A new student to a University may be told to get the building they are looking for they can simply follow the worn path across a field. The student can simply start walking the path that has been created in the grass through the actions of many other students. Unlike an external approach (the map), the student has constant feedback about their actions by seeing whether or not they are on the path. An example of an internal approach can be seen in Figure 2.11, where the scented widgets display traces of interaction and contextual details in the same place that a user interacts with the system.

There may be a trade-off that designers must contend with when selecting an internal or external approach. On the one hand, an internal approach can reduce the burdens on recall and reproduction of behaviour, because traces of interaction history can be presented exactly where it is needed and where it can be used. However, an internal approach introduces new information to a task that may already have well-established processes, and may make completing tasks more difficult or cumbersome. On the other hand, an external approach for displaying model behaviour may be more difficult for a learner to recall and reproduce because it is presented separately from where the information can be used, but does not risk adversely affecting the task processes.

Does there need to be support for self-observation or comparison?

In some situations, a learner will need a better understanding of their own behaviour in order to learn how to complete tasks with more success. To provide this understanding, designers may provide ways for learners to get a realistic idea of what their own behaviour is through self-observation. In some situations viewing one’s own behaviour does not necessarily provide an understanding of what the behaviour means, so comparison of personal behaviour with the behaviour of others may facilitate this understanding, especially when put into the context of success or failure (e.g., “John types faster than you do, because he looks at the keyboard once every hour on average, while you look at the keyboard on average two times per minute.”). Further, comparison of personal behaviour with the behaviour of a model can provide motivation. Motivation can be provided through competing with a model or one’s past self, or through cooperating with a model. Designers must take care not to present aspects of behaviour in such a way that a learner feels it is unrealistic for them to make a meaningful change (e.g., “I will never be able to do that”), or they may become discouraged and opt not to act on the information.

Where will the motivation for learning come from?

Simply by using a social feedback system, a learner may be motivated by the promise of completing a task with more success (an *intrinsic* motivation). Further, comparison may provide *extrinsic* motivations by providing the opportunity for cooperation or competition with others. However, in some cases designers may need to encourage learners by providing motivation through explicit rewards; for example, recognizing and publicizing learner achievements for completing tasks with more success, or through providing material goods (such as prizes). Another possibility is to incorporate elements of games into a task that makes learning or acquiring target behaviour more enjoyable (e.g., [58]).

3.3.1.4 How can potential challenges be contended with?

As previously described (in section 3.2.2), there are two main challenges to the collection and sharing of interaction history on the Web: the overlap challenge and the privacy challenge. In this section, I generalize these challenges to interaction history created in any context, and I describe possible general strategies for dealing with them. This discussion may not be exhaustive, however, and other strategies are likely possible.

How can the overlap challenge be contended with?

The overlap challenge can arise in social feedback systems when previous user interactions or tasks do not share enough similarity or occur to infrequently to provide useful feedback. To contend with the overlap challenge, I identify three general strategies.

First, common task starting points for presenting interaction history can be leveraged. Many activities and tasks have multiple ways in which they can be completed. For example, a sketch of a house may be drawn in many ways, but always starts with putting pen to paper. Common starting points can be used as an opportunity to present recent and relevant feedback about how others have approached a task.

Second, if no common starting point exists, create a common place where overlap can occur. Given that a user of a social feedback system is motivated to learn from others, they may be willing to change the way they start a task in order to encounter or view how others have approached a problem or task.

Third, generalize interactions so that they are free of a particular context or situation. This strategy requires some understanding about the type of task to be supported; including what the important behaviours and techniques that are used in completing a particular sort of task, and how successful behaviour may be objectively identified. For example, if one could identify that all the best drawings of animals (by some objective measure) are started with the artist drawing the head, then providing this information would be useful to anyone drawing an animal.

How can model privacy concerns be managed?

The privacy challenge can occur in social feedback systems when a model's captured interactions reveals information that the model would prefer to keep private. Below, I identify five general strategies for dealing with privacy challenges, which may be used together in combination.

First, data can be anonymized to hide the identity of models from being viewed directly by users of social feedback system. Anonymization may involve attaching a pseudonym or anonymous identifier to a model's behaviour, which would allow a learner to identify the common origin of traces that come from a single model. This strategy might work well in certain circumstances, where the behaviour being viewed does not necessarily suggest the true identity of the model. However, in some situations a learner may have enough contextual knowledge of the model and their activities that viewing representations of behaviour may still suggest the model's identity. Further, some captured activities, by their nature, make the identity of the model clear. While anonymization is a simple step that can be taken to help manage privacy, it may be a limited mechanism in many applications. Designers of social feedback systems should carefully consider whether anonymization of interaction history is a sufficient technique for managing privacy.

Second, interaction history distribution mechanisms may be added to a system, which would allow a user to control or set policies for what interaction history may be collected. In its simplest form, a distribution mechanism may be an on/off switch that users can set to allow interaction history to be collected or disregarded by the system. More complex policies may allow users to define what activities may be collected by the system and under what contextual conditions (including when and where activities take place), and what will be shared with particular people or groups. Such detailed mechanisms, which may include whitelists or blacklists, would allow users fine-grained control over the sharing of their interaction history. However, there are two main problems with more detailed, policy-based approaches: 1) policies require at least some user effort to initially setup and configure; and, 2) for some types of activities and tasks, user views (e.g., what should be shared and with whom) may change over time, which could increase the burden to maintain policies.

Third, a mechanism for identifying and removing interaction history that contains sensitive information may be created. Such a mechanism would work on a spectrum between fully-manual to fully-automatic classification of interaction history as being public or private. On the one hand, a fully-manual mechanism would allow a system user full control over what interaction history is available to the system. However, this may require undue effort on the part of the user to maintain and explore the data that they have created. On the other hand, a fully-automatic approach may reduce the burden on a system user by requiring little effort to manage the collection of interaction history. However, a fully-automatic approach may not result in all sensitive data being disregarded (i.e., false-negatives) or it may identify data as being private that a user would not mind sharing (i.e., false-positives).

Fourth, another approach would be to aggregate, or otherwise combine, the interactions of many users. This method may protect privacy by filtering out individual user activities, only presenting data that is common to multiple users. Aggregation can reduce the likelihood that individual activities can lead to the identification of many users. Aggregation may be appropriate for supporting many types of tasks, because it can also serve to reduce noisy or idiosyncratic behaviour. However, for some applications hiding the details of individual traces may reduce the utility of the approach. For example, if I would like to learn how to swing a golf club like a particular golfer, say Tiger Woods, then presenting data on the golf swings of all professional golfers may not be suitable.

Fifth, and finally, the distribution of shared interaction history can be limited to groups where there is an existing level of trust and understanding (with regards to the particular activities being described). People who are socially close to one another may already have an accurate understanding of one another's behaviour, and would not negatively judge one another based on what can be viewed through shared interaction history. By limiting the use of shared interaction history to describe the activities within a small group, privacy may be less of a concern. The main disadvantage of this approach is that it would limit the number of models available in the system, effectively limiting the opportunities for learning.

3.3.2 Contextualizing the Social Feedback Design Questions to Web-Based Information Seeking

In this section, I narrow the discussion of the design of social feedback systems to the primary use of it in this dissertation: for supporting information seeking on the Web. I will address the previous design questions when further details are required. In this section I refer to the *seeker* as the person who is making use of a social feedback system.

3.3.2.1 What are the tasks or activities to be supported?

Below, I reconsider the task to be supported in terms of the common information seeking tasks that people conduct on the Web. In general, the different types of information-seeking tasks are well understood (see Section 2.2.1). The three core information-seeking tasks (or core tasks), then, are:

1. fact-finding,
2. information gathering, and
3. browsing.

Transactions (online actions such as sending email or banking) will be largely as information-seeking tasks ignored in this work, because they do not necessarily seek to fill a gap in a seeker’s knowledge, they often use specialized systems that do not rely on typical interaction patterns, and they frequently include information that a seeker may prefer to keep private.

What is the basic behaviour used to complete these tasks?

The three core tasks above are most often based on and involve three *basic interactions* on the Web, which include:

1. searching for information using a search engine,
2. reviewing webpages, and
3. selecting links.

The three basic interactions are also the basic points at which a seeker needs to make a decision. For example, during a fact-finding task a seeker might be having difficulty formulating a query that will return information of interest, or while browsing a seeker might wonder which link they should follow to find the best resource. In these examples, the seeker could make use of further information to help guide their decisions. Since the basic interactions are the decision points for seeking information on the Web, providing support at these points can also support each of the core information-seeking tasks.

Can the task behaviour be captured and represented through interaction history?

Browsers can be instrumented to capture a wide variety of user behaviour, from details of scrolling behaviour to the pages users bookmark (see [82] for a review). However, the use of most web browser functionality – including how bookmarking functionality is used, how many tabs are used, or if a page has been printed – is highly idiosyncratic.

I therefore focus my attention on the basic unit of interaction on the Web – the webpage visit. Regardless of how a user makes use of their browser, everyone must visit pages in order to find, gather or browse information. While social feedback need not be limited to interaction history from webpage visits alone, model behaviour should be presented in a way that is understandable by a seeker. Therefore, in this initial work on social feedback I limit the interaction history to what can be collected from time-stamped URLs with a unique user identifier. This simple data can provide a rich awareness of the information seeking process, including: what pages were visited, who visited the pages, when a visit occurred, how long a visit lasted, and the content of search queries (from the query string of a URL).

Time-stamped URLs can be supplemented with external data sources that can further describe the context of interaction history. For example, given a time-stamped page visit, a web service could be queried to find contextual information, such as the currency exchange rate on the day of the page visit, or the weather conditions at the time of the page visit. User-identifiable logging can be, however, a potential privacy concern (as discussed in Section 3.2.2).

3.3.2.2 What information do people need in the task?

Because the type of information that people need in each section is highly dependent on the specific task and activities that are to be supported, I leave this discussion to the specific applications of social feedback (in Chapters 4 and 5).

3.3.2.3 How can the information be displayed?

Below, I reconsider the approaches for displaying information in terms of web-based information-seeking tasks.

How can model behaviour be presented so that it is easy to learn?

Support for loosely-coupled information seeking must be low cost in terms of effort (both to gather and interpret), and should fit within existing work practices. Thus, social feedback should be available within established interaction patterns based on webpages, web browsers, and common tools such as search engines. Creating new tools that change the basic interaction patterns may lead to additional overhead.

How can model behaviour fit into the task?

As introduced above, the information presented in social feedback systems can fit into the task in two general ways: *externally* or *internally*. These general approaches can be made concrete in web-based information seeking by considering whether the models' behaviours will be displayed within a webpage (internal) or outside of a webpage (external) that is part of an information-seeking task. Below, I describe the internal and external approaches for displaying interaction history in webpages.

External versus Internal to the page. Previous work on collaborative information seeking has provided support widgets that display or summarize participants' activities (e.g., [106,107,134]), which allows a large amount of detailed information to be displayed (e.g., using a sidebar). In contrast, work on information scent often augments webpages to provide proximal cues (e.g., [116,45]), allowing users to see additional information scent within the page. However, the internal approach means that only limited information can be displayed without obscuring webpage content. The form chosen should consider both the task and the type of learning (local or global) to be supported.

Presentation External to the Page. If representations of behaviour are to be presented external to the page, there is no limitation on the specific presentation techniques used. However, a designer should aim to make the representation as quick and easy to understand as possible, and consider how and when a seeker would encounter the presented information. Further, keeping the presentations simple may help the seeker retain and reproduce the behaviour when it is needed. This is especially important in external presentations, as designers must assume that the seeker will not be able to see the presented behaviour when making decisions during information-seeking tasks.

Internal Augmentation to the Page: Artifact Encoding vs. Glyph. Interaction history can be provided with new graphical objects (glyphs) placed near the artifacts of interest (i.e., pages and links) [45], or information can be encoded into the representation of the artifacts themselves (e.g., changing the font size of links [116]). Although glyphs can occlude information on the page, they can be consistently displayed without altering the style and functionality of the website. Encoding information in the artifacts themselves, such as changing the color of a page or link, may lead to undesired consequences in the display of the page.

3.3.2.4 How can potential challenges be contended with?

On the Web, the challenges of overlap and privacy may be of particular concern. First, the Web is extremely large and two users with the same task or information need may be able to complete their task without visiting any of the same webpages or websites. Further, a particular webpage may support multiple tasks, so providing social feedback based on shared locations may not guarantee that the support is appropriate. This means that providing useful support based on overlapping tasks could be particularly challenging on the Web. Second, the Web encompasses many different types of activities – including those that relate to personal communications, health, and finance – and it is clear that many people would prefer to keep many of these activities private.

How can the overlap challenge be contended with?

The three general strategies suggested above can be applied directly on the Web. Leveraging common starting points is exemplified in the work on WebWear (in Chapter 6). Creating a new starting point and generalizing interactions so that they are free of a particular situation are demonstrated by the Search Dashboard (in Chapter 5).

How can model privacy concerns be managed?

The concerns of privacy may be particularly challenging on the Web, and there may be many details that must be managed by system designer. For example, viewing individual traces, even if anonymized, may reveal sensitive information, because some URLs contain identifiable information (e.g., user X visited http://example-site.com/login?current_user=scottbateman).

Again, the five general strategies for managing privacy concerns suggested above can be applied directly to the web-based information seeking. The Search Dashboard (in Chapter 5) demonstrates both anonymizing and aggregating interaction history. The WebWear system (in Chapter 6) demonstrates simple mechanism for controlling the distribution of interaction history (through on/off functionality) and a simple mechanism for the identification and removal of sensitive interaction history (through a manual approach to removing history from the system). Further, the studies of WebWear demonstrate the use of social feedback within a group where there is an existing level of trust.

3.4 Summary of Social Feedback

In this chapter I have described the foundation of social feedback, which is social learning theory – the theory of how people learn through observing others. I described the four main parts to the theory: the requirements for learning from others, the identity of a model, the basic forms of models, and the requirements for self-regulation. I then described how the social learning can be instantiated through the mechanism of interaction history. I outlined the main steps and challenges in collecting and using interaction history for describing behaviour. Finally, I described the important design concerns in building general social feedback systems, and I then contextualized these concerns to the case of web-based information-seeking tasks. In Chapters 4 and 5 I will demonstrate how the idea of social feedback can be applied in the design of two different systems: one to provide global learning and the other to provide local learning.

CHAPTER 4
GLOBAL SOCIAL FEEDBACK: THE SEARCH DASHBOARD

Learning how to most effectively and efficiently use a web search engine is a type of global learning, because search engines are a commonly used starting point for many information-seeking tasks. However, it is known that some search engine users are less successful than others (see Section 2.2.3). Therefore, search engine usage is a part of many loosely-coupled information-seeking tasks, because there are often other people who could provide support to a search engine user to help them improve their success.

Although search engines work very well most of the time, many people still experience problems finding what they are looking for. In a recent survey, nearly five percent of users reported completely failing at their most recent attempt to search for something on the Web [54]. When search failures occur, they cost people a lot of time: searchers spend over ten minutes when they fail before giving up, as compared to needing less than five minutes when they are successful [54]. Part of the problem is that when search becomes difficult, many people are unsure of how to change strategies or how to make use of the advanced search engine functionality that could help them [9,110]. People largely use the same search behaviour regardless of the situation and how successful they are [110]. Further, even when people are successful, they still may not have been as efficient as they potentially could have been. This suggests that searchers have room for improvement.

The main approach for improving people's search success has been to improve search engines to deliver the best results for a given query. However, another approach is by educating users to be better search engine users [102]. Previous work has shown that the way people use search engines can predict how successful they are in completing their information-seeking tasks. In particular, if people use of particular search engine functionalities (e.g., the use of search operators) [7] and the characteristics of search engine use (e.g., the average number of terms per query)[158] can predict a seeker's overall success (see Section 2.2.3). This suggests that if users knew more about how to use search engines they would be able to improve their own search performance.

Different methods to help users improve their search behaviour have been relatively uninvestigated. However, it is clear that some search engine users perform better than others (see 2.3.4). One problem is that seekers are not able to get an idea of their own search behaviour and are not able to find out what behaviour works for other people. The lack of facilities for feedback on personal behaviour means it can be difficult for users to learn how they can get better, and what search strategies can lead to improved performance. Further, even if people could get an accurate idea of their search behaviour and how it might be changed, it is not clear whether they would be able to adjust their search behaviour in any meaningful way.

To show that social feedback can have an influence on the attitudes and behaviours of searchers, I created and studied the *Search Dashboard*. The Search Dashboard (or *the Dashboard*) system is the first that aims to positively influence user search behaviour through social feedback, which displays personal interaction history and allows comparison with the behaviour of expert models (such as search or topic experts). I performed a five-week study of 90 users and examined how social feedback affects user attitudes and understanding of how search engines can be used, and how this feedback can lead to observable behaviour change.

The study provides two main results. First, social feedback *can* lead to global learning, as demonstrated through changes in attitudes and behaviour in actual search engine use. Second, providing facilities for comparison with expert models leads to increased interest and engagement over simply allowing reflection on personal behaviour alone, and is critical in learning that leads to behaviour change.

4.1 The Search Dashboard: Social Feedback for Global Learning

I created the Search Dashboard system (Figure 4.1) to provide social feedback, with a secondary goal of understanding what types of behaviour are most useful and influential to searchers in their actual search engine usage (a video figure link is available on page xvii).

Based on the conceptual framework laid out in Section 3.3, I begin by describing the types of interaction history displayed in the Search Dashboard. I next describe how the data for expert models was collected. Finally, I discuss how all of this information was gathered and presented to searchers.

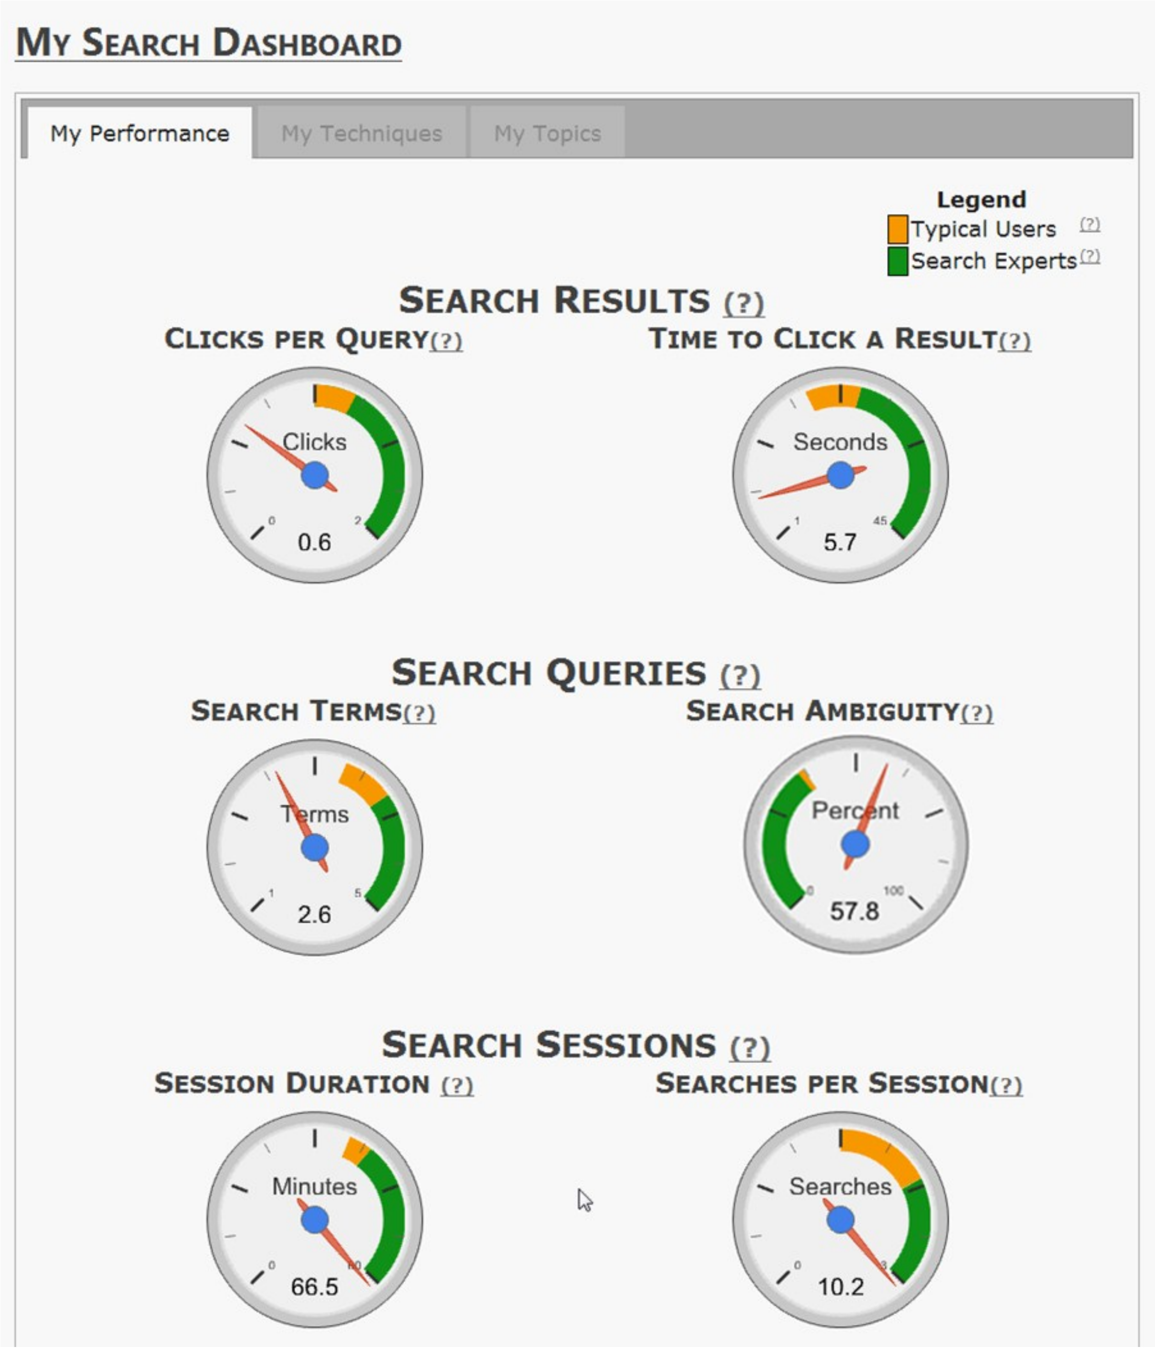

Figure 4.1. The Search Dashboard, displaying the Tendencies section (referred to as Performance in the study).

4.1.1 What tasks should be supported?

The tasks to be supported are the core information-seeking tasks (fact-finding, information gathering, and browsing), that are enabled through global learning of general search engine skills. As previously discussed (see Section 2.2.3), performance in information-seeking tasks using search engines can be affected by knowledge of search engine features [72, 102] and topical expertise [161]. We, therefore, aimed to support seekers by supporting their general search engine usage strategies, their use of search features, and their use of topically relevant information.

4.1.2 What information do people need in these tasks?

To identify what interaction history should be presented for social feedback, I initially conducted a 53-question survey of 75 employees at the Redmond headquarters of Microsoft. Participants were solicited by email and asked to rate their interest in being able to see different aspects of their search history (e.g., the days of the week they search on, or their average query length), and whether they might be interested in comparing that information with other individuals or groups (such as friends, colleagues, or Web search experts). Although the main research questions relate to behaviour change, the initial survey focused on interest based on one of the requirements for social learning: that people must have an interest and perceive value in an observed behaviour in order to have motivation to learn and change it (see Section 3.1). Results from the survey are highlighted where appropriate below.

4.1.2.1 Data Types (Techniques, Tendencies, and Topics)

Based on a survey of related work (cited below and described in Section 2.2.3), I identified three main types of interaction history that may be valuable to learn from (called *Data Types*), since they have been shown to affect search performance when using a search engine:

- *Techniques*: The use of advanced query operators and special search engine features (e.g., [8,102,158]).
- *Tendencies*: Summative actions describing the overall characteristics of search engine use, such as the number of terms used per query (e.g., [72,158,160]).
- *Topics*: The content and subject (or topical domain) of an individual's search engine use (e.g., [9,22,72,90,130]).

Using the formative survey and previous work, I selected 12 specific *elements* of people's search history to display in the Dashboard. These elements represented different points in the space of Data Types (i.e., Techniques, Tendencies, and Topics) and were those most often selected as being of interest by survey respondents (mean selection rate = 47%, as compared to an average of 34% for other aspects). These elements, and how they are calculated, are now described.

4.1.2.2 Techniques Data

Operator use: There are a number of search operators understood by Web search engines, including quotation marks (used to group query terms together), '+' (used to mark a term to be unaltered by the search engine), '-' (exclude a term from a query), and 'site:' (used to restrict results to a particular website domain). A person's search operator use is represented by counting the number queries issued by that person that contain a particular search operator.

Vertical use: In addition to general Web search, there are a number of 'vertical' search interfaces that enable searchers to search particular types of data (e.g., images, maps) or perform particular types of tasks (e.g., shopping, travel). Awareness of the different verticals available is important, so the number of times each vertical is used is counted.

Direct Answer use: Direct answers are information tailored to specific queries, e.g. for weather forecasts. They are shown inline on search engine result pages to address searchers' needs quickly. The number of each type of direct answer a person encountered in search results is counted.

4.1.2.3 Tendencies Data

Query length: An individual may issue long or short queries. The average number of search terms an individual uses in their queries is calculated.

Clicks per query: People vary in the number of search results they select after querying. The average number of results an individual clicks following a query is calculated.

Time to click: After a person issues a query, they may click a result right away, or pause to consider the retrieved results first. To capture this, the average time between issuing a query and clicking on a result is measured in seconds.

Session time: Queries often appear as part of a search session, rather than in isolation. Sessions are identified by looking for search activity with less than 10 minutes between each action, and compute the average duration, in minutes.

Session queries: As a general measure of activity level, the average number of queries that each searcher issues per session is also computed.

Query ambiguity: Query ambiguity is measured by examining the variation in the search results selected for a particular query. Historic search log data is used from many searchers over one month before the beginning of the study, and the average click entropy[11] is calculated, as defined in [48], for each searcher across all queries they issue. The click entropy gives a simple measure of how specific the intent for a given query is (i.e., if users click on the same few results for a given query, then the click entropy is low). Therefore, click entropy is a the opposite of ambiguity, so it is calculated as:

$Ambiguity(q) = (Max(ClickEntropy(Q)) - ClickEntropy(q))\ 1/Max(ClickEntropy(Q))$, where Q is the set of all queries.

4.1.2.4 Topics Data

Categories: The most popular topical categories of the search results clicked by an individual are aggregated. The category of a URL is determined by selecting the highest probability category returned by a content-based Web page classifier (described in [17]), which assigns URLs to Open Directory Project (ODP, dmoz.org) categories. Thirteen top-level ODP categories are used for this labeling: Arts, Business, Computers, Games, Health, Home, Kids and Teens, News, Recreation, Science, Shopping, Society, and Sports.

Domains: The most commonly accessed domains were extracted from the search results clicked by an individual. For example, if a person clicks on many Wikipedia search results, wikipedia.org is a distinctive domain for that individual.

Query terms: To show the most salient search words, the individual terms are parsed from each of a person's queries, stop words are removed, and the frequency of each term counted.

---

[11] Click entropy [48] for a given query (q) is calculated in the following way:

$$ClickEntropy(q) = \sum_{p \in P(q)} -P(p|q)\log_2 P(p|q)$$

Where p represents a page that has been previously clicked in the set of all pages that have been clicked P. $P(q)$ represents all webpages clicked for query q, and $P(p|q)$ is the percentage of times that a page was clicked for the given query among all clicks for the query.

4.1.2.5 Who are the models?

Comparison data was gathered for several different models based on the conceptual framework of social feedback. These models allowed a searcher to obtain a realistic idea of their own behaviour and performance in using a search engine for a given behaviour element, and provide targets of how their behaviour might be changed. Further, about two thirds of the initial survey respondents reported wanting to be able to compare their personal search information to others (65.3% agreed). They were most interested in comparing their data with "people who are experts in topics I am interested in" (68% agreed) and "people who are experts at searching the Web," (51% agreed). They were less interested in comparing to people they know, such as colleagues (41% agreed), or family and friends (25% agreed).

The three models that were created represented *Typical Searchers*, *Search Experts*, and *Topic Experts*. To generate the representative values of the 12 data elements for each model, sets of representative searchers were algorithmically selected from the opt-in logs from a browser plugin widely deployed by Bing, using data from one month starting April 15, 2011. To remove variability caused by geographic and linguistic variation, only entries generated in the English speaking United States were included. Individual searcher history was then aggregated over all searchers to generate a single model profile. I now describe how the searchers were identified.

Typical Searchers: This model allows Dashboard users to compare their behaviour with typical searcher behaviour. The typical searcher model was created to allow users to get a better understanding of how their current behaviour compares to typical usage. To generate the model, 1000 searchers were randomly selected from all non-search experts during the sampling period.

Search Experts: To create this model, the approach of White and Morris [158] was used, who showed that web searchers who tend to be more successful could be identified by the use of search operators. 1000 searchers were selected who most frequently used search operators during the sampling period. Search experts were used for comparison with Techniques and Tendencies data, with the intent of providing searchers with a target for how their behaviour might be changed.

Topic Experts: To create these models topically relevant subsets of searchers were selected from the combined pool of 2000 searchers previously selected as typical or expert. A searcher was considered topically relevant if they had visited at least 10 search results in the category in a one-week period. This approach created 13 topical expert model, one associated with each topic. The number of searchers identified for each topical model ranged from around two dozen for less frequently used categories (e.g., Kids and Teens) to several hundred for the more frequently used categories (e.g., Computers). The 13 generated topical experts were presented with Topics data only. The topical experts model could provide a target for how a searcher might change their search behaviour.

4.1.3 How can the information be displayed?

4.1.3.1 Design Factors

Information about an individual's personal search history and comparison information for the three different Data Types are displayed in the Search Dashboard system. The Dashboard is a standalone web application that searchers could visit in their web browser. The Dashboard is accessed using an intranet URL from a standard Web browser. When searchers visit the website, their personal search history is automatically loaded. Each Data Type (Tendencies, Techniques, and Topics) is displayed separately using a tabbed interface.

It took two to three days for a searcher's personal data to be aggregated and presented in the system. This means that while searchers could refer to the Dashboard during a search task (by visiting the Dashboard in a different web browser window), they could not use it for instant feedback on their actions. The approach of creating a separate website for social feedback would be considered as "external to the page" (see Section 3.3.2.3), because it does not provide details of behaviour and social feedback within the search engine webpage. This approach was taken for two reasons. First, I wanted to have searchers to be able to consider a number of different elements of their behaviour at once and provide judgments on them. It would not have been practical to describe all of the behavioural elements within a single search engine page. Second, the most likely scenario for such a system to be incorporated would be as a replacement or augmentation for current similar systems (e.g., Figure 2.3).

4.1.3.2 Comparison and No Comparison Variants of the Dashboard

Current systems present interaction history from personal search activity only (e.g., Figure 2.3). So, I was interested in the differences between a display that is similar to current systems and how it might compare to a social feedback system. To better understand the value of comparison data in social feedback, I created two versions of the Dashboard (called *Variants*): (i) the *no comparison* Dashboard showed only a searcher's personal data; (ii) the *comparison* Dashboard added the data of models for comparison.

4.1.3.3 Displaying Textual Data

Two types of data are displayed: textual and numeric. Textual data can be found on the Techniques and Topics sections, and are presented as lists in tables (see Figure 4.2). The top five values for a textual attribute are displayed, and the full list (up to 100 values) can be viewed by clicking the "see more…" link below the tables.

When comparison information is included with textual data, two additional columns are added to the table to present archetypal searcher data for the typical and expert searcher. The expert used on the Techniques tab and Tendencies tab is a search expert, and on the Topics tab is a topic expert. Although there are 13 topical expert models, only one was shown at a time. A different topic expert can be selected via a drop down list positioned by each of the topic expert data tables. The topic expert that is most similar to the topic that the user is associated with is selected by default.

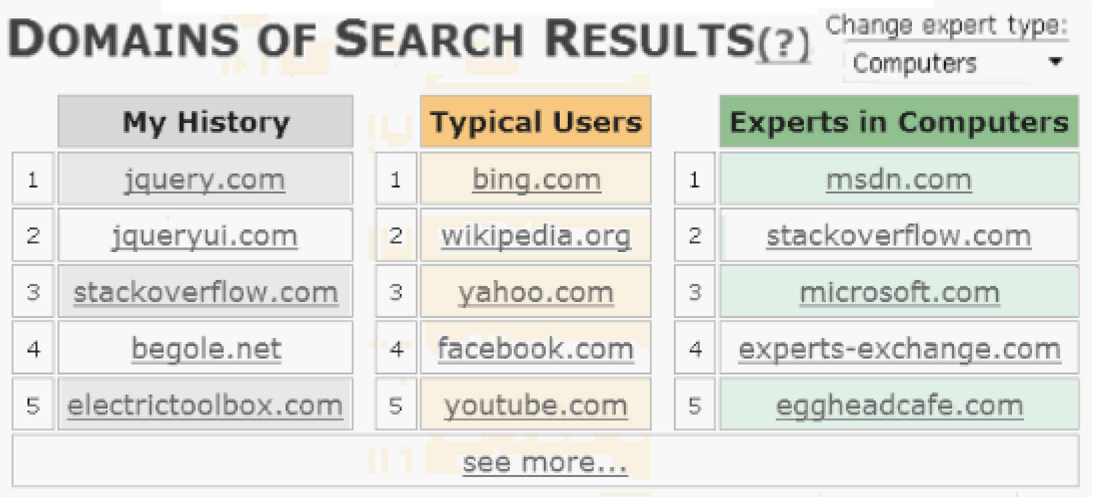

DOMAINS OF SEARCH RESULTS(?)

Change expert type: Computers ▾

	My History		Typical Users		Experts in Computers
1	jquery.com	1	bing.com	1	msdn.com
2	jqueryui.com	2	wikipedia.org	2	stackoverflow.com
3	stackoverflow.com	3	yahoo.com	3	microsoft.com
4	begole.net	4	facebook.com	4	experts-exchange.com
5	electrictoolbox.com	5	youtube.com	5	eggheadcafe.com

see more…

Figure 4.2. The Domains data in the Topics section.

Because some textual data, like domains and query terms, occur more often than others, and thus are more likely to be used without an existing preference, I identified an individual's most distinctive data by normalizing the count of the textual item by the count for Typical Searchers. This is analogous to the TF.IDF method from information retrieval [138], and resulted in data being ordered by the most distinctive items, rather than those that occurred most frequently[12]. This approach is used for a Dashboard user's own data and for expert model data.

4.1.3.4 Displaying Numeric Data

Numeric data, which includes all of the Tendencies data, is displayed in two ways (both are shown in Figure 4.3). When no comparison data is used, only the number and unit of the measurement are displayed in large font. When comparison data is available, numeric data is shown using a gauge chart, which allows a number of data values to be concisely displayed. On each gauge the position of the needle and the number displayed indicates the user's value for the particular measure. Two regions are also displayed, an amber region representing the range of typical searchers (defined as the values between Typical Searcher's value and the Search Expert's value), and a green region representing the expert searcher range (defined as the value from the Search Expert's value to the extent of the gauge). I opted for gauge charts after piloting alternatives (e.g., bullet graphs); they were familiar and required the least explanation.

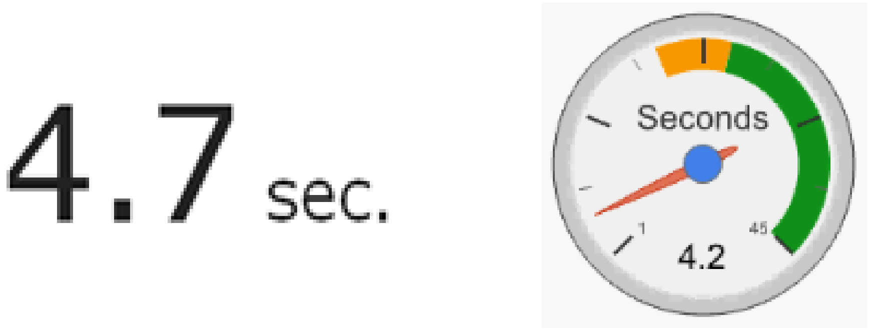

Figure 4.3. Alternative presentations of numeric data: without comparison data (left), and with comparison data (right).

4.1.3.5 Exploring Data and Getting More Information

The Search Dashboard also provides facilities for users to obtain explanations and more details about any of the data displayed. Each data label in the Dashboard provides a brief description of the data element and possible interpretations of values, via a tooltip. These labels are intended to be descriptive rather than instructive. For consistency, the same descriptions are

---

[12] Specifically, the distinctive value for a Dashboard user or expert (*u*) for some information item (*i*), e.g., a domain, a keyword, a topic), was calculated in the following way:

$$DistinctiveValue(i|u) = \frac{UsageCount(i|u)}{UsageCount(i|TypicalSearchers)}$$

used regardless of whether comparison data are shown or not. For example, the tooltip describing "clicks per query" reads: "The number of results you typically click on following a query. Some people are very selective with their search results, leading to very few clicks per search; some people average less than 1 click per result. Other people are happy to explore many search results."

Table text is also hyperlinked, and clicking performs an exemplary action, such as launching a search engine query illustrating the functionality. Hovering over text data also provides more details. For the values in the Domains tables (in Figure 4.2), hovering on a domain name calls a popup with the list of queries that led to that domain being visited. Clicking any query in the popup issues the query to a search engine. Figure 4.4 shows the popup for direct answers.

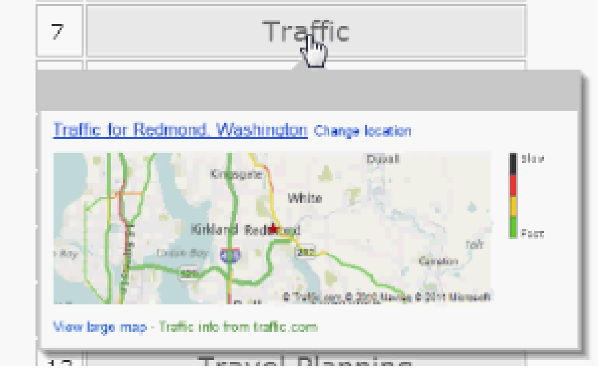
7
Traffic
Traffic for Redmond, Washington
Change location
Slow
Fast
Kirkland
Redmond
White
View large map - Traffic info from traffic.com

Figure 4.4. An example of the Traffic direct answer. When a user hovers over an answer type an example is shown.

4.2 Study

A five-week study with 90 participants was conducted to understand how people make use of different types of information in their search history summary. The study design aimed to address four main questions:

1. Do participants perceive information about their personal search history as valuable?
2. How are the different types of personal search history data perceived and used?
3. Does reflecting on personal search behaviour lead people to change their attitudes about search engines and adopt new search behaviour after seeing the Dashboard?
4. Does the ability to compare one's personal search history with others provide benefit?

The focus of the research questions were on changes in behaviours and attitudes, rather than on search success because success is less accurate to measure, particularly within the context of real day-to-day workplace searches. This is an important first-step, as the ability to effectively impact search behaviour can be valuable regardless of search outcome.

4.2.1 What data was gathered?

To address the four main questions, I collected three forms of data: (i) I logged all participant web browsing activity from which I was able to extract queries issued to all Web search engines and what, if any, advanced search engine functionality was being used. (ii) I collected survey data during each of the three parts of the study (described below). Surveys contained a number of Likert-scale and free-text questions. (iii) Finally, I also collected log data about how participants interacted with the Dashboard.

4.2.2 Conditions

The study involved a mixed 3 × 2 design (*Data Type* × *Dashboard Variant*). Because the first research questions related to how people would value and perceive different aspects of their search history, I grouped the three Data Types (Techniques, Tendencies and Topics) into three tabs (separate sections) on the Dashboard. For the study, the three Data Types were presented in random order, and participants worked with only one Data Type at time, although all participants saw all three Data Types (i.e., within subject).

It was hypothesized that having the data of the models available for comparison would increase how valuable participants found reflection. However, current search history tools provide personal data only. So, I was interested in identifying the differences between personal data and having personal data augmented with the data of archetypal searchers for comparison. For this reason, I developed the two Dashboard Variants (comparison and no comparison) that were studied between-subjects – participants were randomly assigned to use only one of the Dashboard Variants.

4.2.3 Participants

All 90 participants were employees of Microsoft at the company’s Redmond, WA headquarters. Most (73, 81%) were male, which is consistent with the company’s demographics.

All reported searching the Web at least daily, with over half (50, 56%) reporting that they searched more than 10 times a day. Participants were randomly selected from the company directory, and recruited via email. In exchange for participation in all five weeks of the study, they were entered into a sweepstakes for one of four prizes (one $300 gift card, three $100 gift cards).

4.2.4 Procedure

The study was in three parts over five weeks: (i) initial registration, (ii) introduction to the Search Dashboard after three weeks, and (iii) a final exit survey after five weeks.

4.2.4.1 Part I: Registration

At the study outset, participants were asked to enroll using their primary work computer, by completing an initial survey, and configuring a piece of software to log all of their search engine activity for the study period. The registration survey asked basic demographic information, as well as perceptions and attitudes towards Web search.

4.2.4.2 Part II: Introduction to Search Dashboard

Three weeks (on average) after participants enrolled in the study, they were sent an email requesting that they view their personal Search Dashboard. The actual time between enrolling and completing the study ranged from two to four weeks depending on participants' availability.

The search engine activity that was logged during the time between Part I and II of the study was used to populate the personal data in each individual's Search Dashboard.

Participants were randomly assigned to one of the two between-subject Dashboard Variant conditions (either compare or no compare). When participants visited the Dashboard for the first time they were initially presented with a walkthrough that explained their task, guided them through the features of the search dashboard and described how their data would be presented. They were then asked to view each of the three tabs, one at time, for as long as they liked. Recall that each tab represents one of the three Data Types (Techniques, Tendencies, and Topics). To ensure that participants could not simply click through the study without viewing their data, the system enforced a one minute delay on each tab before the participant could proceed. After viewing each tab, a survey was presented that asked about the Data Type they just viewed. After all three surveys were completed, a final survey was presented that asked participants about their overall experience with the Dashboard and asked them again about their attitudes and perceptions of Web search. After the final survey participants were told they could visit the Dashboard at any time.

4.2.4.3 Part III: Exit Survey

Finally, after roughly another two weeks, participants were contacted again. The time between Part II and III ranged from 6 days to 19 days (average 12 days). Participants were asked to visit the other Dashboard Variant – i.e., participants who initially saw the comparison Variant were asked to view the no comparison, and vice versa. Participants were also asked to complete a final survey that was focused on the differences between the two Dashboard Variants.

4.2.5 Data Analysis

Surveys contained 5-point Likert scales (1=strongly disagree, 3=neutral, 5=strongly agree) or 7-point self-rating of skill (1=min, 7=max). Analysis of within-subject factors for survey questions used Friedman's ANOVA for related samples; post-hoc pairwise comparisons used Wilcoxon Signed Ranks Tests for two related samples. For the analysis of questions for between-subject conditions, the Mann Whitney U test was used. In addition to analyzing the quantitative data, the survey comments were used to help explain results found in statistical analysis. The post-hoc tests of survey data and analysis of search log data used many dependent variables, so using Bonferroni corrections to control the experiment-wise error rate, α was set to 0.05 divided by the number of dependent variables.

4.3 Results

I now present the results, organized by main findings. I show that participants found their interactions with the Search Dashboard (and in particular with their Techniques and Tendencies) to be valuable and that it had an impact on their self-assessment of their search skills. The use of comparison data, especially search expert data, led to increased engagement and insights, as well as observed behavioural changes for Tendencies and Techniques.

4.3.1 The Search Dashboard is Engaging

Participants' level of engagement was assessed through the Dashboard interaction logs. Although specific time guidelines were not provided for using the Dashboard, analysis of the Dashboard log data reveals that the participants spent almost half of an hour (28.65 min.) exploring their data. During this time, participants looked at 39.96 tooltips on average. They were also likely to return to the Dashboard on their own after their initial study visit, with 72.53% doing so at least once and the logs recording an average of 2.28 visits per participant.

4.3.2 Techniques & Tendencies More Insightful Than Topics

To understand how insightful the Data Types were, I elicited ratings on "I learned something new", "I was surprised by some aspect of my history", and "Based on what I saw, I will change how I search". All three aspects showed significant differences between Data Types. Statistics are presented in Figure 4.5, and are discussed below.

Overall, participants did not feel the Topics provided much insight in terms of the questions posed, agreeing (below neutral) that Topics would lead to change in how they search. Several participants stated that they were not certain how to make use of topics information: "Interesting information. But, not necessarily actionable." However, a few participants really enjoyed the Topics data. One participant reported, "I love the area of my topics and how it separates them out into categories ... That type of information I find very interesting."

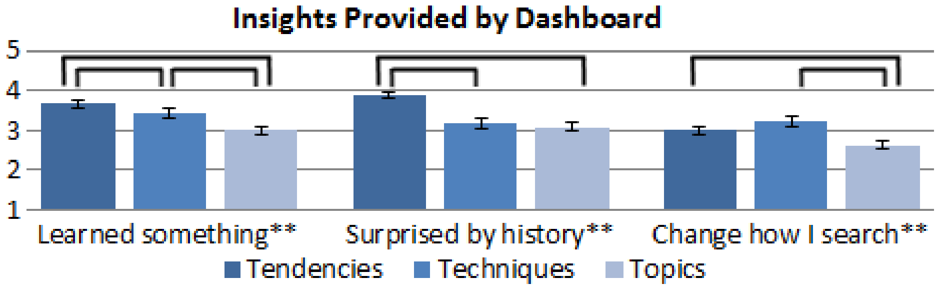

Figure 4.5. Mean ratings and ±SEM of insights provided by the Dashboard for each of the three Data Types. All questions were sig. diff. (p<.001). Sig. diff. pairs are indicated by lines.

The Tendencies and Techniques data were viewed more positively than Topics. While insights were rated higher for both Data Types, participants were particularly likely to be surprised or learn something new when viewing data related to their Tendencies. Again, many participants expressed not knowing how to take action on the data with some commenting:

"probably the [Tendencies] tab [was the most interesting] although nothing was really 'actionable'", and "pretty cool, although not sure how I would modify my behaviour based on that". Despite Topics data receiving low ratings, some participants reported that they found the most value in seeing the data of Topic Experts (such as being able to see the domains that computer experts most often visit). For example, one participant said, "There is good information about what I would use to augment my searches in the future. It's helpful to find… sources of information aside from what I've been searching for in the past. It's like having a friend tell you 'have you checked out this site?'"

Overall, the comments revealed strongly preferred information for which participants could see an immediate application. In these terms, they felt Techniques data would most likely change how they searched. While this was likely due to most people learning something new, being reminded was also valuable: "There are features … that I've forgotten about that this has been a great reminder for."

4.3.3 Self-Assessment of Search Skills Changed

To see how the Dashboard influenced views about search skills, participants were asked to rate their search skills along a number of dimensions when they first agreed to participate in the study (*pre*), and then again several weeks later after viewing the Dashboard (*post)*. As shown in Figure 4.6, through the use of the Dashboard, participants came to believe they were less skilled search engine users than they had thought they were before the experiment began. This suggests that they learned that they have room for improving their search behaviour, whereas initially they saw less room.

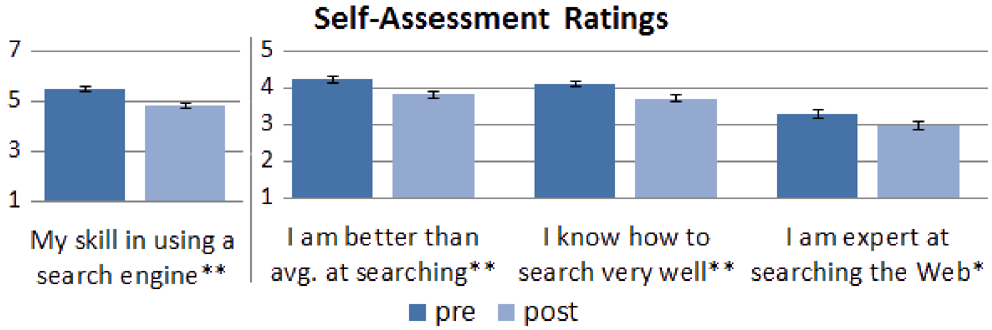

Figure 4.6. Mean and ±SEM for self-assessment ratings before (pre) and after (post) using the Search Dashboard. Sig. diff. were found for all ratings (*=*p*<.01, **=*p*<.001).

4.3.4 Comparison Leads to Increased Engagement and Insights

To see if comparison data led to increased engagement and insights the ratings of Dashboard variants were used as a between-subjects factor (see Figure 4.7). In all cases participants rated the comparison dashboard higher than the non-comparison version. In particular, participants in the comparison condition were more likely to report that the Search Dashboard would change how they search, and as we will see in the subsequent section, this proved to be true. Participants' comments also revealed enthusiasm for the comparison data (e.g., "Interesting. Fun to compare.").

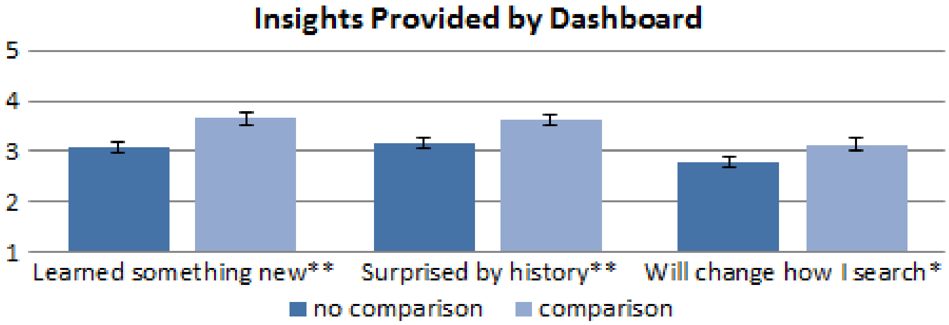

Figure 4.7. Mean and ± SEM for ratings of insights between Dashboard variants. Sig. diff. were found for all aspects (*=*p*<.05, **=*p*<.001).

I hypothesized that participants would explore and spend more time with the comparison Dashboard. A series of one-tail *t*-tests to compare each of the three system-usage variables system usage revealed significant differences (with α=.0167). Participants in the comparison condition spent more time visiting the Dashboard during the study, read tooltips more often, and paid more visits to the Dashboard after their initial visit (see Figure 4.8).

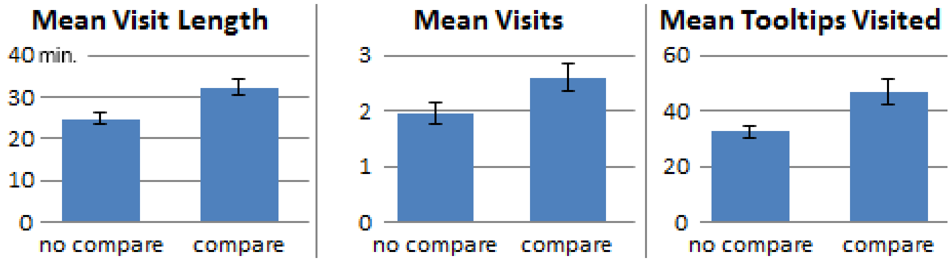

Figure 4.8. Mean and ±SEM usage statistics for both Dashboard variants. For all aspects, usage was significantly higher in the compare condition than the no compare condition (*p*<.01).

4.3.5 Behaviour Change Observed for Tendencies & Techniques

To assess whether or not participants' experiences with the Search Dashboard led them to change aspects of their behaviour, search behaviour *pre* using the Dashboard was compared to behaviour *post*. For this analysis the differences between participants in the Dashboard Variant conditions were examined for each of the data elements presented in the Dashboard. A control group was also created from the browser logs to provide a baseline for each of our metrics and to assess whether there were external factors (such as a new search engine feature) that could have caused any observable behaviour change during the study period.

To build the control group, 1000 searchers were separately sampled for each variable, using the log data described earlier for a time period coinciding with the study. Separate samples for each variable were used to best resemble the mean and variance of that variable across all of our participants' pre-dashboard behaviour. This was done because a simple random sample of searchers would not yield a set of searchers who were sufficiently representative of participants' behaviours and search expertise across all dimensions. Only searchers who performed the actions of a measure were used in analysis (e.g., only searchers who clicked on a result were used to compute time to click).

The means and standard deviations for all measures across all groups and time periods are presented in Table 4.1. Mixed-design ANOVAs were used with group (no comparison, comparison, and control – independent measures), and time (pre dashboard and post dashboard – repeated measures), as the factors. The analysis revealed no significant difference between the experimental groups on any of the metrics for the *pre* dashboard period. There were significant differences between groups in some of the metrics *post* using the Dashboard (all $F(2,560\text{-}1100) \geq 6.24$, $p \leq .002$), where *560-1100* denotes the minimum and maximum degrees of freedom in the error term, based on the total number of searchers, across all groups, who performed the actions for each dependent variable and the number in the control group.

Participants in the comparison group changed Tendencies – taking longer to click on search results and issuing longer search queries – and Techniques – using operators, answers, and verticals more frequently – as compared to the other groups. There was a significant effect of time on behaviour within the comparison group for the same Tendencies and Techniques (all $F(1,561\text{-}1101) \geq 10.94$, $p \leq .001$), and significant interactions for the same variables between group and time for the comparison group (all $F(2,558\text{-}1098) \geq 7.01$, $p \leq .001$). The findings suggest that participants in the comparison group significantly changed aspects of their behaviour after seeing their Dashboard.

Table 4.1. The 11 metrics for the 12 data elements presented pre and post visiting the Dashboard. Significant differences were found *within subject* and *between groups* for the same variables (bolded). For these variables comparison led to significantly higher usage.

	measure	time	control	no compare	compare
Techniques	***Operator use***	pre	3.5% (1.1)	3.7% (1.2)	3.2% (0.9)
		post	3.4% (0.8)	3.3% (1.4)	***7.3% (1.0)***
	Vertical use	pre	2.98 (0.3)	2.97 (0.3)	3.01 (0.3)
		post	2.97 (0.3)	3.12 (0.4)	***3.74 (0.4)***
	Answer use	pre	31.2% (10.0)	30.1% (8.0)	29.9% (8.0)
		post	30.9% (8.0)	31.1% (7.0)	***33.1% (9.0)***
Tendencies	***Query length***	pre	2.91 (0.6)	2.91 (0.4)	2.93 (0.4)
		post	2.95 (0.50)	2.94 (0.4)	***3.07 (0.3)***
	Click entropy	pre	1.51 (0.7)	1.48 (0.7)	1.50 (0.7)
		post	1.50 (0.8)	1.51 (0.7)	1.49 (0.7)
	Clicks per query	pre	0.51 (0.2)	0.50 (0.2)	0.51 (0.2)
		post	0.52 (0.2)	0.51 (0.2)	0.51 (0.2)
	Time to click	pre	16.39 (2.4)	16.52 (3.3)	16.31 (3.11)
		post	16.41 (3.0)	16.43 (3.9)	***17.88 (3.1)***
	Session time	pre	13.6 (8.4)	13.4 (8.5)	13.8 (8.3)
		post	13.9 (8.4)	13.3 (8.3)	13.4 (8.3)
	Session queries	pre	1.42 (1.1)	1.39 (1.0)	1.43 (1.0)
		post	1.48 (1.1)	1.42 (1.0)	1.45 (1.1)
Topics	Domains	pre	85.40 (50.3)	85.29 (58.8)	86.44 (57.4)
		post	86.11 (49.5)	86.12 (59.2)	87.22 (59.2)
	Query terms	pre	77.67 (50.2)	77.93 (50.3)	77.11 (50.2)
		post	80.13 (51.4)	78.04 (49.6)	78.26 (50.3)

4.4 Summary of Results

The study of 90 peoples' experiences with the Search Dashboard provides strong evidence that social feedback allows global learning by helping people learn new skills and technique. Participants in the study demonstrated global learning through their adoption of behaviour that was more in line with that of expert searchers, and their application of their new skills and technique in their real world information-seeking tasks.

The study of the Dashboard provides the following five further results:

- Simply reflecting on personal search engine interaction history with the Search Dashboard, led people to change their attitudes about search engines and about their own search skills.
- Participants who reflected on personal interaction history did not change their search behaviour.
- Participants changed their behaviour for 5 of the 12 data elements presented, but only when comparison data was available.
- Participants preferred, and changed behaviour for, data on Techniques and Tendencies but not Topics; likely due to a strong preference for data the can be easily applied.
- Comparison data also increased user engagement, changes in attitudes, and insight drawn from reflection.

In the Discussion (Chapter 7) I provide explanations for the results from the Search Dashboard, discuss the limitations of the study, and provide considerations for deployment in the real world.

CHAPTER 5
LOCAL SOCIAL FEEDBACK: WEBWEAR

Many loosely-coupled information-seeking tasks on the Web involve local learning. Tasks that involve local learning are set in a social context based on the knowledge and activities of people around us, where we make use of this context to simplify and improve our information-seeking behaviour. For example, we may visit websites that others have mentioned during a conversation, or try a particular social media service because a colleague uses it, or choose subtasks in collaborative work based on what we know of others' current activities.

In these situations, we communicate with tightly-knit groups of contacts (friends, family, or colleagues) to seek advice and obtain the knowledge we need in order to most effectively complete our tasks. This is local learning, where we learn how to proceed from the guidance of others to accomplish specific tasks. This knowledge is not always essential (that is, the work is loosely-coupled and can often continue even without the other person's input [68]), but, when available, it can provide substantial benefits in efficiency and effectiveness.

However, people have few tight-knit contacts, so contacts are not always available to provide information when it's needed – people may be physically away, or too busy to interrupt. In these situations, we lose the benefit of others' experience and guidance. The unavailability of contacts leads to collaborative work being less efficient or effective.

Although CSCW researchers have studied several aspects of web-based collaboration, existing solutions are still insufficient to solve the problems described here (the following points were summarized in more detail in section 2.6.5). First, collaborative information seeking systems (e.g., [66,104,134]) are traditionally oriented towards more tightly-coupled and focused collaboration, and the explicit-collaboration approach often does not match the loose coupling of web tasks [4]. Second, social navigation systems (e.g., [45,156,167]) have traditionally taken two general approaches. One approach, broad social navigation, provides a view of behaviour that is much broader than that of a tightly-knit group, leading to information that is difficult to interpret because it cannot be contextualized through personal knowledge of the provider. The other approach, narrow social navigation, does offer contextual details from tightly knit contacts, but is limited in scope because support is not offered for all of the basic information seeking activities.

I have developed a third approach based on social feedback – called WebWear – that fits between these existing paths: WebWear focuses on support for loosely-coupled information-seeking tasks in which a tightly-knit group of contacts can provide valuable knowledge and advice. WebWear uses a basic idea of social navigation – collecting traces of activity from other seekers – but only gathers information from people within the tightly-knit group, allowing a seeker to better interpret the traces based on knowledge of the people in the group, their preferences and habits, and their current activities. Further, it provides this support for each of the most common information seeking activities (reviewing webpages, selecting links, and using search engines).

I have developed a working WebWear system that gathers, stores, and visualizes traces of shared web activity (Figure 5.1). WebWear works in a standard browser, and passively gathers and displays interaction history – which can support loosely-coupled information-seeking tasks. To test the effectiveness of this approach, a user study of WebWear was conducted on the Web. Participants carried out realistic information-seeking tasks in a hypothetical small-workgroup setting. WebWear was compared with a broad social-navigation system and a standard browser. The findings showed that WebWear led to significant improvements in task success, seeker confidence, and preference, without requiring additional effort. These results show the potential benefit that can be gained from social feedback for local learning: when information comes from known contacts, navigation traces become more interpretable and more useful. Overall, the results of the user study show that systems based on social feedback can provide opportunities for local learning, and allows common LCIS tasks to be completed more efficiently and effectively.

5.1 WebWear: Designing Social Feedback for Local Learning

The goal of WebWear is to support common loosely-coupled information seeking that occurs within a social context. Based on the conceptual framework laid out in Section 3.3, here I describe three aspects of the design of WebWear: tasks, information, and display.

5.1.1 What tasks should be supported?

As previously described (in Section 3.3.1), a system to for supporting information seeking should provide guidance for any of the three core tasks (fact-finding, information gathering, and browsing). Because the three tasks are based on the three basic interactions (reviewing pages, selecting links, and using search engines), supporting the basic interactions would provide support for all three core tasks. In addition, previous research has identified several *strategies* that are commonly used in collaborative information seeking that are relevant to loosely-coupled scenarios:

- *Following*: purposefully repeating someone else's actions to find an object of shared interest [45].
- *Division of labour*: dividing up work so that there is reduced duplication of effort [104].
- *Seeding*: traces left purposefully for discovery by a collaborator (e.g., [120])
- *Serendipity:* unexpected encounters with traces that provide useful information [156]. [13]

Direct communication is another strategy that is also a common part of loosely-coupled work; however, in WebWear, and social feedback more generally, we focus on implicit information from interaction history rather than explicit communication.

5.1.2 What information do people need in these tasks?

Much of what people need for loosely-coupled information tasks has been examined in research on group awareness (see Section 2.5.2). Providing the appropriate level of awareness in collaborative information seeking has been shown to facilitate collaboration [134]; as described above, current social navigation systems may provide information at too high a level for LCIS tasks that are based in a social context. Therefore, supporting information-seeking in tightly-knit groups involves showing traces from group members, at all stages of a person's information seeking activities.

5.1.3 How can the information be displayed?

5.1.3.1 Design Constraints

Support for loosely-coupled collaboration must be low cost (both to gather and interpret), and should fit within existing work practices, due to the discretionary nature of loosely-coupled

---

[13] While serendipity may not be applied willfully, as one might commonly expect with a strategy, one must be willing to pursue traces of interest when they are encountered unexpectedly.

information seeking. Thus, WebWear must work within established interactions based on webpages, web browsers, and common tools such as search engines.

5.1.3.2 Design Factors

Location: External versus Internal to the page. I chose an internal approach to allow better integration with existing tools such as search engines. Further, in many instance no trace information would be available, using the internal approach allows seekers to fall back to using webpages as they normally would.

Style of Augmentation: Artifact Encoding vs. Glyph. Trace information can be provided with new graphical objects (glyphs) placed near the artifacts of interest (i.e., pages and links) [45], or information can be encoded into the representation of the artifacts themselves (e.g., changing the font size of links [116]). I decided to use glyphs to encode the information: although glyphs can occlude information on the page, they can be consistently displayed without altering the style and functionality of the website.

Information to Display. I aimed to display only information relevant to making choices in the information seeking process [167]; therefore, I included information about the page currently visited, the links that can be chosen on the page, and, when search is available, information that supports query formulation. I show two basic elements in the visualizations: whether or not a contact has visited a particular page, and a coarse indicator of interest in that page.

5.1.4 WebWear

The WebWear system (Figure 5.1) is an add-on for the Chrome browser that collects and adds trace visualizations to web pages (a video figure link is available on page xvii). WebWear provides glyphs that encode the number of contacts who visited a page or a link, and the recency of the visit. If no glyph is present, it means that none of the seekers' contacts have visited the page.

Page-Level Glyphs. The basic glyph used in WebWear is a small colored bar; our informal piloting (alternative glyphs displayed on paper) showed that these were easy to learn and could encode two variables that were easily discernible. Bars representing visits to the current page are colored purple. The fill of the bar represents the number of unique contacts who visited a page (see Figure 5.2). The bars fade away over time to indicate the recency of the last visit. The bar is darkest if the visit occurred in the last few days, and fades to its lightest if the visit was over a month old (see Figure 5.2). Glyphs are (8x10 pixels) for page-level visitation and placed at the right of each link (Figure 5.1). A larger version of the glyph (300x15 pixels) is shown at the top of the page to indicate visitation to the page currently being viewed (Figure 5.1).

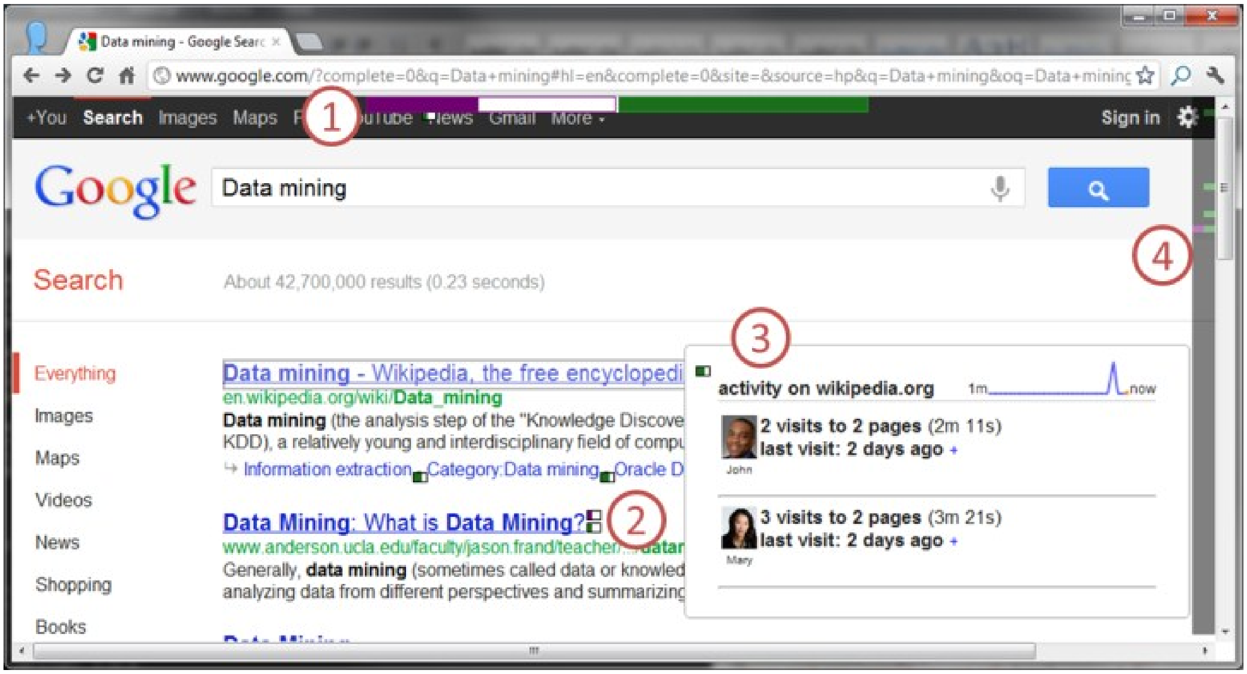
Data mining - Google Searc
www.google.com/?complete=0&q=Data+mining#hl=en&complete=0&site=&source=hp&q=Data+mining&oq=Data+mininc
+You
Search
Images
Maps
YouTube
News
Gmail
More
Sign in
1
Google
Data mining
4
Search
About 42,700,000 results (0.23 seconds)
Everything
Images
Maps
Videos
News
Shopping
Books
Data mining - Wikipedia, the free encyclopedi
en.wikipedia.org/wiki/Data_mining
Information extraction
Category:Data mining
Oracle D
Data Mining: What is Data Mining?
2
www.anderson.ucla.edu/faculty/jason.frand/teacher/...datar
3
activity on wikipedia.org
1m
now
John
Mary

Data mining (the analysis step of the "Knowledge Discove KDD), a relatively young and interdisciplinary field of compu

Generally, data mining (sometimes called data or knowled analyzing data from different perspectives and summarizing

2 visits to 2 pages (2m 11s)
last visit: 2 days ago +

3 visits to 2 pages (3m 21s)
last visit: 2 days ago +

Figure 5.1. WebWear being displayed on a search results page for “Data Mining”. 1) Glyph presenting visitation to the current page (purple) and website (green). 2) Activity glyphs on a particular link. 3) Details of activity that appears from hovering over a glyph. 4) Overview, providing the location of glyphs on the page.

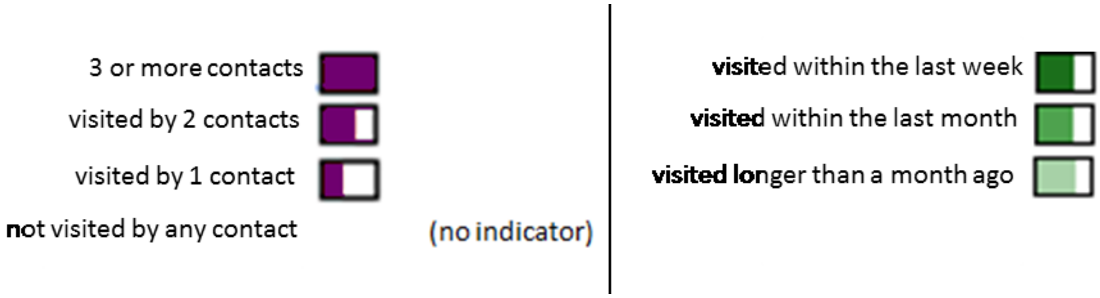
3 or more contacts
visited by 2 contacts
visited by 1 contact
not visited by any contact
(no indicator)
visited within the last week
visited within the last month
visited longer than a month ago

Figure 5.2. (left) The visual glyphs displayed beside a link indicating the number of unique visitors to the specific webpage that the link leads to. (right) The glyphs displayed beside a link indicating the timeframe for the most recent visit to the website the link leads to.

Website-Level Glyphs (Increasing Encounters). To increase the likelihood of encountering useful information, WebWear provides a second indicator that aggregates the activity on all pages visited on a given website. Website-level aggregation was chosen because visits to a particular website are likely to have some tasks or information needs in common. We used the same visual encoding as the page-level glyphs, but colored these bars green (Figure 5.1).

Details-on-Demand. All glyphs are interactive, and can provide details around the represented activity. Mousing over a glyph produces a popup that shows a recency-ordered list of visitors, and an overview of their visitation for the page or website. The popup displays the total number visits on the website, the number of pages visited, the total amount of time spent, and when the last visit occurred. A list of the occurrence and duration of each individual visit can also be displayed. Popups relating to website visits, include a recency-based list of pages visited on that site. Further, a sparkline visualization in the popup allows seekers to get a sense of the activity of all contacts visits to a page or website in the last month (see Figure 5.3).

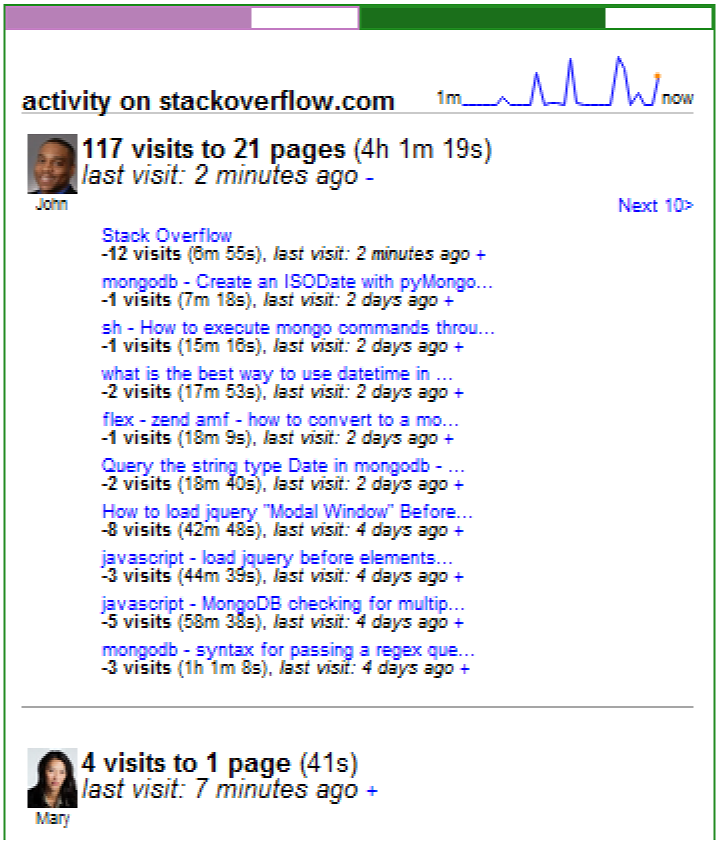

activity on stackoverflow.com

117 visits to 21 pages (4h 1m 19s)
last visit: 2 minutes ago -

Stack Overflow
-**12 visits** (6m 55s), *last visit: 2 minutes ago* +
mongodb - Create an ISODate with pyMongo...
-**1 visits** (7m 18s), *last visit: 2 days ago* +
sh - How to execute mongo commands throu...
-**1 visits** (15m 16s), *last visit: 2 days ago* +
what is the best way to use datetime in ...
-**2 visits** (17m 53s), *last visit: 2 days ago* +
flex - zend amf - how to convert to a mo...
-**1 visits** (18m 9s), *last visit: 2 days ago* +
Query the string type Date in mongodb - ...
-**2 visits** (18m 40s), *last visit: 2 days ago* +
How to load jquery "Modal Window" Before...
-**8 visits** (42m 48s), *last visit: 4 days ago* +
javascript - load jquery before elements...
-**3 visits** (44m 39s), *last visit: 4 days ago* +
javascript - MongoDB checking for multip...
-**5 visits** (58m 38s), *last visit: 4 days ago* +
mongodb - syntax for passing a regex que...
-**3 visits** (1h 1m 8s), *last visit: 4 days ago* +

4 visits to 1 page (41s)
last visit: 7 minutes ago +

Figure 5.3. The details popup for activity on the wikipedia.org website. Shown are details of all contacts' visits to the Wikipedia.org website.

Overview. Because some pages are long, and seekers may not be able to see all glyphs at once, we added an overview display showing the vertical position of all glyphs on the page. A grey overlay next to the scrollbar displays purple and green marks indicating the vertical position of page-level and website-level glyphs. Clicking any of the marks scrolls to that glyph (see Figure 5.1)

Contact Filtering. Seekers can filter the traces by selecting and de-selecting contacts in the WebWear control menu located on the browser toolbar (see Figure 5.4).

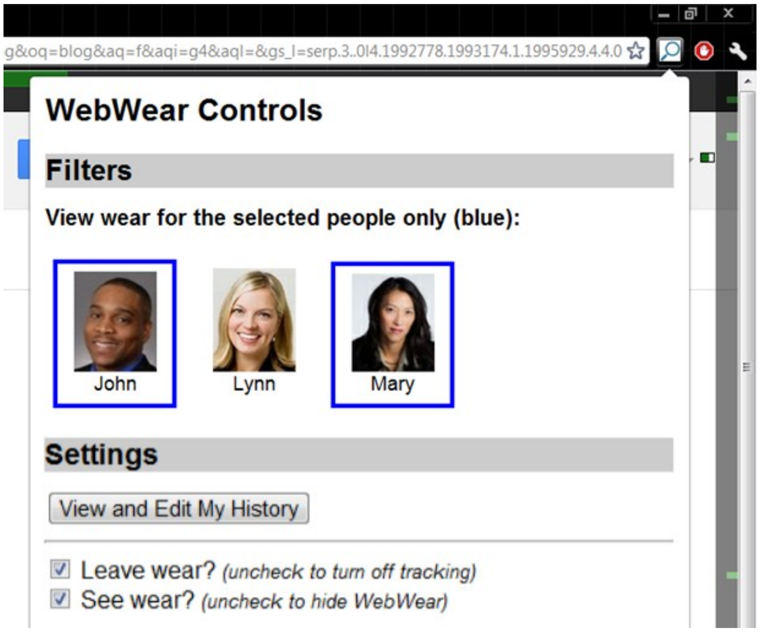

g&oq=blog&aq=f&aqi=g4&aql=&gs_l=serp.3..0l4.1992778.1993174.1.1995929.4.4.0

WebWear Controls

Filters

View wear for the selected people only (blue):

Settings

☑ Leave wear? *(uncheck to turn off tracking)*
☑ See wear? *(uncheck to hide WebWear)*

Figure 5.4. The WebWear Controls Menu, which provides access to filtering the traces presented by user (currently the WebWear will present traces for John and Mary, but not Lynn). WebWear users can view and selectively delete any history collected by the system. Users can also selectively change whether the system is collecting or displaying traces.

Group-based Query Completion. For any page that uses a standards-based query field (i.e., with an HTML name attribute value of “q”), WebWear also provides a recency-based query completion list based on the query history of contacts (Figure 5.5). The group-query completion overrides existing query completion that is offered by many search engines. Clicking on the search field displays the most recent queries executed by all contacts, providing a simple way to see the activity of all contacts. Letters typed in the search field act as a filter in the list. The query completion works effectively on Google, Bing, and Yahoo, and also in site-specific search fields for sites such as stackoverflow.com. Group search activity is shared across all search fields, regardless of the site where it came from (i.e., group search activity that was created on Google, is visible on Bing or Yahoo).

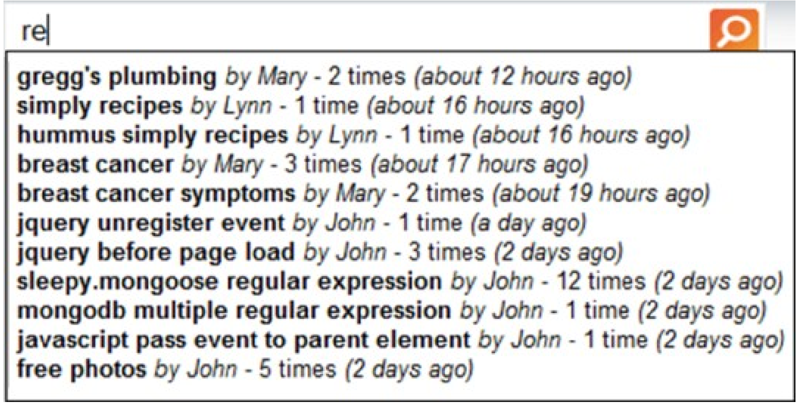

re

gregg's plumbing by Mary - 2 times (about 12 hours ago)
simply recipes by Lynn - 1 time (about 16 hours ago)
hummus simply recipes by Lynn - 1 time (about 16 hours ago)
breast cancer by Mary - 3 times (about 17 hours ago)
breast cancer symptoms by Mary - 2 times (about 19 hours ago)
jquery unregister event by John - 1 time (a day ago)
jquery before page load by John - 3 times (2 days ago)
sleepy.mongoose regular expression by John - 12 times (2 days ago)
mongodb multiple regular expression by John - 1 time (2 days ago)
javascript pass event to parent element by John - 1 time (2 days ago)
free photos by John - 5 times (2 days ago)

Figure 5.5. Group-based query completion displayed on Bing; a recency-based based on the queries of contacts.

Controlling WebWear. Because seekers may want to turn off the sharing of traces in particular situations, WebWear can be turned off through checkboxes located in the WebWear menu (see Figure 5.4). Seekers can selectively turn off sharing traces and seeing others' traces. The system is turned off completely when the browser is placed in privacy mode. If sharing is enabled, WebWear collects webpage visits for the current seeker, and stores them on a remote server.

5.2 Empirical User Study of WebWear

To evaluate the primary hypothesis – that social feedback from tightly-knit groups can improve success in information seeking – we carried out a user study that asked participants to perform web-based information tasks in a variety of different conditions. The study was designed to answer two main questions:

- Whether the social feedback in WebWear improves task success (effectiveness and efficiency) compared to either no history, or broad social-navigation information; and,
- Whether the WebWear visualizations (as described above) are easy to access and interpret for users, when used in a real web browser and a realistic task context.

5.2.1 Interface conditions

We tested three different interfaces in the study: a normal browser, a broad social-navigation display, and WebWear.

- *Normal browser* was an unchanged version of the Chrome web browser (version 19.0.1084.52).
- *Broad social navigation (BroadNav)* is a visualization of link popularity that was similar to the simple bar visualizations used in WebWear, and appeared on all pages (Figure 5.6). Popularity information was simulated with Google PageRank data (obtained by querying the Google Toolbar webservice for each link). PageRank approximates the likelihood of randomly visiting a page. While it is only weakly correlated with real world page visitation [97], as an approximation it provides a reasonable stand-in for true link visitation counts.
- *WebWear* was implemented as described above (see Figure 5.1). Information traces were pre-determined for the tasks used in the study; traces were pre-recorded and made visible for three members of the participant's task workgroup (described below).

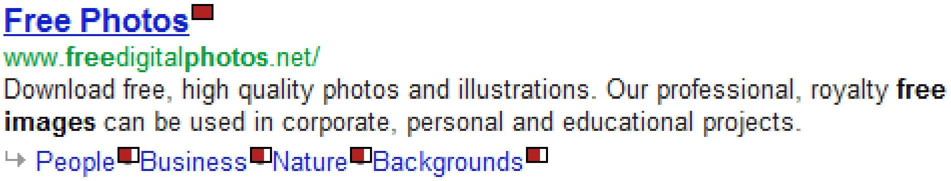

Free Photos

www.freedigitalphotos.net/

Download free, high quality photos and illustrations. Our professional, royalty **free** **images** can be used in corporate, personal and educational projects.

Figure 5.6. Glyphs as displayed in the BroadNav condition of the study. Mousing over the glyphs in the BroadNav system did not provide any additional information.

5.2.2 Tasks and scenarios

For the experiment, participants were asked to imagine working with a hypothetical group, which included: Mary, a close friend of the participant; John, a friendly colleague and a computer programmer; and, Lynn, the group's manager and mentor. Participants carried out two kinds of tasks that were designed to employ two different types of collaborative strategies (described above), both of which are common to real-world workgroups. First, *follow* tasks involved following up on a lead or comment provided by a member of the person's scenario workgroup, and finding a specific article or site that had been mentioned or discussed by the other person. In follow tasks, we were interested in whether participants could successfully find specific items that had been partially identified by another person. For example, one of the follow task descriptions was:

> *John tells you that he saw a great book of bird photographs on Amazon that the office should buy for Mary's birthday. Mary has a lot of bird books, but John says he knows she doesn't have any 'coffee table' books like this, and it only costs about $38. You are in charge of buying the present, so you decide to try and find the book John had seen.*

Second, *division* tasks involved completion of a group task, where the workgroup investigates a particular topic or area; in these tasks there were no specific answers (as in follow tasks), but rather the group's job was to cover the area efficiently. In division tasks, we were interested in whether people would be able to organize and coordinate the task, and whether people would be able to avoid duplication of work. For example, one division task description was:

> *Your colleague Mary has a history of breast cancer in her family. Over lunch last week, you decided to research breast cancer symptoms together. You are trying to get a wide coverage from different sources, so you are trying to each collect URLs for three introductory articles that mention common breast cancer symptoms. You plan to look at these together tomorrow. Of course, it would be useful to find different articles than Mary does.*

These tasks were presented in the context of a hypothetical scenario in which the participant was a member of a workgroup, and that they had a set of tasks on a to-do list that they wanted to finish at the end of the day. The scenario indicated that the other members of the workgroup were not available for direct communication, meaning that the participant had to complete the tasks using only the information available from the scenario and the interface (a complete list of the study tasks and the task scenario can be found in Appendix A).

Two minor constraints were imposed on participants: first, they were asked to use only a single tab to reduce the variability in behaviour between participants; second, they were asked to start each task from the Google search page.

5.2.3 Participants and Apparatus

Twenty-four participants (14 women, 10 men, *mean age*=26.13 years, *std.dev.*=7.67, *max*=37 years, *min*=19 years) were recruited from a local university. Eighteen were students from a range of disciplines, others were researchers (5) or software engineers(1), and none were familiar with the research project. Participants all reported searching the Web at least once every day and spending at least 2 hours/day using a computer, during a normal workday.

The study used a standard version of the Chrome browser, running on a recent Windows 7 PC, and a 1920x1080 LCD monitor. The WebWear condition used the actual system as described above, but with small modifications. For example, simulated traces were added for the workgroup, which included activities that contained overlapping and non-overlapping traces to the experimental tasks; and, gathering of new interaction history was turned off.

5.2.4 Design and Procedure

The study used a 3x2 within-participants design with two factors: *Interface* (*normal* browser, *BroadNav* for broad social navigation, or *WebWear*) and *Task* Type (*follow* or *divide*). Participants carried out six tasks in total: both task types in each of the three interfaces.

Participants first completed a demographics survey, and were given a handout that introduced them to the hypothetical scenario and collaborators. They were then given a practice task using the normal browser. They next completed one of the three interface condition blocks (balanced with a Latin square), which included both a *follow* and a *divide* task (alternated ordering between participants). Before each of the BroadNav and WebWear interface blocks, participants were given a short tutorial (approximately 5 minutes), where they were introduced to the new interface and asked to perform a random query to show how the system displayed information. They were then given an additional practice task. The experimenter interacted with the participant during the practice tasks and pointed out the displayed information in the context of the task. Participants then completed the two experimental tasks. After each task they completed a short questionnaire concerning their experiences, and after the WebWear and

BroadNav interfaces they completed a second questionnaire that asked their opinions of the interface. At the end of the session, participants were asked about their preferences.

The study system gathered two types of data for each task. *Performance data* included completion time for the task and an accuracy score for their final task result (participants were given 1 point if they correctly identified the webpage in follow tasks, and 1 point if they did not duplicate any answers in divide tasks). *Subjective data* came from the three questionnaires, which were composed of 7-point Likert-scale statements and free-text questions (and can be found in Appendix A).

5.2.5 Data Analysis

Performance data was analyzed using RM-ANOVA, and post-hoc tests used Bonferroni corrections for multiple comparisons. Because the accuracy score was a binary value, I initially transformed the accuracy score data using the Align Rank Transform [169]. The Aligned Rank Transform is a method to deal with such study designs and allows parametric tests (i.e., ANOVAs) to be conducted on non-parametric data. Subjective responses were analyzed using Friedman's test, and post hoc comparisons were done using the Wilcoxon signed-rank test. However, because Friedman's is a one-way test only, I split the task-related questionnaire data by task type, and looked for differences between interfaces under the different task conditions. I did not look at the differences in task type data because I was mainly interested in participants' subjective responses to interface, and not task type.

5.3 Results

I present the results organized by the main hypotheses about effectiveness and efficiency in completing the tasks. I first present performance data and subjective responses, and then report on experiences and preferences.

5.3.1 Were participants more effective with WebWear?

Participants completed their work more effectively when using WebWear (see Figure 5.7). There was main effect of Interface on accuracy score ($F_{2,46}$=13.76, p<.001) and task type ($F_{1,23}$=6.29, p<.05) but there was no interaction effect between interface and task type ($F_{2,46}$=2.11, p>.05). Pairwise comparisons show that participants were significantly more accurate in completing the task with WebWear than with either BroadNav ($p<.005$) or the normal browser ($p<.001$). There was no statistically significant difference for accuracy between BroadNav and normal ($p>.05$). Pairwise comparisons showed that participants were significantly more accurate in completing the divide tasks than the follow tasks ($p<.05$).

Participants also felt more certain about having successfully completed the task when using WebWear. In follow tasks, participants felt more certain about having found the correct webpage (Figure 5.7). In divide tasks, participants felt more certain about having not duplicated any answers (Figure 5.7). There was a significant effect of Interface on certainty about divide tasks ($\chi^2(2)=39.78$, $p<.001$). Pairwise comparison show that participants were more certain with WebWear than with BroadNav ($z=-2.94$, $p<.005$) and normal ($z=-3.63$, $p<.001$). Participants were also significantly more certain with BroadNav than normal ($z=-2.27$, $p<.05$).

There was a significant effect of Interface on certainty for follow tasks ($\chi^2(2)=18.25$, $p<.001$); pairwise comparisons show that participants were significantly more certain with WebWear than BroadNav ($z=-4.07$, $p<.001$) or normal ($z=-3.32$, $p<.001$). Participants were more confident with BroadNav than normal ($z=-2.78$, $p\leq.005$).

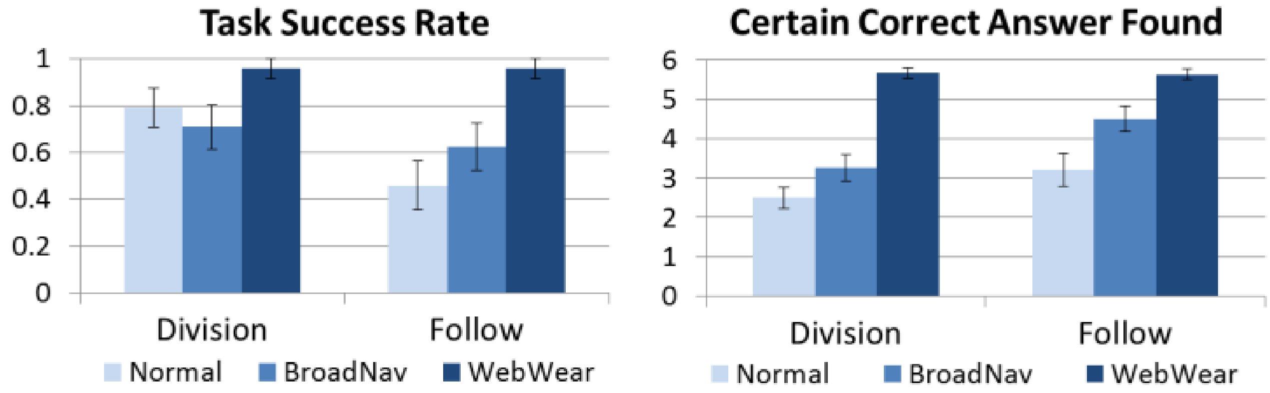

Figure 5.7. (Left) Task success rate.
(Right) Subjective certainty about task success.

5.3.2 Were participants more efficient when using WebWear?

To consider efficiency, we looked at the time taken to complete the task, subjective perceptions of efficiency, and subjective need for communication with a collaborator.

Completion times are shown in Figure 5.8. There was a main effect of Interface on completion time ($F_{2,46}=20.89$, $p<.001$). Pairwise comparisons show that BroadNav was slower than both WebWear (p<.001) and normal (p<.001), but that there was no difference observed between WebWear and normal (p>.05). There was no effect of task type on completion time ($F_{1,23}$=.32, p>.05), with divide tasks (*mean*=2.9 min., *SEM*=.19) taking roughly the same amount of time as follow tasks (*mean*=2.7 min., *SEM*=.35). There was no interaction effect observed ($F_{2,46}$=.56, p>.05).

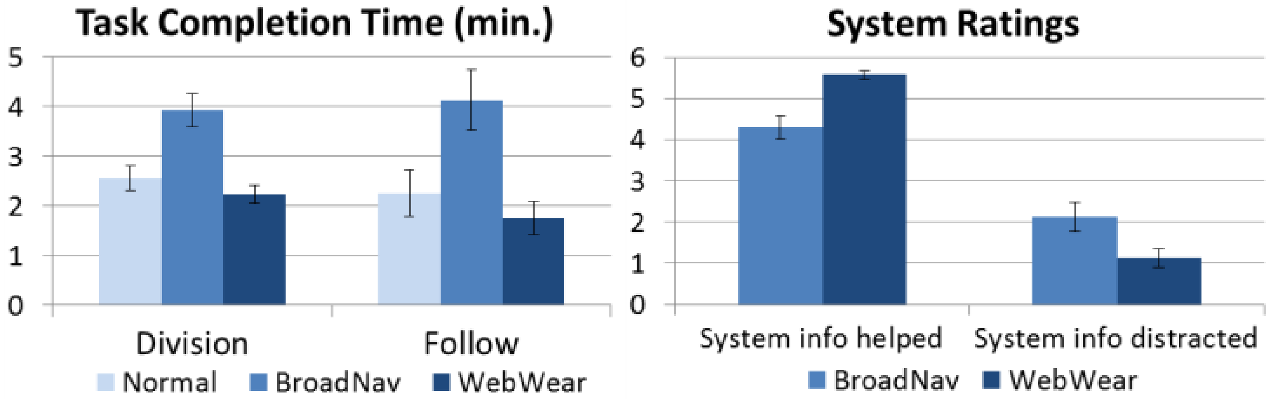

Figure 5.8. (Left) Average task completion time. (Right) Participant ratings of the system being helpful and distracting.

Participants felt they were most efficient when using WebWear (see Figure 5.9). There was a significant effect of Interface for both divide tasks ($\chi^2(2)$= 10.84, p<.005) and for follow tasks ($\chi^2(2)$= 13.92, p<.001). For divide tasks, participants felt they were more efficient with WebWear than both BroadNav (z=-2.65, p<.01) and normal (z=-2.81, $p \leq$.005). There were no differences between Normal and BroadNav (z=-.84, p>.05). For follow tasks, participants again felt they were more efficient with WebWear than both BroadNav (z=-3.39, $p \leq$.001) and normal (z=-3.21, $p \leq$.001), but there was no observed difference between Normal and BroadNav (z=-.24, p>.05).

Participants felt less need for communication when using WebWear (see Figure 5.9). There was an effect of Interface for both divide tasks ($\chi^2(2)$= 25.42, p<.001), and for follow tasks ($\chi^2(2)$=11.17, p<.005). For division tasks, the perceived need to communicate was lower for WebWear than both BroadNav (z=-3.31, $p \leq$.001) and normal (z=-4.01, p<.001). Participants also felt they needed to communicate less with BroadNav than normal (z=-2.07, p<.05). For follow tasks, again the perceived need to communicate was lower for WebWear than both BroadNav (z=-2.86, p<.005) and normal (z=-3.84, p<.001), but there was no difference between BroadNav and normal (z=-.68, p>.05).

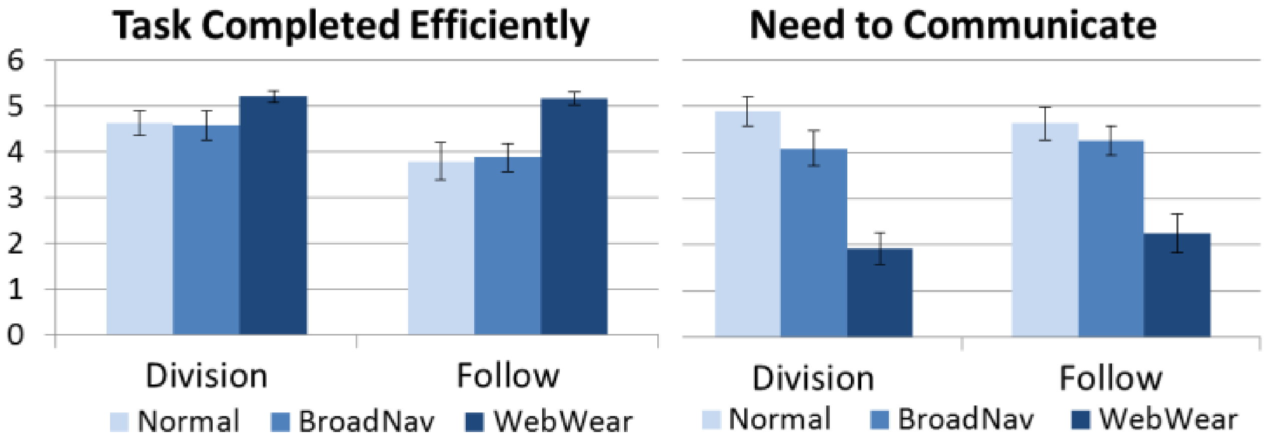
Task Completed Efficiently

Need to Communicate

Figure 5.9. (Left) Participant agreement with having completed the task efficiently. (Right) Ratings of needing to communicate to be successful in the task.

<u>5.3.3 Perceived Usefulness of WebWear and BroadNav</u>

Participants felt that the information available in both WebWear and BroadNav was useful in completing the tasks (Figure 5.8), but they rated the information in WebWear as being significantly more helpful (*z*=-3.447, *p*<.001). Participants also felt that the information displays in both BroadNav and WebWear did not get in the way or distract participants from their work (see Figure 5.8); however, they felt that WebWear was significantly less distracting than BroadNav (*z*= -2,55, *p*<.05). Finally, 22 of the 24 participants felt that WebWear was the best tool for the type of tasks seen in the study, and 22 agreed that they would use WebWear for information seeking activities (*mean agreement*=5.13, *SEM*=0.26, *max*=6, *min*=0).

5.4 Summary of Results

The empirical user study provides three main results:

1. *WebWear allowed people to be more effective*: People were significantly more accurate in finding answers with WebWear than with either a normal browser or with a broad social-navigation system, and were significantly more confident in their answers.
2. *WebWear allowed people to be efficient:* The increased accuracy did not come at the cost of additional effort – tasks were completed faster in WebWear than in BroadNav and no slower than using the normal browser. Participants also felt

they were most efficient with WebWear and reported a reduced need to communicate compared to the other systems.

3. *WebWear was preferred:* People subjectively felt that WebWear was most appropriate for the task, that the information it presented was helpful and did not distract from their tasks, and that they would like to use it in appropriate contexts.

In the Discussion (Chapter 7), I provide explanations for these main results and discuss limitations of the findings and issues in the real-world deployment of WebWear.

CHAPTER 6
REAL WORLD CHALLENGES FOR SOCIAL FEEDBACK

Social feedback can provide support for loosely-coupled information seeking by allowing seekers to learn new skills (Chapter 4) and improve their efficiency and effectiveness in completing many common tasks (Chapter 5). The solution of social feedback relies on interaction history to represent behaviour that can be shared between information seekers. However, there remain several questions about whether the approach can scale to real-world usage. There are two outstanding challenges with regards to the use of interaction history that need to be better understood: *the overlap challenge* and *the privacy challenge*.

First, the *overlap challenge* can occur in social feedback systems because individual information seekers may not share similar tasks, may not be willing to share their interaction history, or may visit the same webpages too infrequently. These three issues mean that when support is needed that social feedback may be available too infrequently or not at all. The main reason that the overlap challenge might occur is the breadth and depth of the Web – that is, an information seeker could potentially spend most of their time visiting unique webpages and performing distinct tasks. This is of particular concern to social feedback systems based on local learning (such as WebWear) which provide support based on a small group. A better understanding is needed of how realistic it might be for social feedback systems to rely on overlap to occur in a way that allows support and guidance to be available.

Second, the *privacy challenge* can occur in social feedback systems because people’s web activity encompasses many personal and possibly private details. Social feedback systems rely on people being willing to share their interaction history for others to learn from, but many seekers may be unwilling to share their interaction history due to concerns about what others may think of their Web activity. However, in certain situations people are very willing to share their activities, and actively do so (e.g., social networking sites such as Twitter[14] or foursquare[15]). A better understanding is needed about the contextual details of information-seeking tasks and how they may relate to people’s real-world privacy concerns. If privacy concerns could be understood and reduced, people would be more willing to readily share their interaction history.

---

[14] http://twitter.com

[15] http://foursquare.com

To provide further insight into the scalability of social feedback to real-world conditions, I conducted two studies that examine these challenges in two separate ways: a *field trial* of the WebWear system, and a *diary study* of the contextual details around people's information seeking activity.

First, I conducted a week-long field trial of the WebWear system with a 16-member workgroup. The field trial showed that the system effectively provided many opportunities to encounter the traces of colleagues in day-to-day information seeking practices. Further, the field trial also showed that traces actually provided support for several real LCIS tasks. However, though participants did not report any bad experiences from having traces of their interaction history shared with colleagues, several participants still expressed some concerns about privacy.

Second, I conducted a two-week diary study that allowed participants to periodically describe the context of the information-seeking task they were currently working on, and provide ratings describing their privacy concerns. The study uncovered previously uninvestigated factors that may affect or predict people's privacy concerns about sharing interaction history. The diary study shows that people are much less concerned about sharing their interaction history when they are actively engaged in work-related information-seeking tasks and that they have low levels of concern about sharing their interaction history with their colleagues. The concerns that were reported were expressed mostly in non-working tasks and by the same individuals.

Together the two studies provide further evidence that social feedback is a feasible and appropriate solution for supporting real world information-seeking tasks.

6.1 Field Trial of the WebWear System

I conducted a one-week trial of WebWear where it was used by 16 members of a real workgroup as they carried out their daily tasks. The study explores the challenges that arise if pages and task do not overlap. Additionally, the study allowed people's privacy concerns about sharing interaction history and using social feedback in real information-seeking. WebWear was selected for the field trial for two principal reasons. First, WebWear fundamentally relies on overlap to provide support – therefore, people who are using the system as part of their real tasks can report on whether or not they are able to encounter the traces of others, and, if so, whether these traces can provide any benefits. Second, people who use WebWear as part of their real tasks can report on privacy concerns that naturally occurred during information seeking. Further, because WebWear makes use of a small pool of users to contribute interaction history and because individual identities are shared, it is not at all clear that the system can overcome the privacy and overlap challenges.

The field trial focused on three main questions:

- Would people's information-seeking activity overlap, providing opportunities for WebWear to be used?
- Did encountering the traces of others while using WebWear lead to any useful experiences?
- What were peoples' impressions of sharing interaction history during real-world information seeking, and did they have major privacy concerns about sharing?

6.1.1 Participants and Procedure

Sixteen members of my research lab participated in the field trial (1 professor, 1 visiting researcher, 10 graduate students, and 4 undergraduates). Typical work activities for this group cover a wide range of HCI and computer science topics. Participants were solicited at a lab meeting, where a 10-minute tutorial about how WebWear worked was given; none had previously used or seen it.

Interested lab members enrolled, by installing WebWear on their lab machines for the one-week (four and a half workdays) of the study. The deployed WebWear system was the same as previously described (see Section 5.1.4), but additionally provided a facility for participants to view and permanently delete any of their own traces from the system. Further, participants were informed that they were free to turn off the sharing of interaction history or the presentation of traces whenever and for as long as they like.

Participants optionally completed a survey four times during the week via email, which asked four general questions about their experiences. The four free-text questions were:

1. How did you use WebWear today? For example, this might include: followed someone's "wear", learned something new, just browsed people's activity, or any other way you used it.
2. Was what you saw useful or interesting to you, or can you imagine it being potentially useful? How?

3. If you didn't notice it or use it today, was there any reason?
4. Do you have any ideas of how the system could be improved or could be more useful?

At the end of the week, I held an hour-long group interview, where participants' experiences with WebWear were explored in more detail. The interview was semi-structured, in that there was a pre-set list of twenty-two questions; however, I followed-up on interesting responses and allowed participants to describe their experiences in detail. The questions that were asked were variations based on four high-level questions:

1. How did you use WebWear and what was it useful for?
2. Were you concerned about privacy?
3. What situations do you think using WebWear would be useful and appropriate?
4. What would be your thoughts on using WebWear over a longer term?

6.1.2 Results

The results from field study are organized by three main research questions described above. The surveys, interviews and high-level usage statistics are presented together.

6.1.2.1 Was there overlap?

Participants were active during the week of the trial, sharing approximately 85 page visits[16] and 14 searches per participant per day (6140 page visits and 1022 searches in total). In the group interview all participants reported encountering traces every day, and 88% of the surveys reported actively exploring other users' traces.

Co-visitation of Pages and Websites

I also examined the occurrence of visitation overlap – that is, when two or more people actually pay a visit to the same page (represented by a URL) or to the same website (represented by a domain, such as google.ca). Overall, co-visitation of pages occurred 3% of the time (for 76 of 2583 unique URLs visited), and co-visitation of websites occurred 14% of the time (for 68 of

---

[16] A visit in the WebWear field trial system was defined as a continuous period in which a webpage was visible in the web browser. This means that webpage visits would begin with a page being loaded or a tab being switched to, and end with a new webpage being loaded or by switching to another tab. However, a webpage visit did not end if the web browser lost focus and was not the active application.

the 479 pages visited). Recall that this does not mean traces were only seen 3% of the time for pages and 14% of the time for websites. Social feedback traces are displayed on links going to the previously visited pages, which can occur anywhere, and a page typically displays many links. The actual percentage of pages where traces were displayed is unknown, but would be much higher. When traces are encountered, they allowed participants to explore one another's activities without having to leave the page or task they were currently working on. These results show that overlap does occur, and even occurred within the short period of the trial. Also, the simple technique of displaying traces at the level of website is an effective mechanism for substantially increasing the amount of direct overlap participants encounter. Further, given that participants visit so many pages, they are likely to encounter co-visitation of a page once every day or two, and visit the same website several times a day.

Co-Occurrence of Searches

WebWear displays a recency-based list of previous searches, which can be filtered by a user's query (see *Figure 5.5*). I also looked at the co-occurrence of exact search matches. Exact searches would co-occur if people happened to use the exact terms (which would be rare, even if people have the exact same information need), or if they used someone else's WebWear search traces to follow someone's past activities. Exact search duplicates occurred for only 2% of searches between individuals (in 10 of the 426 searches). This was despite the fact that one conversation in the group interview discussed how frequently the search traces in the group query completion was used (this is described in more detail below).

Roughly half (6 of 10) of the repeated searches occurred because others had followed someone else's previous search (the use of the "follow" strategy). For example, in the group interview one participant described wondering what a long string of characters was doing in the search field (the search they was "105020*4*3*4"), and clicked it to find out what it meant. This revealed to the participant that other people were using Google as a calculator.

The other 4 searches appear to have occurred by chance. These searches were all navigational [23] – that is, they were searches conducted in order to get to a specific page or website, and user were employing the search engine as a way to provide a shortcut for a longer URL. Repeated navigational searches included: 'google calendar', 'google scholar', and 'facebook'. In these cases support is not typically needed, because a searcher knows that the query will return exactly what they want: a link to a specific website.

While the co-occurrence of searches was rare, it appears that there were cases where people used it for following others. The infrequency in search overlap does not mean that the search traces provided were not useful though, as the group-query completion was used frequently by participants for maintaining an awareness of group activities (described in more detail below).

Summary of Overlap Findings

Because traces were available on common information seeking starting pages (such as google.ca, where all 16 participants visited at least 7 times), encountering traces during information-seeking tasks was common. Further, presenting website-level traces provided roughly five times the frequency of webpage overlap, meaning that direct overlap was more common. Overall, these results suggest that WebWear provided ample opportunity to encounter traces during many information-seeking tasks. However, this does not mean that the traces were necessarily useful or provided support when it was needed.

6.1.2.2 Was the overlap useful?

Trial participants found WebWear useful in four main ways:

1. *Group Awareness:* For maintaining a general awareness of the work activities of group members (see Section 2.5.2);
2. *Division of Labour:* For avoiding duplication of work (see Section 5.1.1);
3. *Following:* For coordinating actions by following the traces left by others (see Section 5.1.1);
4. *Serendipity:* For unexpected encounters with traces that provided useful information (see Section 5.1.1).

Further, participants used WebWear in one further way:

1. *Seeding:* For bringing the attention of potential collaborators to some behaviour, through purposively leaving traces for others to see and use (see Section 5.1.1)

I will discuss each of these five uses of WebWear in more detail.

The only way to access data about the usefulness of traces was through participants' subjective reports, taken in the group interview and from the daily surveys. Examples of at least one of each of the above were observed, even though it was a short one-week (four and a half workday) trial. These examples demonstrate that local learning can occur in real world environments.

Group awareness

In the group interview all participants agreed that WebWear was useful for maintaining an awareness of the activities of others. Awareness is important for providing opportunities for collaboration in the workplace (see Section *2.5.2*), and to facilitate collaborative information seeking [134]. One specific example of group awareness was reported by a participant in a daily survey, "While looking for tutorials on … XNA I did notice, thanks to WebWear, that [P12] had also been searching for this topic. I didn't even realize that [P12] was using XNA in his project. Now I guess if I run into anything major, I at least know someone who may have also been dealing with the same topic."

Five other participants reported using group-query completion as a tool to maintain awareness. Participants would simply visit a search engine as they normally would, but while they were there, they would first activate the query completion list without typing anything. This would give them an up-to-the-moment view of the most recent search activities of other group members. Three of the participants reported changing their search behaviour so that they could encounter these search traces more often – that is, rather than searching by typing a query in to the address bar in Chrome as they normally would do, they would first visit the Google homepage, so that they could see WebWear's group-query completion (Figure 5.5).

Division of Labour

Work can be completed more efficiently by dividing up tasks to allow collaborators to complete task components individually. During the short trial, there were two instances of divide tasks. An example of a division task arose when P8 asked P10 if he knew how to solve an issue with a Bluetooth dongle. P10 did not have an answer and returned to his desk, deciding only later to search for an answer to try to help P8. As P10 started his search he saw P8's queries, and selected different search terms to explore a different path, now knowing that P8 had already searched without success.

Following

Following allows people to narrow down the space of possibilities by selecting a course of action based on decisions made by others. In addition to the six occurrences of following observed during search (as described above), there was one further example of following described in the daily surveys: “[P15] had visited Dexsoft, a 3D model website, and had discussed with me the category of models that she would like me to look at. When I visited the Dexsoft website I hardly had to look around the website to find the models she was talking about because I just followed her WebWear ‘footprints’.”

Serendipity

Serendipity refers to the unexpected but useful encounters with trace information. Serendipity would have occurred in each of the previous examples above. Here the use of the traces was not planned (verified with participants through follow-up emails). The occurrence of serendipity also shows that WebWear provides appropriate support for LCIS, because people are able to conduct their tasks as they normally would, but are also able to make use of the support when it is available, thus allowing work to be completed more efficiently and effectively without additional effort.

Seeding

Seeding was also exemplified in the study; however, its use wasn’t necessarily useful, it was playful. Several participants reported leaving searches not to satisfy an information need, but to be encountered by others. The seeding of searches was mostly seen on the first day of the study, and was used by participants to test out WebWear and to see how it worked. For example, one participant who does a lot of programming searched for “what is a variable”. Once participants noticed this behaviour they began playing off one another’s queries. In another situation that followed, participants began poking fun at another one of the participants (whose name has been changed to “Will Jones” for this example).

- [p11 searched for]: what does [Will] do
- [p12 searched for]: what does [Will Smith] do
  [(note that p12 references a public figure with a name similar to Will Jones’ name)]

- [p13 and Will Jones searched for]: what does [Will Smith] do today
  [(note that p13 makes references to the same public figure as p12 had, and Will Jones repeats the same query)]
- [p14 search for]: [Will Jones] is amazing
  [(p14 makes the first reference to Will Jones' full name after viewing the other's traces)]

Again, while these examples are not necessarily useful other than for the participants to establish how the system is working and to have a playful experience, they still demonstrate seeding behaviour. Through these playful acts, participants in the study became aware of seeding behaviour as a potential use for WebWear. In the group interview participants described the use of seeding in the example above as being like an internet messaging service (IM).

When Overlap was Difficult to Use

A few participants found that in some cases there was too much overlap presented by the system. These situations arose with traces of activity for other websites (the green glyphs in Figure 5.1), and would have likely been the result of visiting a page that had many outgoing links to a common website, such as google.ca. A few participants reported that when these scenarios arose, the website-level information was less useful, because everyone visited google.ca regardless of their task. While the use of website-level traces on google.ca provided information that was too broad, there were also clear examples of where website-level interaction history was more narrow, and more interesting and useful to participants as a result. For example, 9 of 16 people had visited stackoverflow.com (a programming help website) within the trial (over only 4.5 working days), and many of these participants liked seeing what other people were looking for help on.

Another participant noted that she was spending a lot of time exploring the traces presented in WebWear rather than on completing her tasks. The participant described doing this because she was "too curious" rather than because she needed support.

Summary of the Findings for Usefulness of Overlap

The examples above illustrate real-world cases and four distinct ways in which WebWear (section *5.2*) was useful. Despite lasting a very short period, the field trial provides strong evidence that real world loosely-coupled tasks can be supported through social feedback. Further, all field trial participants agreed that there was potential value in being able to share traces, both to maintain awareness and to help coordinate work.

6.1.2.3 What were the impressions and were there privacy concerns?

Despite having mainly positive experiences, during the group discussion, many people still reported being wary of possible issues with privacy. The privacy concerns related mostly to the fear of unwanted activity accidentally being entered into the browser, or other participants misinterpreting the context of search activity. Not all participants experienced these concerns, though, and some reported forgetting about the data collection because they were not overly worried.

A different way in which privacy was a concern was expressed in the way that people would, on occasion, use language with negative connotations. In a few cases participants used terms such as “spying on” or “creeping on” others, to describe observing one another’s traces. The use of language with negative sentiment suggests that participants may have felt that they were violating other people’s privacy despite the fact that those people willingly shared their traces. Even though some people had these feelings, all participants reported that they actively checked out one another’s traces.

All participants agreed that the work context was suitable for sharing traces, and while some concerns with sharing existed, no one reported having a bad experience in using WebWear, and were open to using it in the future. Further, participants’ concerns were not serious enough that they changed sharing activities. Participants were told during their introduction to WebWear that they were free to turn off sharing or viewing traces at any time, and that they should do so, if they felt concerned about their privacy. However, participants shared a large number of visits throughout the field trial. For each full day after installation, participants shared at least 853 page visits each day. This shows that privacy concerns were not serious enough that it affected participation.

6.1.3 Summary of Results for the Field Trial

The initial field trial of WebWear in real use provides three findings:

1. There was enough overlap in people’s real-world browsing that participants were able to make use of WebWear on numerous occasions over one week;
2. Participants reported several different ways where they found WebWear information useful, both for specific loosely-coupled information-seeking tasks, and for maintaining awareness of others in the workgroup;
3. While several participants said that privacy from passive tracking remained a concern, those participants also felt that using WebWear did not unduly intrude on their privacy, and there were benefits to sharing in the work context. Further, participants concerns were not serious enough to prevent them from sharing their interaction history throughout the week.

6.2 Diary Study

During the lab study and field trial of WebWear, participants consistently expressed the view that the workplace was an appropriate context for the use of personally identifiable social feedback. However, a few still expressed concerns about privacy, despite the fact that WebWear had actually been used in the context of workplace activities. This suggests that, even though people feel the use of social feedback is appropriate, there are still contexts in which privacy concerns occur during workplace information seeking.

Perceived risks (such as those associated with privacy) can be a powerful influence on the decisions we make, and often outweigh perceived benefits (called *negativity bias*) [128]. This means that social feedback systems must minimize perceived privacy risks as much as possible, in order for people to willingly share their interaction history (which helps to address the overlap challenge). However, little is still known about the contexts factors within workplace information seeking that lead people to have higher or lower levels of concern for their privacy.

To identify what particular contextual factors may raise privacy concerns, I designed a diary study that explored people’s concerns with sharing their interaction history within the workplace. I initially identified three previously minimally explored or unexplored contextual factors that can provide important insights into people’s privacy concerns: 1) *work relatedness:* whether the task is a work or non-work task, 2) *task type:* the type of task that is being conducted, and 3) *collaboration:* whether the task is part of an established collaboration or not.

Work Relatedness: First, the Web allows us to easily perform many different tasks regardless of where we are. Therefore, it could be that many people conduct tasks that are based in their personal lives rather than work lives. In information-seeking tasks that are not related to work activities (or *non-work tasks*), people may prefer to keep the details of the task private from their colleagues. To explore if work relatedness may be connected to people's privacy concerns we asked people to classify the task they were currently working on as being either a *work* or a *non-work* task. Previous work has suggested that people believe they would be willing to share their work related web browsing with close colleagues (see Section 2.6.5), but these judgments have not been made within the context of real information-seeking tasks.

Task Type: Second, the type of information-seeking task being conducted varies greatly, and can present very different privacy concerns (e.g., online banking activity is more likely to be kept private than the weather forecast). To investigate if high-level task type may be a good predictor of privacy concerns in workplace information seeking, participants classified their current task as being one of fact-finding, information gathering, browsing, transactions or other (based on Kellar's information-seeking task types, section 2.2.1).

Collaboration: Third, previous work has suggested that collaboration in information seeking is common [4]. In collaborative situations one might expect that since the information being sought will be shared at some point, people will have fewer concerns if their activities are public. However, there is little information on whether the details about collaborative tasks are less of a privacy concern than personal tasks.

The questions in the diary study allowed me to examine six main issues with regards to the privacy concerns about sharing interaction history:

- Are there at least some webpages that people visit where they do not have important privacy concerns?
- Do people differ in how many situations they identify as being a concern?
- Do people frequently conduct non-work related tasks during their workdays, and if so, how do people's privacy concerns differ between work tasks and non-work tasks?
- Does the type of information-seeking task being conducted result in different levels of concern about sharing interaction history?

- Are information-seeking tasks that involve collaboration less of a privacy concern than individual tasks?
- Overall, what are people's specific concerns with sharing interaction history?

6.2.1 Participants

Eighteen people participated in the study (12 male, 6 female); the average age was 27.6 years (sd. 5.0), and all participants were either graduate students or undergraduate research assistants. All spent most of their working time in large open workspaces with individual work stations. 15 people reported using a computer at least 7 hours/day on typical workdays, and the remaining four at least 3 hours/day. All participants used their main work computer to complete the study, of which five were laptops that were used both at home and work. Participants were given a $20 honorarium for enrolling in the study. Fourteen of the participants were members of the research lab, while four others members of other labs.

Twelve of the participants had previously participated in the WebWear field trial, and so were aware of the privacy concerns that came up in the group interview. For these participants, it was made clear that some of the questions in the diary study related to the issues raised in the group interview; however, they were told that they should not consider their experiences with WebWear when responding to the survey.

6.2.2 Procedure

Participants were solicited through an email invitation that explained the details of the study, provided a link to a walkthrough video, and a link to download the diary study system. The walkthrough video demonstrated all stages of the study, including the installation of the study system and an introduction to the surveys. Upon installation in a participant's Chrome web browser, the study system presented a consent form and demographics survey, and once these were completed, the system would start periodically making requests to the participant to complete a diary entry.

Participants were asked to complete a total of 30 surveys within two weeks of enrollment. Each survey required 3-5 minutes to complete. The study system would display a request popup to ask participants to fill out a survey every 15 minutes, but would only make a request if the Chrome browser was currently the active application and a webpage was being displayed (see

Figure 6.1). The request popup displayed the current number of surveys completed and optionally allowed the participant to mark the page as a “do not ask” page, so that the system would not make future request while the participant was visiting that particular URL. To encourage a wider range of responses and tasks, the system would also not make a request on URLs for which a participant had already completed a survey.

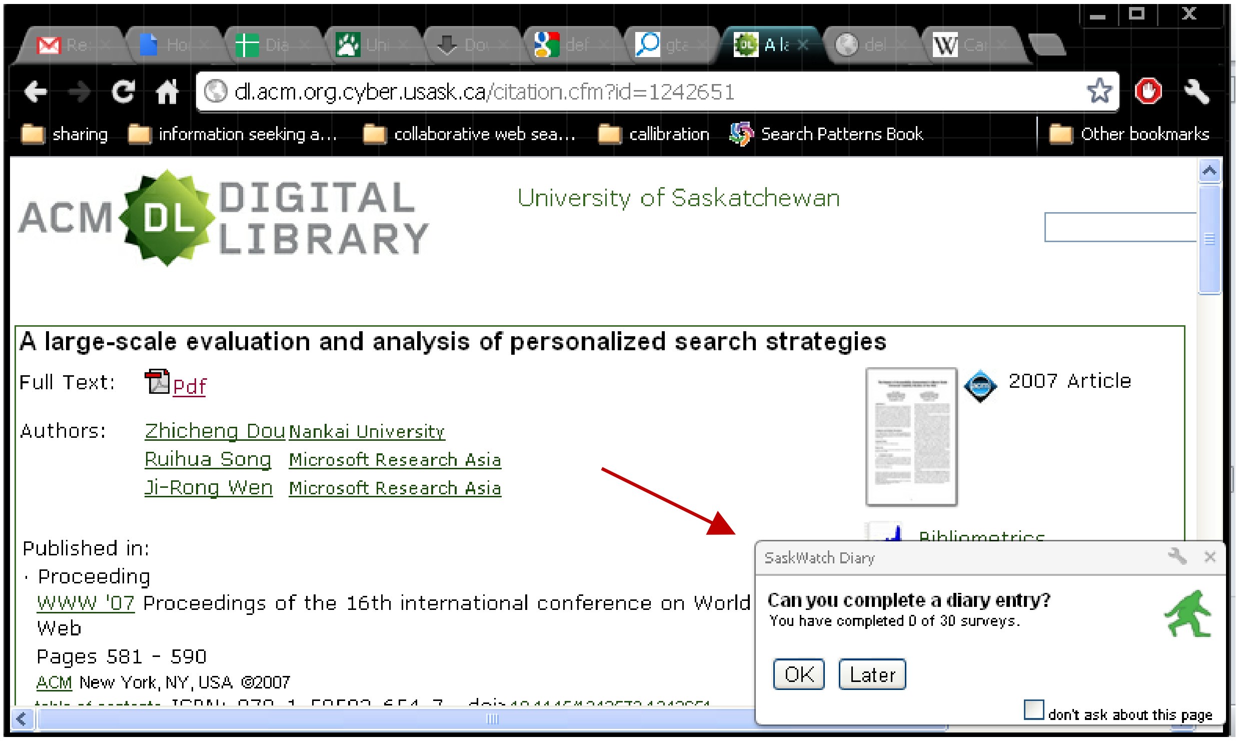

dl.acm.org.cyber.usask.ca/citation.cfm?id=1242651

University of Saskatchewan

A large-scale evaluation and analysis of personalized search strategies

Full Text: Pdf

2007 Article

Authors:
Zhicheng Dou Nankai University
Ruihua Song Microsoft Research Asia
Ji-Rong Wen Microsoft Research Asia

Published in:
· Proceeding
WWW '07 Proceedings of the 16th international conference on World Wide Web
Pages 581 - 590
ACM New York, NY, USA ©2007

SaskWatch Diary

Can you complete a diary entry?
You have completed 0 of 30 surveys.

☐ don't ask about this page

Figure 6.1. A request to complete a survey made by the diary study system (indicated by the red arrow).

If participants opted not to complete a survey at the time of a request, another request would not be made for at least 15 minutes. If participants agreed to complete a survey, a survey form was embedded in the page they were currently viewing, to allow participants to consider their responses within the context of the page and their current task (see Figure 6.2). Once a survey was completed a survey, the system would wait at least 30 minutes before making another request.

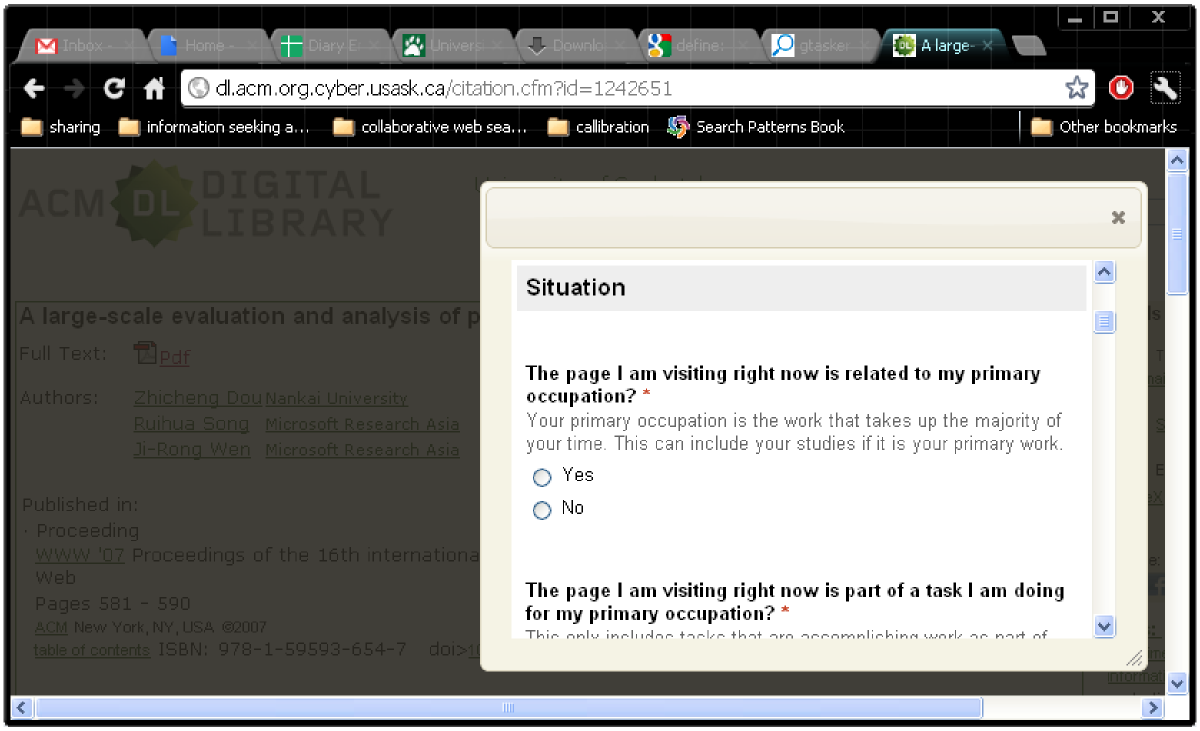
Inbox
Home
Diary E
Univers
Downlo
define:
gtasker
A large
dl.acm.org.cyber.usask.ca/citation.cfm?id=1242651
sharing
information seeking a...
collaborative web sea...
callibration
Search Patterns Book
Other bookmarks
ACM DL DIGITAL LIBRARY
A large-scale evaluation and analysis of p
Full Text: Pdf
Authors: Zhicheng Dou Nankai University
Ruihua Song Microsoft Research Asia
Ji-Rong Wen Microsoft Research Asia
Published in:
· Proceeding
WWW '07 Proceedings of the 16th internationa
Web
Pages 581 - 590
ACM New York, NY, USA ©2007
table of contents ISBN: 978-1-59593-654-7 doi>1

Situation

The page I am visiting right now is related to my primary occupation? *

Your primary occupation is the work that takes up the majority of your time. This can include your studies if it is your primary work.

○ Yes

○ No

The page I am visiting right now is part of a task I am doing for my primary occupation? *

This only includes tasks that are accomplishing work as part of

Figure 6.2. A survey from the diary system embedded in the current webpage being visited by a participant.

6.2.3 Data

The data collected in diary study consisted of the survey data only. A full list of the questions can be found in Appendix A. Surveys consisted of 11 questions, which fell into two broad categories: *situation* questions asked about the situation and task surrounding the participant's visit to the current page; and, *views towards sharing* questions, which asked about participants feeling towards having their interaction history made visible to their colleagues, and their interest in seeing the interaction history of their colleagues for the page.

The six situation questions were forced-choice questions, which in some cases allowed participants to select "Other" and provide a textual description of the situation. Situation questions allowed participants to describe the three situational factors: work relatedness, task type, and collaboration (described above).

The five questions about participants' views towards sharing were Likert-scale questions on either a 5-point or 7-point scale. These questions asked participants to rate how concerned they would be if their colleagues could see the interaction history describing their visit to the particular page and what those concerns might be. Colleagues were defined of the other "members of your primary workgroup." A free-text space also allowed participants to further describe their concerns about having their interaction history shared for the page they were currently visiting.

While I was also interested in whether or not participants would be *willing to share* their interaction history about the specific pages they visit, in this study I focused on participant ratings of privacy concerns only. In the field trial of WebWear, it was observed that participants were willing to share their workplace interaction history with their colleagues; however, they still expressed some privacy concerns. So, it was already established that people would be willing to share when using a social feedback system. For the diary study, I felt that if participants had even minor concern about privacy for a given page, they would tend towards judgments of not sharing. This would be due to the fact that the questions were asked outside the context of any real benefits to sharing (like those offered by the Search Dashboard and WebWear) [128].

To reduce any privacy concerns and to encourage participants to make honest reports about their activities during workdays, the study system did not collect any identifiable data and did not collect the URLs of the pages a participant had visited or reported on. While name and contact information were collected in the consent form, this was not connected to participant responses. Upon installation of the diary system, a random unique ID was generated, and allowed survey responses from the same individual to be connected together.

6.2.4 Data Analysis

The main criteria used to judge participants' level of privacy concern, were agreement ratings to the statement: "I would be concerned if the people in my primary workgroup saw details about my visit to this page. The details would include that I visited the page, when I visited and for how long." Participants rated their agreement with this statement on a seven-point scale (strongly disagree, disagree, slightly disagree, neutral, slightly agree, agree, strongly agree). Other Likert-scale questions are used in analysis to reveal further details about participants' privacy concerns. They are described with the results where appropriate.

Responses to Likert-scale questions were initially converted to numeric values to facilitate comparison, for both 7-point scales (0=strongly disagree, 1=disagree, 2=slightly disagree, 3=neutral, 4=slightly agree, 5= agree, 6=strongly agree) and 5-point scales (0=strongly disagree, 1=disagree, 2=neutral, 3=agree, 4=strongly agree). All questions were phrased in the same way relating to 'concern', so that the mean values could be directly compared for the overall question, and for each of the three contextual factors (work relatedness, task type, and collaboration). In this way, higher mean values meant that there were higher levels of concern about privacy, and lower mean values that there were lower levels of concerns. In the analysis, participants were deemed to be 'concerned' if they responded, at neutral or on the agreement side of the scale, and to be 'not overly concerned' if they rated a question below neutral (on the disagreement side of the scale).

The three contextual factors investigated in the study had the following levels:

- *Work relatedness:* work task or non-work task.
- *Task type:* fact-finding, information gathering, browsing, transaction, or other.
- *Collaboration:* collaborative or non-collaborative.

The diary study surveys contained two questions relating to whether or not the task was part of collaboration. The first question asked whether or not the participant had communicated with someone about the task they were conducting before or during their visit to the current page. The second question asked if participants planned on sharing any of the information on the page with anyone, either directly (through some form of communication) or indirectly (i.e., it would be used to inform future communications). In this analysis, I combine these two questions together to indicate whether or not a task was part of some form of collaboration.

The 2-week sample of 30 survey responses provided a limited scope to collect individual information seeking behaviour. This limited scope of data collection means that an individual participant's information seeking might be dominated by a particular set of contextual factors during the diary period (e.g., the participant conducted more information gathering tasks than normal because they were working on a literature review during the diary study). To avoid the bias that this may cause in overall results, the values for Likert-scale responses are first averaged over an individual for a particular contextual factor, and then averaged over all participants.

6.2.5 Results

All eighteen participants completed the study by answering at least 30 surveys within the 2 week period, and three participants completed 1 to 3 additional surveys; in total 545 surveys were completed. The diary survey underwent some initial refinement after the start of the study, and while the changes were not major, 12 survey responses were initially dropped to maintain consistency through all of the data. The final data set used in analysis consisted of 533 surveys.

For questions relating to task type and whether the task was collaborative, participants could select "Other" and describe the situation they were in as it related to the contextual factor. These descriptions were used to recode the response to the best value based on their description, if possible; otherwise, the response was left as "Other". Recoding occurred for a total of 37 questions. For example, one participant described the type of task they were conducting as "Searching for a specific resource", which was recoded to as a "Fact-Finding" task.

6.2.5.1 Overall Privacy Concerns

Overall, participants were not overly concerned with their interaction history being shared with their colleagues (*mean agreement* = 1.58, *std. dev.*=0.95). Figure 6.3 shows the mean ratings by participant. As can be seen, only two participants expressed overall concern, and 13 of the 18 participants rated their level of concern below 2 (or disagree).

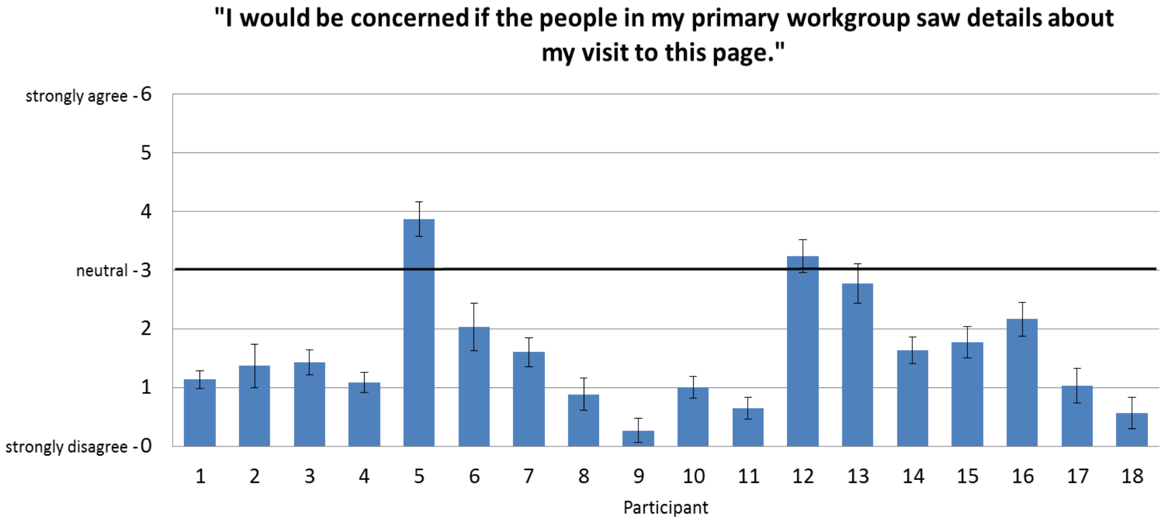

Figure 6.3. The mean agreement with being concerned if interaction history was shared with colleagues for the page currently being visited (by participant). The higher the value the more concerned a participant is on average, the lower the value, the less concerned.

Examining the individual ratings revealed that participants were not overly concerned about privacy for 76% of the pages. However, two individuals did express concern with this type of sharing and accounted for roughly one third of all incidents where concerns were expressed (participant 5 and 12). These two individual expressed concerns in over half of the tasks they reported on. These results suggest that peoples' overall level of concern is quite low with regards to sharing interaction history in the context of workplace information seeking.

Participants also rated their agreement with the statement "I would be concerned if the people in my primary workgroup saw anonymous details about my visit to this page. Anonymous details include that someone in the group visited the page (but not necessarily me), when the visit happened and for how long." When considering this form of anonymous interaction history participants' level of concern was reduced (*mean agreement* = 1.03, *std. dev.*=1.1). However, this did not reduce the level of concern for participants 5 and 12.

6.2.5.2 Relationship between Work Tasks and Privacy Concerns

In the diary study, work tasks were defined as only including "tasks that are accomplishing work as part of your primary occupation or studies." Overall, participants' information-seeking tasks completed during the workday were evenly split between work and non-work tasks – 268 were deemed as work tasks and 265 as non-work tasks. However, between participants there was a fair amount of variation in the proportion of work and non-work tasks (see Figure 6.4). These findings show that people commonly mix work and non-work information-seeking tasks throughout their day.

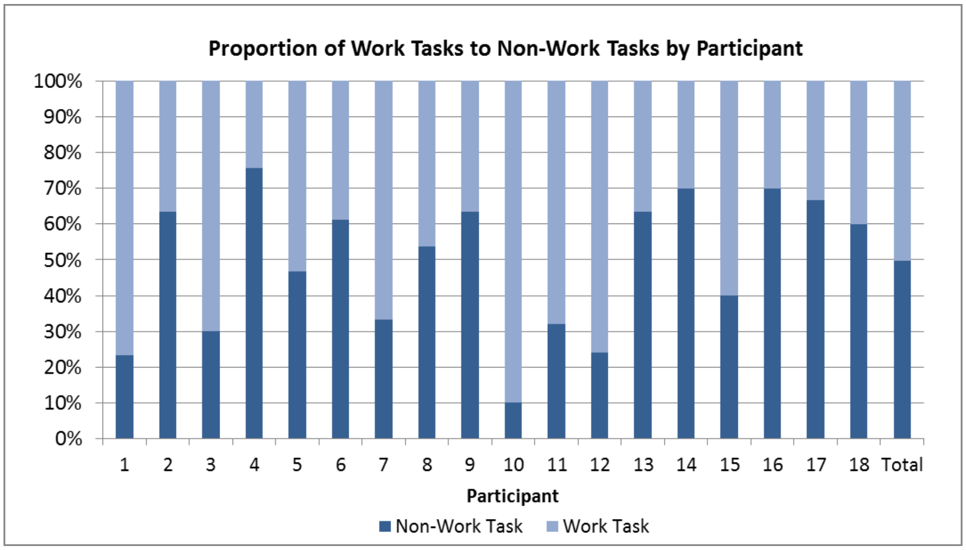

Figure 6.4. The proportion of work tasks to non-work tasks by participant.

There seems to be a clear relationship between participant ratings of concern about colleagues viewing interaction history, and whether a task is a work task or not. Figure 6.5 repeats the contents of Figure 6.3, but splits that data by the work relatedness (work task or non-work task). Overall, participants reported substantially lower level of concern for work tasks (*mean agreement* = 1.10, *sd* =1.19) than for non-work tasks (*mean agreement* = 2.16, *sd* =1.02). It should be noted that overall participants' were no overly concerned in either work tasks or non-work tasks. However, participants' level of concern was reduced for work tasks, and participants were not overly concerned in 86% of the incidents. Further, when considering work tasks, only 6 of the 18 participants reported specific instances where they had any concerns. The large majority of work tasks where there were concerns came from participants 5 and 12, who accounted for 22 of the 28 instances where concerns were reported. When not considering the responses of participant 5 and 12, the remainder of participants reported not being overly concerned in 95% of tasks reported.

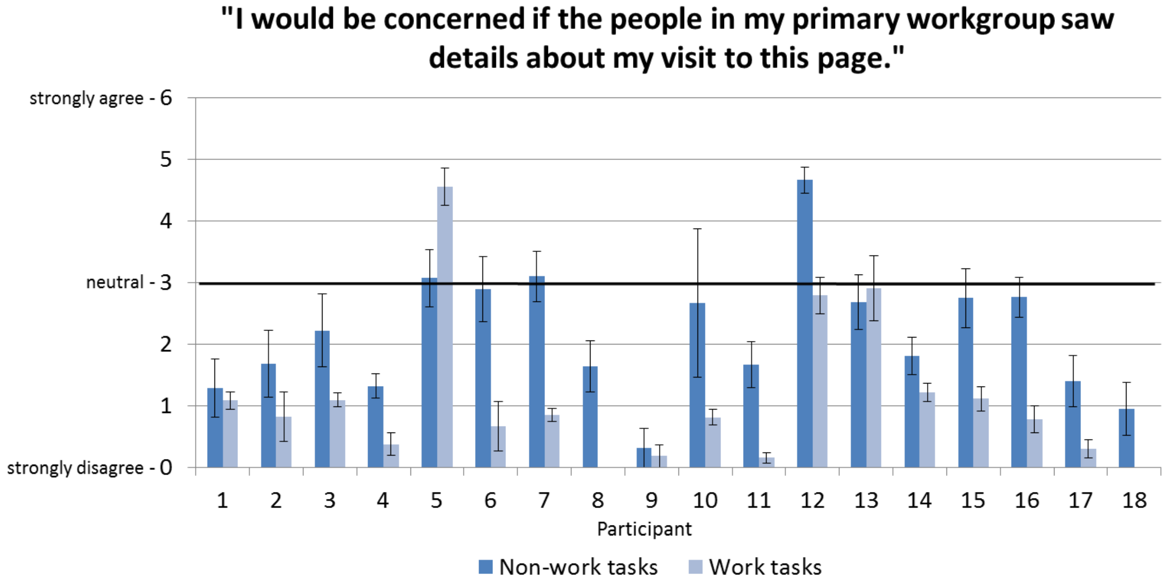

Figure 6.5. The mean agreement with being concerned if interaction history was shared with colleagues for the page currently being visited (by participant), split by whether it was a work or non-work task.

Overall, participants had a lower level of concern for work tasks. However, for a few individuals the level of concern remained the same or increased slightly for work tasks. Participant 13's mean rating went up slightly from non-work to work tasks. Participant 5, who was one of the two participants who had concerns overall, also expressed a much higher level of concern for work tasks than non-work tasks. The reason behind participant 5's increased levels of concern for work tasks is unclear. In comparison, participant 12, who also expressed concerns overall, was not overly concerned about work tasks. Overall, these results show that sharing interaction history with colleagues is less of a concern to participants when conducting work related tasks than when conducting non-work tasks, with a couple of exceptions.

6.2.5.3 The Relationship between Task Type and Privacy Concerns

In each survey, participants categorized the type of task that led them to the current page they were visiting as being one of fact-finding, information gathering, browsing, transactions, or other (see Section 2.2.1). Regardless of task type, participants were not overly concerned about their colleagues seeing their interaction history (see Figure 6.6). Participants expressed slightly more concern in transaction and browsing tasks. Recall that transaction tasks include the use of some online system such as email or online banking, and browsing includes tasks where there is no specific goal in mind. It could be that these tasks tend towards being personal (and therefore, non-work tasks). Looking at the individual rating, though, reveals that 93 of the 129 tasks (72%) where there were concerns occurred in browsing or transaction tasks.

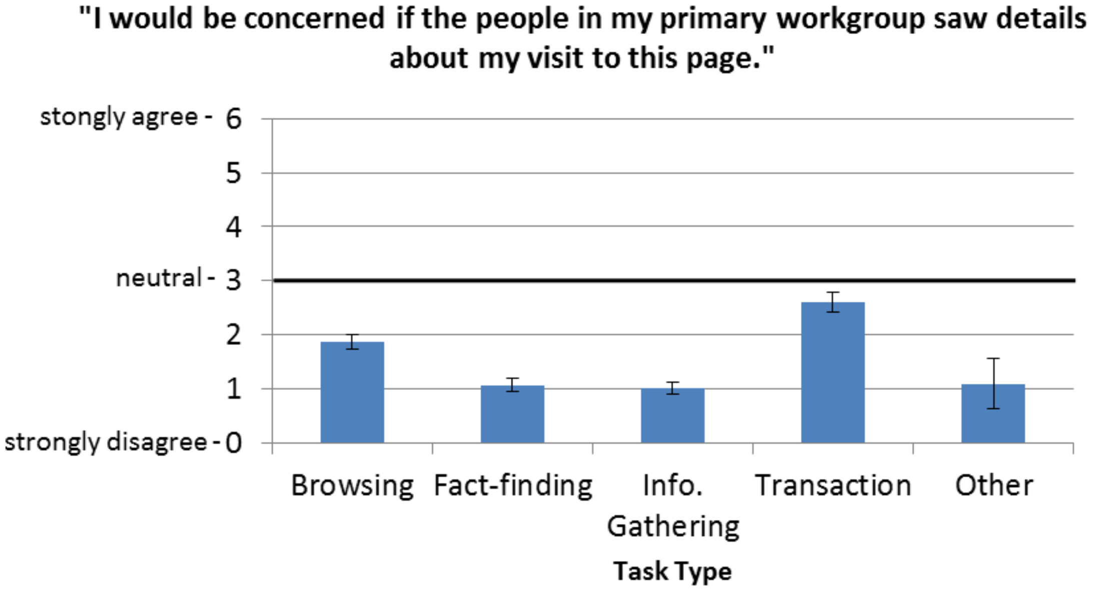

Figure 6.6 The mean agreement of being concerned with interaction history by task type.

Because of the privacy differences observed in work and non-work tasks, I looked at how frequent each task type was in work and non-work situations (see Figure 6.7). Browsing tasks were roughly six times more likely to be conducted in non-work than in work situations. This may indicate that browsing tasks are most often conducted as a leisure activity, which people may have a slight preference to keep to themselves. Transactions were also more frequently non-work tasks, whereas information gathering and fact-finding tasks were mostly completed in work situations. Tasks classified as “other” and that did not fit into the other four categories were only reported in 11 instances, making it difficult to report on observable patterns. Again, note that concern levels were rated as being below neutral for all tasks. Overall, it seems people are more concerned with task types that tend towards more personal, and that have less to do with workplace activities.

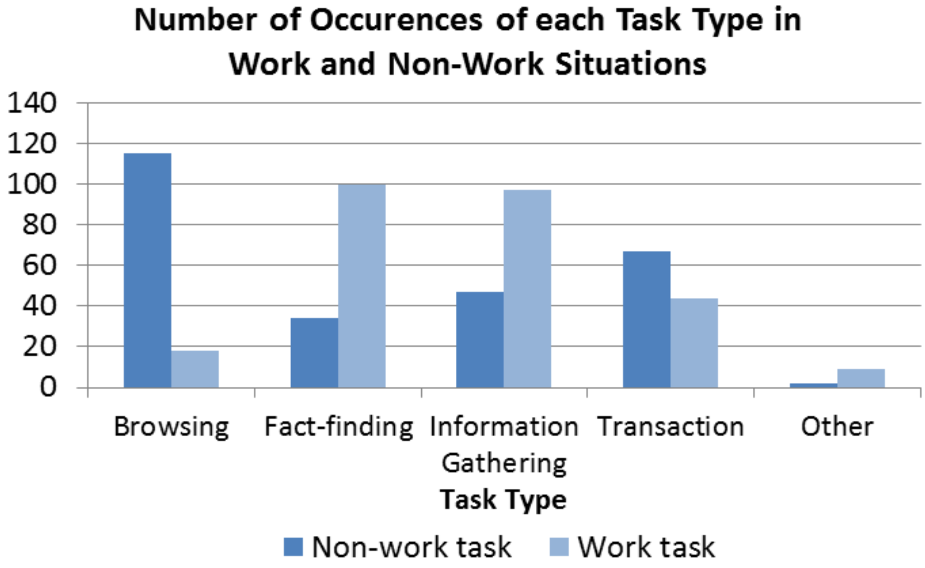

Figure 6.7. The number of occurrences of each task type reported in the diary study in work and non-work situations.

6.2.5.4 Relationship between Collaboration and Privacy Concerns

Collaboration was frequent among participants’ information-seeking tasks, with some form of collaboration reported in 58% of surveys. Collaboration was more common in work scenarios, with over two-thirds of tasks involving some form of collaboration. Non-work tasks, in contrast, were more frequently individual (see Figure 6.8).

I expected that people would have a lower level of concern when collaboration had already been established, because there is already some information sharing. However, this was not the case: there was no substantial difference between privacy concerns when collaboration was already established (see Figure 6.9).

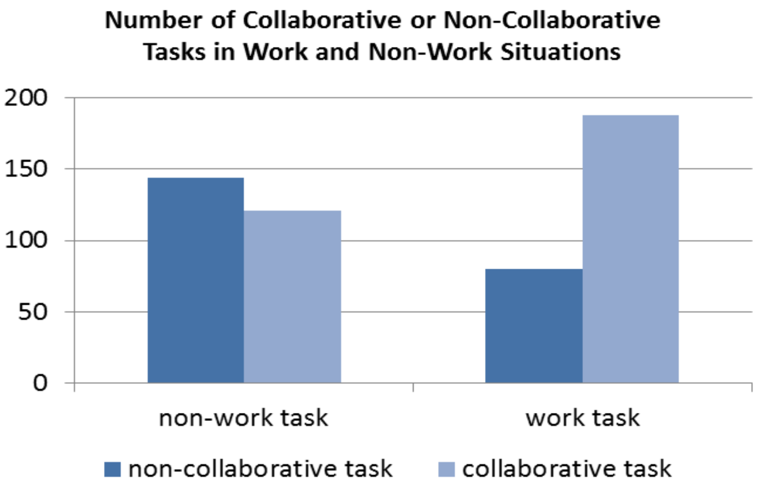

Figure 6.8. The total number of information-seeking tasks that involved some form of collaboration, split by work and non-work situations.

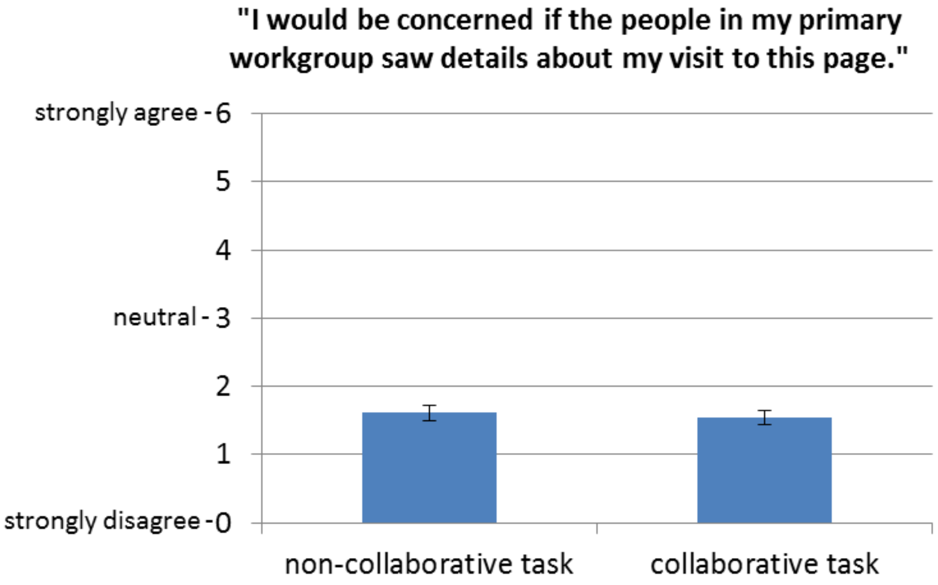

Figure 6.9 The mean ratings of agreement with being concerned if colleagues could view interaction history, split by non-collaborative and collaborative tasks.

6.2.5.5 Specific Privacy Concerns

For each survey, participants rated their agreement with three specific worries (on a five-point Likert scale). The three specific worries had been identified in the group interview of the WebWear field trial: 1) *Misinterpretation of the activity:* "Someone might misinterpret what I am doing"; 2) *Perception of time-wasting:* "Someone might think I should be doing something else"; and, 3) *Negative judgment of activity:* "Someone might not like or might not respect what I am doing."

Overall, participant ratings were low for each of the specific worries, and there were no large differences between each worry: misinterpretation of the activity (*mean agreement* = 1.04, *std. dev.* = 1.08), perception of time-wasting (*mean agreement* = 1.44, *std. dev.* = 1.28), and negative judgment of activity (*mean agreement* = 1.08, *std. dev.* = 1.09).

Again, ratings for specific worries were split into work and non-work tasks in order to explore how participants' level of concern may change in work and non-work situations. Consistent with previous findings, participants were not overly concerned about any of the three specific worries, and levels of concern were again higher in non-work tasks (see Figure 6.10). However, participant ratings of the 'perception of time-wasting' were substantially higher in non-work tasks than in work tasks. This suggests that during work tasks an important worry is that others might feel that time spent on personal or leisure activities would be better spent on activities that are part of a work task.

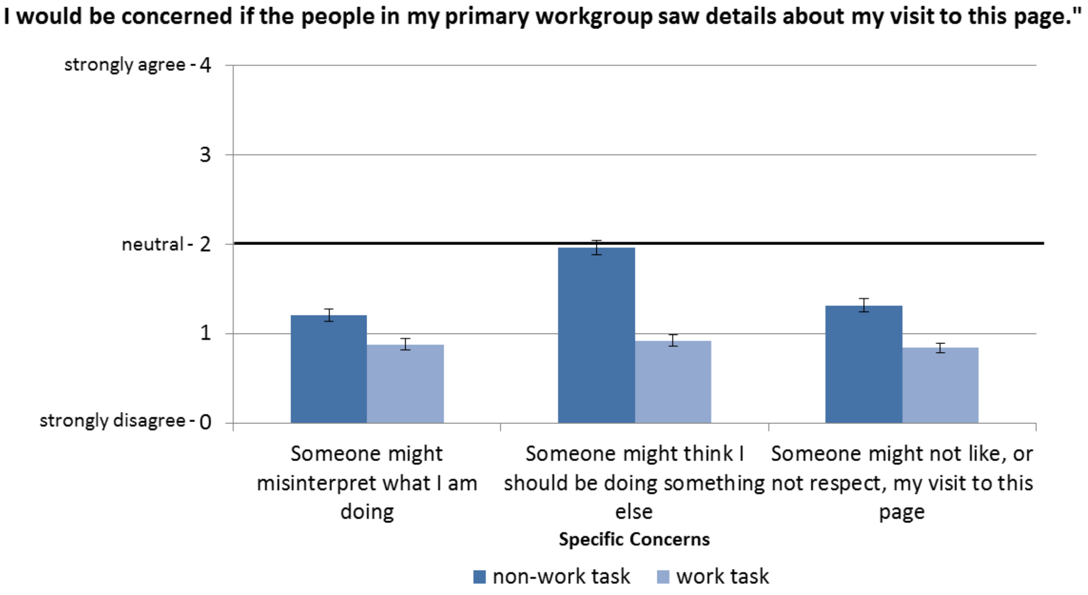

Figure 6.10. Participant ratings of agreement with specific concerns for sharing interaction history with colleagues.

6.2.6 Summary of Results for the Diary Study

The diary study of people's privacy concerns with sharing interaction history in the workplace provides six results:

1. Overall, people are not overly concerned with sharing their interaction history from workplace information seeking with colleagues. Suggesting that the workplace is an appropriate place for sharing interaction history.
2. When a task could be considered a work task, people felt less concern about sharing their interaction history than for non-work tasks (which would include personal or leisure activities).
3. People rated task types slightly differently. People expressed more concerns in transaction and browsing tasks than in fact-finding and information gathering tasks. Transaction and browsing tasks were more frequently non-work tasks, whereas fact-finding and information gathering tasks were most often work tasks.
4. Although information-seeking tasks frequently involve some form of collaboration (58% of the time). People were equally low concerns about privacy, regardless of the collaborative scenario.
5. Some individuals express having privacy concerns, even while most others have low levels of concern. Two of the eighteen participants expressed concern when others did not.
6. Of the three specific worries explored, people were most concerned in non-work situations that their colleagues might feel they are misusing their time. However, overall ratings of concern for this situation remained below neutral.

6.3 Summary of Results

The two studies in this chapter provided further evidence that social feedback could work effectively in real world deployments. The work in this chapter investigated the two primary challenges: overlap and privacy.

Through a real world field trial of the WebWear system, I showed that the overlap challenge can be contended with by broadening the scope of interaction traces and providing support at common starting points, such as search engines. The field trial also showed that people had mainly good experiences from the use of WebWear. While all participants agreed that WebWear was appropriate for use in the workplace, many still expressed having some privacy concerns. Despite any concerns, participants demonstrated that WebWear did not unduly intrude on their privacy and were not important enough to prevent active sharing of interaction history during real information-seeking tasks.

The diary study followed up on the results found with WebWear. The diary study focused on uncovering the potential contextual factors that may exist in workplace sharing, which could be used to understand when privacy concerns may arise. Overall, the diary study showed that people's level of concern about sharing interaction history in the workplace was low. Further, simply identifying a task as a work task consistently identified situations where participants had even lower levels of concern.

Taken together, the results of the two studies demonstrate that social feedback is feasible in real world situations. The overlap challenge can be overcome through the appropriate display of traces. The privacy challenge does not prevent most people from using social feedback systems, and particular contexts do exist where privacy concerns can be reduced even further.

In the Discussion (Chapter 7), I provide possible explanations for these main results, discuss implications for designing future social feedback systems, detail further issues in the real-world use of social feedback systems, and identify future research directions for exploring the challenges associated with making use of social feedback.

CHAPTER 7
DISCUSSION

In this chapter I analyze and discuss the findings of the previous three chapters, while making reference to the conceptual framework of social feedback presented in Chapter 3. I first summarize the main findings of the dissertation, and then discuss and provide explanations of the results for each of the last three chapters. Next, I provide an overall discussion, and finish with a reflection on broader motivations and how this work has addressed them.

7.1 Summary of Findings

In the previous three chapters I have described two systems and four experiments which test the new idea of social feedback. The two systems cover the two broad ways in which social feedback systems allow people to learn how to accomplish loosely-coupled information-seeking tasks more successfully: global learning and local learning.

First, I described the Search Dashboard (Chapter 4), which provides global learning of new search engine skills and techniques through comparison with the behaviour of expert search engine users. To achieve this, the Dashboard presents aspects of personal behaviour alongside the behaviour of expert models that were created by aggregating the interaction history of many different individual expert search engine users. As described in the presentation of social learning theory and social feedback (sections 3.1 and 3.3), being able to get a realistic picture of one's own behaviour and being able to compare that behaviour with a model provides clear targets for behaviour change and can facilitate learning. Through a study of the Dashboard and its effect on real world search engine use, I showed that comparison with experts was essential for behaviour change to occur. This was validated through a control condition that provided a Dashboard without expert model behaviour for comparison. The study of the Search Dashboard shows that global learning is possible through systems based on social feedback.

Second, I described the WebWear system (Chapter 5), which provides local learning to support seekers in accomplishing common information-seeking tasks. Based on social learning theory and social feedback, WebWear provides access to the past behaviours of those who are socially close to a seeker. To achieve this, WebWear collects, shares, and presents identifiable traces of interaction history within tightly-knit groups of seekers. These traces of interaction history allow the seeker to use the contextual knowledge they already have about tightly-knit contacts' activities and interests to interpret the presented behaviour to accomplish tasks with more accuracy, in less time, and with more confidence. Through a controlled comparison of WebWear, a standard browser, and a broad social navigation tool, I showed that WebWear leads to improved success in completing common loosely-coupled information-seeking tasks. Through a one-week field trial, I also showed that WebWear supports real-world tasks such as those evaluated in the controlled study. These studies of WebWear showed that social feedback can provide local learning and can allow information seekers to accomplish common information-seeking tasks with more success.

Finally, I investigated two principal challenges to the use of social feedback in the real world: overlap and privacy (see Section 3.2.2). The overlap challenge refers to the possibility that useful interaction history may be available too infrequently or not at all. Through a field trial of WebWear (Section 6.2), I showed that visitation overlap can be increased by broadening the scope of shared behaviour, and that opportunities for useful support do occur regularly. The privacy challenge refers to the possibility that because people perform tasks on the Web they wish to keep private, they may be unwilling to share their interaction history. The field trial showed that participants were willing to share their interaction history and did so with enough frequency that overlap occurred in at least several situations where it was useful for conducting a task. In addition, through a diary study (Section 6.2), I showed that people's overall levels of concern about privacy remained low for workplace information seeking, and that tasks that are focussed on work had reduced levels of concern. Together these results argue for the feasibility and usefulness of social feedback in the real world.

7.2 Global Learning and the Search Dashboard

Below I provide explanations for the results of the study of the Search Dashboard (results were presented in Section 4.3, and summarized in Section 4.4). I then discuss how the results may be generalized and discuss issues for real world deployments.

7.2.1 Search Dashboard: Explanations for Main Results

7.2.1.1 Why did the Search Dashboard work?

The Search Dashboard successfully helped people learn the behaviour of expert search engine users. Participants in the study demonstrated global learning through their adoption of behaviour that was more in line with that of expert searchers, and through their application of their new skills and techniques in a number of task scenarios. There were two main reasons why the Search Dashboard allowed participants to learn, and change their behaviour: *comparison with expert models* and the *ease of retention and reproduction* of the behaviour presented.

Comparison with expert models.

In the study, the main reason that the Search Dashboard worked was due to the ability to compare personal behaviour with the behaviour of the expert models. This was demonstrated through the use of the non-comparison variant of the Dashboard, which was identical to the original Dashboard, but with the comparative models removed. Comparison data from the expert models improved participants' perceptions of the value of the data and increased the degree of insight they were able to draw. Without the clear targets that were provided by seeing how other people behave, participants found the data less useful because it was not obvious how they could improve. Further, participants may not have known that there was room for change, but seeing that experts or even typical searchers used behaviour that was different from their own behaviour showed how they could change. For example, a participant seeing the Techniques data that shows how direct answers are commonly used by experts, would have easily been able to infer that experts use this technique to save time (because answers to common questions are shown with the search results). For Tendencies data, comparison provided clear targets for how behaviour could be changed, which would not have been available in the non-comparison condition. For example, showing a seeker that they take eight seconds on average to select a result does not necessarily suggest they should slow down. However, when participants saw that Search Experts take longer than Typical Searchers to select a result, they were able to infer that experts must study and select results more carefully than typical searchers (as suggested in survey responses).

Ease of retention and reproduction.

The ease with which data presented in the Dashboard can be acted upon affected participant views of the data's value, and affected how likely they were to adjust that aspect of their behaviour. This relates back to social learning theory, and the learning requirements of retention and reproduction. Together retention and reproduction imply that behaviour must not be overly complex in order for it to be remembered and for people to recall and execute it when it is needed.

The ease with which data can be retained and reproduced was one reason why some aspects of search behaviour were changed and others were not. Recall that five of the twelve aspects of behaviour led participants to change that aspect of their searching. Even with comparison data, some aspects of behaviour do not provide a clear path towards change. In general, the Topics data did not provide a clear use case to all participants, whereas data on Techniques and Tendencies were more straightforward.

Data relating to Topics was more difficult to apply since this data largely relates to the queries people use. Topics data can still provide useful information about topically relevant search terms, topic categories, and websites; however, because the range of potential information needs is so large, the data presented in the Topics section of the Dashboard may have been unrelated to a person's current need. For example, knowing that computer experts make frequent use of the stackoverflow.com domain might be interesting, but would be of little help when buying a laptop. This suggests that the topic data was too broad to be immediately useful in most cases, and that that topics data needs to be narrowed in order to provide support. Topic data would automatically be accounted for in a system like WebWear, because it displays traces from specific activities – for example, whether a closely-knit contact is searching for a laptop or programming help can be inferred from the traces. Another approach that would be somewhat broader than WebWear, but still fairly narrow would be to use collaborative filtering algorithms to provide personalized recommendations on topically relevant information for a user.

There are also several examples of Techniques and Tendencies data that did not lead to behaviour changes, because a means for reproducing the observed behaviour was not clear. For example, query ambiguity (an aspect of Tendencies that measures the specificity of queries, see Section 4.1.2.3) is an abstract concept that is not observable in search results or while formulating a query, so it is not apparent exactly how to influence one's own query ambiguity rating.

Finally, the ease of retention and reproduction was also essential because of the use of an *external to the page* approach for displaying behaviour (described in Section 3.3.2.3). This meant that people had to view their behaviour in one interface, infer what the behaviour meant and how it could be changed, decide if there would be value in changing the behaviour, and finally recall that behaviour when it could be applied in a real search task. In contrast, an *internal to the page* approach can display representations of behaviour that are in the location where modified behaviour can be used, meaning that retaining and reproducing the behaviour can be greatly simplified. Because there were many steps required from when the behaviour was viewed and acted upon, behaviour in the Dashboard needed to be simple and understandable for participants; otherwise, it simply would have been forgotten or disregarded, and never used.

7.2.1.2 Is comparison needed for learning?

Presenting people's personal interaction history and descriptions of the different aspects of search behaviour provided participants with information they may not have previously considered as important for their search performance. Seeing their behaviour likely caught their interest, and led them to explore and consider the different behavioural aspects carefully. In this way, viewing personal interaction history without expert models was sufficient to provide reflections on participants' behaviour, by providing new insights and interaction possibilities, and also by showing them that there may be more to search engine use than they had previously thought.

In the "no-comparison" version of the Dashboard, participants did demonstrate learning: they came to believe that they were not as skilled with search engines as they had previously thought. While the intent of social feedback is not to have people feel poorly about themselves or their abilities, having a realistic view of one's own abilities and behaviour is an important aspect in learning (see Section 3.1.4). So, reflecting on personal behaviour through the Dashboard was sufficient in this regard for people to demonstrate learning. However, people did not learn new search behaviour from the no-comparison Dashboard.

The need to compare with experts for behaviour change to occur may also suggest that learning from observation works at a more fundamental level. People in the non-comparison condition may have seen all of the same information that people in the comparison condition saw. However, due to the perceived credibility of an "expert searcher," people may naturally put more stock in the information that is provided and feel it has more value, which leads them to be more likely to pay attention and remember it. This mechanism could even be completely subconcious, where people do not know that they are retaining more because an expert has demonstrated the use of some unknown knowledge or behaviour.

Finally, that data that does not lead to behaviour change can still have value. The Search Dashboard helped to change users' views of Web search overall because a complete picture of different search aspects was given, even without comparison. Further, different people valued the data elements differently: some people preferred topics data and others prefered techniques data. Providing a fuller range of behaviours could serve to engage a wider audience, and lead people to pay closer attention to all of the behaviour provided.

7.2.2 Search Dashboard: Generalizing the Results, and Deployment Issues

While the results clearly show that social feedback *can* lead to behaviour changes, they must be interpreted remembering how participants encountered the Search Dashboard. The surveys that were provided in the study likely aided in inducing closer consideration of the presented behaviour, because they created an artificial requirement for people to look at and think about their search history in a way they might not have otherwise done. Further, it is possible there was an effect of novelty. While there may have been confounding influences, I believe their role was minimal compared to the interface itself, because behavioural changes occurred for the comparison condition and not for the non-comparison condition.

Another important consideration is how experts were identified. Our approach of selecting the users who most often employed search operators likely misidentifies some experts and misses others. However, the approach was correct often enough that the aggregated behaviour patterns of expert and typical user models were consistent with previous findings (e.g., [158]). Regardless of how experts are characterized, our work makes clear that searchers can learn new search behaviour through comparison with the created expert models.

It is also not clear whether the Dashboard necessarily helped participants become better searchers or to improve their search performance. For example, users may try to use more query terms than necessary, leading a search engine to return results that are less relevant. This would mean participants are worse off because of their changed search behaviour. In this initial study of global learning, the primary goal was to understand if social feedback can allow searchers to learn new search behaviour and attitudes. However, related work has suggested that the knowledge and use of particular search engine functionality and other characteristics of search engine use can predict improved searching performance [7,158], and that people can infer the thought process of others by imitating behaviour [38]. I hypothesize that people may have been able to *think* more like search experts – i.e., people were able to narrow down the space of possible behaviours to those that are more in line with experts. This would mean that people can apply their new knowledge appropriately and that they did improve their search performance.

Social feedback can lead to changes in terms of behaviour and attitudes about search; therefore, it should be used by search engines where possible. Given recent interest in providing feedback on search behaviour and in educating searchers (see Section 2.2.3), I believe the Search Dashboard can be used in some situations with little change. A Search Dashboard-like system would fit in well with existing search engine history facilities. Further, because all of the data used in the Dashboard are currently collected by search engines, the development cost of such a feature would be low.

It is also interesting to consider how users would encounter and use a Search Dashboard system if it were part of a search engine. Current search history functionality may not be widely used (except in specific situations such as trying to refind a previously viewed website). However, the Dashboard offers a very different, and likely more engaging, view of personal search behaviour. Upon learning about a new Dashboard-like system, many people would be inclined to test its functionality; and, as I demonstrated in the study, even a single session is sufficient to affect people’s attitudes and search behaviours. After an intial encounter, people would use the Dashboard as a reference when they have difficulties or when an opportunity arises. This being said, it is also interesting to consider how aspects of the Dashboard may be inserted directly into existing search engine interfaces for more frequent contact with searchers, and as a way to reduce the requirements of retention and reproduction on learning. Reminders should be personally meaningful to capture user interest and suggest a target for change, rather than an impersonal standard item such as a “Tip of the Day”. For example, after a period of observed search engine use, a message could be added directly to search results that says: “You have used fewer terms in your queries this month than last month. Search experts average more query terms than novices.”

7.3 Local Learning and WebWear

In this section, I provide explanations for the results from the user study and field trial of WebWear (these results were presented in Section 5.3, and summarized in Section 5.4). I then discuss how the results of the work on WebWear may be generalized and provide considerations for real world deployments.

7.3.1 WebWear: Explanations for Main Results

The results observed in the user study show that social feedback can provide local learning, and can allow loosely-coupled information-seeking tasks to be completed more effectively and more efficiently. Further, I showed that trace information is more useful for local learning in these tasks when it is contextualized to a tightly-knit group of contacts.

7.3.1.1 Why did WebWear work in the user study?

WebWear provided three main tools that led to the successful results in our user study. First, group-query completion allowed users to see if their search terms aligned with those of their contacts. This was useful in both ‘follow’ and ‘divide’ tasks. In follow tasks, users could repeat queries to try to get the same results as a previous collaborator; in divide tasks, they could formulate unique queries to get results that had not been seen by their collaborators. Second, the activity glyphs displayed next to links (see Figure 5.1) allowed users to quickly identify the pages that they should visit or avoid. Third, the popup details on the glyphs provided appropriate information that allowed users to confirm if an activity was of interest or not. Most often, simply seeing *who* had left the trace was all that was needed (and the user portraits available in the system gave this information at a glance). Further, the details for websites (Figure 5.3) provided a useful way to quickly find a known target. For example, to complete the “bird book” task (described in Section 5.2.2) a user could simply visit amazon.com, look at the website popup, and see the list of pages John had visited on Amazon. This form of contextual history provides a natural way for people to quickly find history of interest. Together these three tools allowed seekers to explore the previous actions of their closely-knit contacts. When trace information was available, it could be quickly considered, and if useful, it could be applied immediately by users updating their planned actions based on the new information.

7.3.1.2 Did WebWear lead to new information-seeking behaviour?

In both the user study and the field trial, I found that people adopted behaviour that was different from their normal search behaviour. First, they used the trace information – in the user study all participants actively made use of the trace information; this was demonstrated in their increased success in completed the study tasks. Second, people worked in ways that would allow them to encounter traces as part of their normal search practices. For example, recall that some field trial participants actually started using the Google search webpage rather than their browser address bar to execute searches; this allowed participants to encounter the search traces of others. Third, I observed several cases where people would formulate queries in the search box without executing them, to see if they could find someone else's past query. Finally, I also saw a few people become momentarily confused when activity traces didn't appear where they were expected on search result pages. In these cases participants seemed to have incorporated expectations of WebWear-style information into their use of information scent on web pages. This problem arose due to inexperience with the system, but highlights the importance of making the source and display of trace information simple and easy to understand.

7.3.1.3 Why did broad social navigation not work?

The broad social navigation system (BroadNav) tested in the user study was no better than a normal browser for most tasks, and users completed tasks significantly more slowly than in the other interfaces. The reasons for these results may arise from the issue of introduced cost. While I intended the design of WebWear and BroadNav to have a low cost (small glyphs that were easily ignored), it is likely that participants needed more time to consider them. BroadNav offered information that was too general for our tasks (e.g., participants wanted to find John's bird book, not the book that was most popular); therefore, time spent considering BroadNav information was not beneficial. This effect could diminish with more training – that is, users could learn to only attend to the glyphs for certain tasks. Attending to the information of WebWear also introduced a cost, but this was offset by the benefit provided.

7.3.1.4 Were the results of the WebWear user study expected?

The results of the user study were not overly surprising given the scenarios that were tested. However, this specific type of task (loosely-coupled information-seeking tasks) has not been previously examined; the literature suggests that loosely-coupled information seeking may be a more common use case than explicit forms of collaboration (as described in the introduction of Chapter 5). WebWear combines lessons from social navigation and collaborative information seeking, and the studies show that this combination provides good support for loosely-coupled information-seeking tasks.

Further, there were two previous issues in using interaction history on the Web that pose potential problems: *ephemerality* and the *loss of context* (introduced in Section 3.2.2.1). Since these are issues with interaction history in general, whether they would cause problems for users of WebWear is unknown.

Ephemerality refers to the fact that much of the Web is constantly changing. The same search engine query produces different results for different people, at different times, and in different places. Further, the content of the pages themselves are always changing: new pages are created, new comments are added to blog posts, the top stories change on news sites, and old content disappears. Our user study used artificially created traces, but these traces were created on the actual Web, which is subject to ephemerality. The study took place over the span of four weeks, and while the experimenter would notice changes in the content of pages from time to time, participants were not greatly affected by it. In some cases, this meant that the traces displayed by links for a given query were available for some participants and were not visible for others. However, as was intended, participants would continue with their tasks and make use of traces when they were available.

A loss of context refers to the fact that the actual intention behind displayed behaviour cannot be observed by a seeker viewing traces. Intention is not captured or represented by social feedback systems, but rather must be inferred by a seeker making use of traces. Participants in the user study were able to use the information from outside the system to interpret the traces inside the system, and understand the context. The participants in the field study did have the advantage of having access to the task descriptions, which included the background knowledge needed to successfully infer about traces. However, participants in the field trial did not have the details of context close at hand; instead they were able to successfully infer the context based on the knowledge they already had of their closely knit contacts.

These issues of context loss and ephemerality mean that the results of the user study were not a foregone conclusion. However, the simple design of WebWear allowed these issues to be avoided, while still effectively providing support.

7.3.1.5 Did people really learn with WebWear?

The form of local learning that WebWear provides may not conventionally be thought of as learning. It is clear that WebWear led participants to modify their behaviour based on the information they observed, which fits into the definition of learning (described in Section 1.1.1).

In local learning, new knowledge is highly contextual and possibly short-lived – it is based on a particular situation, is only needed for a brief period of time, and may only be used once. These characteristics can best be described as *advice*, which has been recognized as an extremely important form of learning, because it allows people to learn about courses of action in specific situations where getting personal experience can be difficult [111]. So, WebWear did provide learning in the form of implicit advice from observing the previous actions of other people (social navigation systems have previously been described as providing implicit advice [142]).

7.3.2 WebWear: Generalizing the Results and Deployment Issues

There are three main issues in generalizing our results to the real world: usefulness, overlap, and privacy. Below, I will discuss all three issues, but limit the description of the two main challenges of social feedback (overlap and privacy) as they relate to the deployment issues of WebWear; I discuss these issues in Section 7.4.

7.3.2.1 Providing Beneficial Information

WebWear shows that traces of interaction history can be successfully collected, presented, and used for local learning, but there are likely many situations where traces may not provide valuable information. One of the known challenges of social navigation systems is the problem of *herding* (i.e., blindly following others' trails) [63]. Herding is a problem if a trail leads a task to be completed poorly.

Second, there is a cost to using trace information during information seeking – it takes time to consider the traces, decide if they are useful, and decide a course of action. Therefore, the benefits of making information abundant and available must also be considered in terms of the quality of the information – this could be considered as *signal-to-noise ratio*. For example, I did not expect a great deal of overlap in the field trial, due to the wide range of activities that people had and the short length of the trial. However, the (green) website glyphs that can appear next to each of the links on a page (see Figure 5.1) actually appeared quite often for some participants on certain search result pages, which they found distracting and not overly useful – which means that there was too much overlap and the signal-to-noise ratio was low. Systems like WebWear must balance between broadening the possibility of encountering traces and ensuring that the information is still valuable.

There are, however, several design possibilities that can be explored to address this specific issue. For example, a simple approach would be to narrow our website-level traces. Currently, website activity is presented for all traces on a top level domain (e.g., google.com); this could be narrowed by including the full domain (e.g., the website would be separated for traces on sites like news.google.com and maps.google.com). However, there is some evidence from the field trial that curiosity leads people to investigate more indicators than necessary (which could partially be an effect of novelty). With more familiarity, and because loosely-coupled collaboration is discretionary, people can decide on a case-by-case basis whether the information provided by WebWear is useful. Further, the interpretation context will help people to determine whether information is valuable or spurious (e.g., "I know Mary is not an expert in biochemistry, so I won't follow her trail in this area").

7.3.2.2 Challenges of Social Feedback for WebWear: Overlap and Privacy

The challenges of social feedback may be of particular concern to social feedback systems like WebWear. First, the sparseness of traces is an important potential issue for small workgroups – given the size of the Web it could be that traces are encountered too infrequently to be useful. Second, WebWear passively collects a seeker's interaction history and attaches their identity to it, making it clear to all people who view it what the activities of the seeker were. Revealing such activities can be worrying to people. These issues are of particular concern to WebWear and are discussed below.

7.4 Real World Challenges for Social Feedback

The field trial showed that overlap of traces was a common occurrence in the real world and that people felt that sharing their traces with the members of their lab did not unduly intrude on their privacy. The diary study found that the overall levels of concern were low, and workgroups provide an appropriate context for sharing interaction history, especially when considering tasks that are focused on work. In this section, I provide explanations for the results from the field trial of WebWear and diary study together (summarized in Section 6.3), as they relate to the challenges of overlap and privacy.

7.4.1.1 Real World Challenges: Explanation for Main Results

Why did overlap occur?

The results of the WebWear field trial suggest encountering other people's traces is possible, and can lead to useful support. However, it is still difficult to get a picture of how frequently a group of users might encounter one another's traces for supporting information seeking.

Overlap of visitation is important for a local learning tool like WebWear, which relies on the traces from a small group to be encountered during information seeking – e.g., a seeker must reach a specific page where traces are presented. For WebWear, encountering the behaviour of others is essential to the utility of the system. In comparison, for the Search Dashboard overlap is less of a concern, because general tools like search engines and web browsers are commonly used by all information seekers on the Web, and so overlap can occur commonly.

Two principles provide some indication as to why there was reasonable overlap observed in the WebWear field trial, and why it is reasonable to expect overlap in future deployments: the *use of common starting points* and the *characteristics of tightly knit groups*.

First, there are a set of common starting points (e.g., search engines) that increase overlap and provide opportunities for shared encounters. WebWear's group-query completion was designed with this in mind, and the field trial showed that this feature was frequently used, and used in the ways that were intended. Further, some people picked up on this as a place where they could frequently encounter the activities of others; these people changed their behaviour to make sure overlap would occur.

Second, a defining characteristic of tightly-knit groups such as workgroups is that they have shared goals, interests, and tasks (even when working in a loosely-coupled fashion). This common ground should translate into people visiting some of the same regions of the Web. In the field trial, the participants (who are members of my research lab) all have shared interests in computer science, human-computer interaction, and computer programming. During the one-week trial, there was strong evidence that the common ground in the group led to common visitation: all participants used Google to search the Web, fourteen visited the University website to access administrative services, nine people used stackoverflow.com for programming advice, five people visited the departmental website to find information on courses, and seven visited acm.org to review literature. These common places represent shared interests that are the basis for the group's relationship, and are likely to occur over and over again as work tasks take people back to the same sites.

Why was the overlap useful?

The overlap was useful in two main ways: it provided a general awareness of shared activities, and it enabled support to be provided in loosely-coupled information-seeking tasks.

Having a general awareness of the activities of others was mentioned several times in the field trial group interview as being an important outcome from using WebWear. Awareness is an essential part of workplace interactions (see Section 2.5.2). People recognized that searching the Web and browsing webpages are a core part of their daily activities and directly coincide with the tasks that they have. This means that information seeking activity is an ideal channel to provide updates on people's general work activities.

Overlap also provided several opportunities to support loosely-coupled collaborative tasks during the field trial (including follow and divide tasks). These occurred for the same reasons that visitation overlap occurred – the characteristics of tightly-knit groups mean that group members will also share tasks. When tasks are shared and overlap occurs there is an opportunity for support to be given. In the field trial, support was available often enough that it was useful in at least several situations. However, it is clear that in many situations support will not be available. One potential reason for this is the cold-start issue (see Section 2.6.2.2), where the first user to conduct a task will not have the benefit of support. However, the goal of social feedback systems does not need to be complete coverage of all tasks for there to be benefit. Given how common web-based information seeking is, providing any form of support can lead to substantial overall gains.

Why did privacy concerns not prevent sharing in the field trial?

Concerns about privacy have often been the main barrier to the deployment of some kinds of awareness systems [74]. Because WebWear passively collects people's interaction history, it could in some cases feel like surveillance. In addition, WebWear attaches names to particular actions, meaning that there is no anonymity as there is in the Search Dashboard. Therefore, systems like WebWear must be deployed very carefully. Despite the potential for privacy concerns, all sixteen of the people who initially volunteered for the field trial completed the one-week exercise, indicating that there are people who are amenable to this degree of sharing (within a small workgroup). In the responses gathered from these participants, no one complained about being tracked (and no one withdrew from the study or stopped sharing interaction history for long periods). Privacy was a strong issue of interest in the group discussion – but the conversation was equally about how concerns can be mitigated in this context while maintaining low effort in sharing. This conversation also highlighted directions for future work (studying low-cost methods for activity sharing) and has implications for other areas where user behaviour is actively collected (discussed further in Section 8.2.1.2).

In the group discussion held after the field trial, people often made reference to not minding their activities being shared in this case, because it was with people in their workgroup whom they know and trust. This suggests that, although many of the relationships in this group do not extend outside of the workplace, there was enough trust and understanding of the other members of the group that participants felt reasonably safe. Some participants also felt that the simple control mechanisms – to stop sharing activities and to permanently delete accidentally entered behaviour – were sufficient for this type of sharing. Finally, several participants stated that they do not see their work setting as providing a great deal of web privacy anyway, and so they already carry out browsing tasks at work with the expectation that others will be able to see their activities. These factors all likely played a role in people's relative comfort with participating in the field trial, and led them to actively share their behaviour without serious worries.

Why were levels of privacy concerns low in the diary study, and particularly low for work tasks?

Based on the comments and the overall focus on privacy in the group discussion of the WebWear field trial, I expected that overall levels of privacy concerns would have been higher, and incidents where there were concerns to be more frequent. However, when people actually made judgments about their concerns relating to privacy they were not overly concerned most of the time. Privacy concerns are likely subject to negativity bias [128], and people feel that the few negative incidents that they do experience are more common and widespread than they actually are. While the majority of concerns that were expressed came from a few individuals, all participants expressed having at least one incident in the diary study where they had concerns with sharing. This means that privacy concerns are not an uncommon experience in general, but for the most part, concern levels remain low because people have an established relationship with their workplace colleagues. Because of the trust that exists with colleagues, most have few concerns about sharing their interaction history in the context of work. This means that colleagues are an appropriate tightly-knit group for workplace sharing of interaction history.

Levels of concern were also lower when participants were actually performing a work task. The diary study also showed that roughly half the tasks conducted at work were non-work tasks. It is clear that participants had a tendency towards wanting to keep the details of personal tasks private, because they are outside the basis of relationship with colleagues, may contain confidential details, and are more likely to be misunderstood (colleagues lack the contextual knowledge to accurately interpret the activities). Sharing information about work activities is clearly less of a concern, because colleagues often know about the activities already.

Will low levels of concern about privacy lead people to share their interaction history?

Whether people would share or not is likely a function of several factors, including the perceived risks of unwanted viewing and the cost of effort in managing interaction history. Participants in the field trial did not mind sharing interaction history with other lab members in the short term. They felt that there was minimal risk because they trusted their colleagues, they were guaranteed that unwanted history in the system could be permanently removed, and that no collection was happening when they disabled the system. Further, the effort of managing sharing was low, because it was done automatically (although it could be turned off). However, several participants in the field trial still expressed some concerns about sharing. Over a longer term and with fewer controls, these concerns may lead people to opt out of sharing their interaction history or using a social feedback system.

While the results from the diary study do not show that people’s concerns about privacy can be ignored, they shed important light onto the issue by showing important situational factors where sharing interaction history may be more acceptable. First, participants were not overly concerned about sharing their information seeking activities at work. Second, participants did carry out personal tasks in the workplace. Third, participants had lower levels of concern about their work tasks than personal tasks, for sharing with their workgroup. Fourth, certain individuals may be much more likely to express concerns about sharing their interaction history.

It is clear that the lower the privacy concerns, the more likely people are to be willing to share. The diary study has shown that the workplace can be appropriate for sharing interaction history, especially if sharing is limited to work-related activities. However, how work tasks can best be targeted and personal tasks ignored for sharing interaction history is an open issue. Finally, some people may prefer not to participate under any conditions, because their privacy concerns are too prevalent. Additional study might reveal the reasons and other contextual factors that can lead certain people to have stronger concerns than others. Further, with more experience with social feedback systems and understanding of how they can be used, many people would become more comfortable with having their information-seeking activities shared with trusted contacts.

7.4.2 Real World Challenges: Generalizing the Results

The studies provide two main findings that can be generalized in further work in information seeking and social feedback. First, people's tasks on the Web do often overlap within in small tightly-knit workgroups. This means that future research can take advantage of the common occurrence of shared activities and visitation to investigate new ways for information seekers to work together. Second, the finding that privacy concerns are lowest when work tasks are shared with work contacts may suggest a more general principle that can be applied when sharing interaction history – that is, the basis for a personal relationship within a particular tightly-knit group also describes the tasks that are appropriate for sharing. For example, it is likely that people would feel sharing *family photos* with *family members* is more appropriate than with their colleagues. This being said, functionality that allows this kind of targeted sharing of interaction history would require additional effort on the part of a user. Further, automated solutions are likely not feasible – the automatic identification of tasks and tightly-knit groups with whom to share task details will not be solved easily. However, as was seen in the field trial and the diary study, the workplace does seem to provide ample opportunity for sharing interaction history. Further, it also seems likely that people would be willing to share their work tasks while at work if their colleagues were also doing it. Even though personal tasks appear to make up a large share of workplace information seeking, a second major portion is work tasks, which means that there are ample opportunities for widespread sharing of interaction history within workgroups.

7.5 Overall Discussion

In this section, I discuss four themes raised in this dissertation in more detail: the importance of selecting an appropriate model in social feedback systems, the important differences between global and local learning, the importance of loosely-coupled information seeking, and the importance of information seeking more generally.

7.5.1 The Importance of Appropriate Models: Tightly-Knit Groups and Role Models

As described in social learning theory, the identity of a model determines the value people perceive in viewed behaviour. Therefore, identifying appropriate models may determine whether a social feedback system can be successful or not. In the Search Dashboard study, search engine users valued the skills and techniques displayed by expert search engine users. In the WebWear study and field trial, participants valued being able to see the relevant activities of their closely-knit work contacts.

The use of appropriate models also allows the challenges of social feedback to be overcome. The Search Dashboard overcame the privacy and overlap challenges because experts could be identified from a large pool of users without needing to reveal individual identities. Before the studies of WebWear, it was not clear that the system would be able to overcome the privacy and overlap challenges, because the pool of users contributing interaction history was small and because individual identities were shared. However, WebWear was used in a context where the properties of a tightly-knit group were able to overcome these challenges. First, the overlap challenge was overcome because tightly-knit groups share enough in common that they visit similar parts of the Web and conduct similar tasks. Second, the privacy challenge was overcome because enough trust already existed within the group, so participants were not overly concerned that their activities would be misused or interpreted in an overly detrimental way. It is possible that with a different group, where the social connections were looser, that overlap would have occurred less frequently and that privacy concerns would have been more frequent and more serious.

The appropriate use of a model by a social feedback system also means that an information seeker must already have some of the important details about the model in their mind. For example, a WebWear user must know (or be able to make reasonable a guess) about the activities of their colleagues, in order to determine the value or relevance of encountered traces. In contrast, previous work that has made use of interaction history (see Section 2.4) has focused on creating computational representations of users based on interaction history, and uses a system to make the important inferences about what information is most relevant to a user or not. This dissertation has shown that at least some of the work of selecting and making use of interaction history can be done by the user of a social feedback system, since users have the knowledge to assess the value and relevance of model behaviour.

7.5.2 What are the differences between the approaches to supporting global and local learning?

Global learning involves acquiring new skills and behaviour that can be applied in many situations, whereas local learning involves modifying behaviour or receiving new information that allows a specific task to be complete more successfully. Beyond these characterizations, different approaches may best lend themselves to supporting each form of learning. Below I contrast the approaches of providing aggregations of behaviour to providing individual traces of behaviour.

The Search Dashboard aggregates the behaviour of many anonymous individuals to provide global learning. There are two main advantages of aggregating behaviour: to make high-level behavioural patterns clear, and to reduce privacy concerns. First, individual search behaviour is idiosyncratic and can be difficult to understand. People’s behaviour is idiosyncratic because their behaviour, in part, is dictated by the specific tasks and goals they have. For example, some tasks result in shorter search engine queries, whereas others tend towards longer more verbose queries. Further, it can be extremely difficult for a user to learn how people make use of tools when looking at individual traces. For example, seeing a single interaction history event that says a seeker searched for “facebook” does not tell me if searching for facebook is common. By aggregating the behaviour of many individual search events, overall patterns of behaviour can be revealed and high-level questions can be asked (e.g., does this person search for “facebook” often? Or, do all people search for “facebook” often?).

Second, aggregated behaviour can reduce privacy concerns, because information that is uncommon can be filtered, and when behaviour is shared (such as visitation to a particular webpage) individual traces of interaction history are less likely to identify an individual. Further, as seen with the Search Dashboard, aggregation can allow individuals to be aggregated into group-level profiles of behaviour (e.g., expert search users), which hides individual identity.

A disadvantage of aggregated presentations of behaviour is that they require a lot of individual seekers to create a representative model of behaviour. Another disadvantage is that a model is only as accurate as the means used to characterize the model. In the work on the Search Dashboard, I used the simple metric of the use of advanced search engine functionality to identify ‘experts’ from search engine logs. The method was effective because the behavioural characteristic of expert searchers was already well understood, and the technique itself has been shown to accurately identify expert searchers [158]. However, if some other technique had been used, and expertise was misrepresented, then searchers using the Dashboard may have been misled.

Even though aggregated behaviour can lend itself to global learning, the identifiable traces used by WebWear may support global learning as well. An example of global learning was described in the field trial of the WebWear system. One participant encountered a search query for something seemingly strange: “105020*4*3*4”. By following the trace for this query, it was revealed to the participant that the query was making use of the search engine as a calculator. However, using WebWear to get a sense of the overall behavioural patterns of expert searchers – for example, how many words experts use in their search queries – would have to be accomplished by examining each query individually and would be cumbersome if not impossible.

WebWear lends itself to local learning because the specific details of interaction history are available, and this information is essential for working out the context of observed behaviour. However, the identifiable traces raise privacy concerns. This contextual information would be hidden in the Dashboard, which also makes it unsuitable for local learning because the information it offers would be too broad.

While the Search Dashboard and WebWear have focused on supporting only a single type of learning each, a single system could support both global and local learning. To do this, the system would need to support both approaches: the aggregation of high-level behaviour for revealing patterns of use, and providing the details of lower-level traces available to reveal the contextual details needed to infer advice in specific situations.

7.5.3 The Importance of Loosely-Coupled Information Seeking

In this dissertation it was necessary to provide a definition of loosely-coupled information seeking tasks (Section 1.1.1). In this section I reflect on loosely-coupled information-seeking tasks, their importance in daily work, and how social feedback was shown to support them.

Outside of information seeking, most of the interactions we have with other people are informal in style, in the workplace or otherwise: they are unscheduled communications without prescribed formats and they occur opportunistically [89]. Work and communication occur in this way because it is human nature. As described by the principle of least effort (see Section 2.2.2), people often prefer the simplest course of action to acquire information that is minimally acceptable to complete a main task. This is why people view receiving guidance from others as discretionary and also likely why people prefer loosely-coupled collaborations in many situations – they are quick and easy, which allows the focus to remain on an overarching task.

Loosely-coupled information-seeking tasks occur when other people have experience or knowledge that could help in completing the task. Outside of web-based information seeking, there are many scenarios in which other people have experience or knowledge that could help us accomplish tasks more successfully. From choosing a healthy meal to changing a car's windshield wipers, everyday tasks can be facilitated by the knowledge of others. The challenge is providing people with the guidance they need, because people often do not look for guidance otherwise.

The characterization of LCIS (Section 1.1.1) extends current understanding of the social setting for information seeking by recognizing that guidance from others can be useful even when it is not asked for, and that the support offered must require minimal effort. In many tasks people prefer the risk of inefficiency or failure in information seeking, rather than incurring the overhead of asking for help and guidance from other people on the off chance that they could provide useful information. When considering LCIS against more formal styles of collaboration, the focus naturally changes from explicit forms of communication and collaboration to more passive forms that better fit the discretionary nature of help in such tasks. Social feedback can successfully support loosely-couple information seeking situations because it is a low-effort solution.

7.5.4 Information Seeking is Essential

Marchionini stated that "information seeking is a fundamental human process closely related to learning and problem solving… thus [it is] a natural and necessary mechanism of human existence ([95], p. 5-6)." In this dissertation, I have focused on an extremely familiar form of information seeking – web-based information seeking using a standard web browser. However, information seeking is a much broader activity. Often times in the process of learning or solving problems we must engage in some form of information seeking, to fill the gaps in knowledge that we encounter. These information seeking tasks can make use of the Web, but they also involve other people, books, work materials, or the physical environment. Therefore, social feedback could be generalized to other scenarios where support is required to improve information seeking success (possibilities for generalizing social feedback are discussed in Section 8.2.3).

7.6 Scope and Limitations

The research conducted in this dissertation has been necessarily limited by the constraints of time, breadth, depth, and the methods selected. I have shown that social feedback works as designed, by supported learning in information seeking activities. However, the question remains whether the approach is feasible over a longer period of time. It is unknown whether newly learned behaviours will continue to be used and whether people will continue to find value in the systems after months of use.

With systems like WebWear the main question is whether the system will continue to be used. In WebWear there needs to be a balance between seekers being able to use the information when it is needed and disregarding it when it is not. If the traces presented by WebWear are not useful frequently enough, people may stop using the system. However, over time and with more familiarity people may be able to attend to information selectively, and begin to rely on WebWear as an important part of their information seeking and collaborative work practices.

The Dashboard is intended to be used infrequently (because that is all that is needed for meaningful learning to occur), and so the main question is whether learned behaviour continues to be used over time. It could be that over time people fall back into their original style of search engine use. However, people may also use the system for feedback on how the new skills and techniques they have applied have changed their behaviour over time, and how they compare to expert searchers after changing their behaviour. To understand if, and how, the Dashboard would be used over a longer term will require further study.

The path forward for answering these questions is potentially challenging, because it involves deployment of systems in real use over a long period of time. Over a long period, new questions would arise. For example, in the case of WebWear there may be a relationship between the signal-to-noise ratio and the amount of data available in the system. Over a longer period of time, WebWear would collect much more data than what had been seen in the one-week trial. It could be that the amount of data actually decreases the signal-to-noise ratio, making it difficult to find useful behaviour among all of the presented traces. Further, privacy concerns may change over time. On one hand, the behaviour that is shared today may be taken out of context by other people in a year's time. On the other hand, with more familiarity and experience with the technology, people may have fewer concerns about sharing their behaviour, and realize that the actual risks are few. Findings such as these would emerge over time and lead to new research questions, such as: How long is interaction history relevant? For how long will people recall the contextual details of other's behaviour? Should the details of interaction history, such as identity, fade over time? When should a social feedback system "forget" interaction history? This dissertation has provided a basis for social feedback systems, and shown that they are feasible and useful under real-world conditions, at least in the short term; but longer term research is now warranted.

7.7 Reflection on Broader Motivations

My work in this dissertation has also been motivated by several wider questions. First, and most broadly, we currently have unprecedented abilities to capture data in a networked environment. In the past, the technical challenges associated with the volume of interaction history and the limitations of network bandwidth made conducting research in this area difficult. However, data storage and network connectivity are no longer major hurdles to systems that rely on interaction history. Given these new abilities, what can we do with this type of information (implicitly created traces of interaction history)? This work makes modest steps towards the original visions from Hill et al.'s work on Edit Wear and Read Wear [70], in which people's interaction history could provide guidance and support to others.

Second, we now have tools that allow us to easily share information with others, to keep track of work that needs to be done, and to remind us of important events. But the sheer volume of these purpose-specific tools creates a new problem – the need to spend time managing the tools themselves. This takes effort and time, and does not fit into the discretionary use of such tools. The research in this dissertation provides some insights into how new tools may be built to better aid individuals and groups in carrying out work, to capture and represent knowledge, and to share experience; with minimal effort.

For the communities of human-computer interaction, information seeking, and computer-supported cooperative work, this work investigates a challenging question: how can traces of work be incorporated into digital systems and used in the real world? There have been several compelling examples of the use of social navigation that have afforded users a better understanding of the digital objects they are exploring and interacting with. Further, such tools have been shown to provide performance improvements for common refinding and navigation tasks (e.g. [2]). However, these have all either been in a controlled setting with a specific task and/or over a controlled set of digital objects. It is critical for continued research in this area to have a better understanding of how interaction history works in less controlled and natural settings, such as real information-seeking tasks on the Web.

CHAPTER 8
CONCLUSION

Information seeking on the Web is an extremely common and important activity that is used to support many tasks. Despite the fact that people use the Web every day to look for information, information seekers are often not as successful as they potentially could be in completing their tasks. Many of the situations where people have difficulties are loosely-coupled information-seeking tasks – where other people have experience or knowledge that could help in the task. Such help is discretionary, since we could carry out the task in some fashion on our own, but without help there are often missed opportunities where tasks could have been completed more effectively or efficiently.

To address this problem I have described the idea of social feedback, which is based on social learning theory and instantiated through the mechanism of interaction history. In this dissertation I have shown that social feedback provides a way of supporting loosely-coupled information seeking, by enabling information seekers to learn from the experience of others. I have identified two sources of information from whom we can learn: experts (people who we admire or have some prestige with regards to a task we have) and tightly-knit contacts (people with whom we share interests, tasks, and goals). Social feedback allows learning in two ways: global learning and local learning. Global learning refers to learning new skills or acquiring new behaviour that can be applied in many different information seeking situations. Local learning refers to learning that occurs within a specific context or situation, and can lead to specific task instances being accomplished more effectively or efficiently.

In this dissertation, I first presented a conceptual framework that shows how social feedback systems can support information seekers in learning how to complete their information-seeking tasks more successfully. I then described the design and evaluation of the Search Dashboard, a social feedback system that demonstrates global learning of search engine skills. Next, to demonstrate how a social feedback system could support local learning and allow seekers to complete common tasks with more success, I described the development and evaluation of the WebWear system. I then explored the two main challenges to the successful deployment of social feedback systems: *the overlap challenge*, which concerns the overlap between information seekers that provides opportunities for social feedback to be given, and *the privacy challenge,* which concerns people's willingness to share interaction history to enable social feedback. Finally, through two separate studies I provided evidence that the overlap and privacy challenges will not prevent the successful real world deployment of social feedback systems.

8.1 Contributions

The main contribution of this work is the idea of social feedback, which provides support from the knowledge and experiences of others in information seeking. Social feedback is based on social learning theory – which describes how people learn from observing the actions of others – and is instantiated through the mechanism of interaction history – which is the passive collection and use of interactions with technology to describe behaviour. Social feedback supports loosely-coupled information seeking by allowing information seekers to learn how to complete their tasks more successfully through observing how other people have behaved in similar situations.

This work has also provided several minor contributions:

1. A conceptual framework that can be used in the design of social feedback systems.
2. The Search Dashboard system, which demonstrates how social feedback systems can be designed to support global learning.
3. Empirical evidence that the Search Dashboard can enable the acquisition of expert skills and behaviour in the real world use of search engines.
4. The WebWear system, which demonstrates how social feedback systems can be designed for local learning.
5. Empirical evidence that information seekers can learn how to accomplish loosely-coupled collaborative information-seeking tasks more successfully.
6. An initial exploration of the overlap challenge indicating that overlap is not a major concern for the long-term deployment of social feedback systems in the real world.
7. An initial exploration of the privacy challenge indicating that privacy is not a major concern when social feedback systems are used in the appropriate contexts, and that the context of workgroups is appropriate for the use of identifiable interaction history.

8.2 Future Research

This research opens up several directions for future research. Below, I organize these future directions by the following themes: social feedback system features, further evaluations, and the broader application of social feedback outside of information seeking.

8.2.1 Social Feedback System Features

There are four main directions for exploring new social feedback system features: the visual design space for social feedback, mechanisms for sharing interaction history, mechanisms for broadening the display of traces, and finding new tightly-knit contacts.

8.2.1.1 Exploring the Visual Design Space for Social Feedback

The visual displays selected for the Dashboard and WebWear were based on identifying simple visualizations that met the basic information requirements for the particular situation and the learning goal. However, many other visual encodings for the information are possible. A systematic exploration of the design space of visual encodings seems appropriate for future research on social feedback. It could be that certain visual encodings simply do not work for conveying elements of behaviour, while others could improve the learning outcomes.

Further, because the deployments of both the Dashboard and WebWear were relatively short, temporal information was not needed for learning in the tasks. WebWear does encode the age of activity within the glyphs: as the activity becomes older the color of the glyphs begins to fade (see Figure 5.2). The Search Dashboard does not show change in behaviour over time, but it may be particularly useful for people to see how their behaviour has changed over time. Encoding such temporal information should be considered for longer term deployments of the systems.

8.2.1.2 Interaction History Sharing Mechanisms

As described in Chapter 6 (on the challenges of social feedback), privacy concerns are an important challenge facing social feedback systems. While I have shown that there exist current contexts where individually-identifiable interaction history can be appropriately shared with low levels of concern for privacy (in tightly-knit workgroups), there are technological approaches that may further reduce the overall level of concern. However, in exploring new mechanisms for sharing interaction history a balance needs to be struck between the effort of managing the distribution and the perceived risks to privacy.

WebWear has two simple mechanisms for controlling the distribution of interaction history: collection can be stopped and started, and past interaction history can be permanently deleted. These two mechanisms are fairly simple, but seem to have been effective.

Further privacy-protection mechanisms could be explored, such as whitelists and blacklists for websites (e.g., I may be willing to share activity on Wikipedia, but keep all my YouTube activity private). Lists of people could also be used for directed sharing. Contacts could be organized into useful groups for sharing and might be combined with whitelists and blacklists of websites (e.g., I am willing to share my activity on Wikipedia with my colleagues because it is most often related to work, but I am only willing to share my activity on YouTube with my wife because she won't judge me by my viewing habits). Similar approaches exist in social networking services, such as Google+[17], for the explicit sharing of information. However, such a solution introduces a substantial cost: the need to define and maintain the group lists which represent who would receive interaction history and when they might receive it. In recognition of such overhead, research has proposed semi-automated systems to facilitate the identification of relevant contacts for sharing information [6].

Another interesting privacy-protection mechanism is providing complete control of interaction history, including its storage, on a user's computer. In both WebWear and the Search Dashboard, interaction history is stored on a remote server. This means that the data is ultimately out of the control of a system user. Recently there has been an effort to provide general purpose information storage platforms that allow data, such as interaction history, to be stored on the user's computer, and for which the user can control distribution by granting access to individual applications and people [152].

8.2.1.3 Mechanisms for Appropriately Broadening the Display of Traces

WebWear broadens the situations where interaction history can be encountered by displaying interaction history describing website-level activity and not just activity on particular webpages (see Section 5.1.4). Providing website-level support allows social feedback to be available in more situations, because a single website may contain many pages that have been

---

[17] See https://plus.google.com/

interacted with. Websites were selected as the means by which to broaden support because it is a simple means that can keep the activity displayed relevant to a particular task (since entire tasks are often completed on a single website). However, as discussed in the report on the field trial (see Section 7.4.1.1), there were instances in which this simple approach was too broad. Different mechanisms that allow the appropriate broadening of displayed traces should be investigated.

One approach to this problem would be to create a list of websites where the current website-based approach is appropriate. For example, in the field trial we found that sharing activities in this way was not particularly useful for google.ca (because too many different tasks are completed through searching Google), but it was appropriate for stackoverflow.com (information seeking on this website was always for programming help). Other approaches for narrowing the activity of websites could be explored such as using subdomains to restrict the activity (e.g., activities on google.ca could be narrowed by using sub-domains, such as news.google.ca or maps.google.ca).

Further, semi-automated approaches could be explored that would attempt to identify and display only activities that are associated with the task. This would require some overhead for a user to confirm that activity is indeed associated with a particular task. However, it would provide the further advantage that people could explore all activities associated with a completed task.

8.2.1.4 Finding New Tightly-Knit Contacts

The studies showed that an essential part of the success of future social feedback systems for local learning will be due to the properties of tightly-knit groups. However, there were likely other opportunities where the WebWear field-trial participants could have used some guidance, but none was available. Future work should consider how social feedback systems could address can appropriately maximize the size of tightly-knit groups to increase availability of support. While automated approaches may be considered – for example, by providing the traces of some unknown person who had the exact same task at some point in the past – this seems a challenging problem for the short term. Another way that could be explored is through encouraging people to expand the basis of what they would consider to be tightly-knit through recommendations of people who share similarities or who have tightly-knit contacts in common. Because trust is such a central property within tightly-knit groups, the system would likely have to provide important details that would allow a user to decide if they can or should trust a potential new contact. Exploring the ways that tightly-knit groups could be grown could be a fruitful direction; a WebWear user with a larger set of contacts would have a higher likelihood of receiving support when it was needed, by having the knowledge and experience of more people available.

8.2.2 Questions for Evaluation

This research has shown that social feedback is feasible; however, there are additional evaluations that can provide further evidence that the approach will be successful, and that it can be applied to other scenarios in information seeking. In addition, further study could shed light on how the learning of behaviour could be strengthened.

8.2.2.1 Longer Term and Wider Deployment of Social Feedback

As previously discussed, a main limitation of this work was the length of the deployment (see Section 7.6). I see two further questions that should be pursued in future research with regards to the Dashboard. First is the question of whether or not learning and adopting the behaviour of expert searchers leads to improved search success. This question is difficult to answer, because success is difficult to measure in real world information seeking. However, previous work has used human annotators to manually score the success of a searcher by examining summaries of the interaction history [158]. This approach could be employed to determine whether the Search Dashboard actually leads searchers who adopt expert behaviour to improve their search success.

The other way that the Search Dashboard study was limited was in that it guided study participants through the Dashboard, and may have induced closer consideration of the information than would have otherwise occurred. This being said, the effect of the study was minimal compared to the effect of comparison: the comparison interface (based on social feedback) was shown to be what led to behaviour change, because participants in the non-comparison condition did not change their behaviour. However, the question remains about whether unguided interactions with the Search Dashboard could still result in global learning of search skills and behaviour.

8.2.2.2 Application and Evaluation of Further Global Learning Scenarios

The work on local learning (demonstrated through WebWear) was shown to provide support for all three basic interactions that are part of information-seeking tasks (using a search engine, selecting links, and reviewing webpages). However, the work on global learning with the Search Dashboard demonstrated support for using search engines only. Further work in global learning and information seeking should target the two other basic interactions: selecting links and reviewing webpages. Selecting links could be supported by helping seekers learn how experts assess the information scent of a link (e.g., by using proximal cues such as the text displayed in the link, the text surrounding the link, and the URL of the link – see Section 2.1.2.1). Similarly, reviewing webpages could be supported by assessing the credibility of the content on a page (e.g., [60,132]). Again, a system could help a seeker learn how to assess the credibility of webpage content as an expert would. Like the Search Dashboard, these two new systems would have to display information based on a solid understanding of the behaviour that experts employ and what the important factors are in assessing the information scent of a link and the credibility of webpage content.

8.2.2.3 The Effect of Extrinsic Motivation on Learning

In both the Search Dashboard and WebWear studies, participants were not offered any rewards based on learning or performance. The studies instead relied on participants to be intrinsically motivated to make use of and learn from the behaviour presented by the systems. The fact that there were positive results in both real world search engine use (after viewing the Search Dashboard) and in the field trial of WebWear shows that people can be sufficiently motivated to learn new behaviour by the promise of improved success in information seeking. However, social learning theory suggests that extrinsic motivation could strengthen learning in some cases or for some individuals (see Section 3.1.1 and 3.1.4). Various theories of learning have proposed that extrinsic motivation can be offered through competition, cooperation, and recognition (see Section 3.1.5). These are also a subset of motivating elements that have been identified in gamification research (how elements of games can be added to encourage completion of otherwise mundane tasks) [58]. The ways that extrinsic motivation could be added to local and global learning systems, and what effect it might have on learning and task success, is an interesting topic for future research in social feedback.

8.2.3 The Application of Social Feedback Outside of Information Seeking on the Web

In this work, social feedback has been described within the context of information seeking on the Web. However, the same basic idea of social learning theory instantiated through interaction history could be applied in many other situations where learning from observation could be beneficial. Below, I provide three further applications of social feedback: in computer programing, in general user interfaces, and in ubiquitous computing.

8.2.3.1 Support in Learning Expert Programming

Computer programmers must keep abreast of many new technologies and languages, each with their own set of conventions, best practices, and support libraries. This means that programmers must dedicate a great deal of time to researching and staying on top of this information in order to be proficient. Further, gaining expertise at programming often benefits from contact and guidance from others, so that new techniques can be learned and current approaches can be discussed and critiqued. Many programmers do not have the opportunity to discuss their code or learn from others who are more proficient; much of a programmer's working time is spent alone. Part of the problem is that, like web tools, programming environments (IDEs) do not offer many facilities to learn from others. Social feedback could be applied directly to IDEs both for global learning and for local learning.

Part of the inspiration for this work came from a local learning tool for a programming editor: the Edit Wear and Read Wear tool described in Section 2.6.2 (see Figure 2.8). Edit Wear and Read Wear were based on the zmacs programming editor, and allowed collaborators to answer awareness related questions such as: Who edited this line? Which part of code has been edited most frequently? Which part of the code is the newest? Future work would examine how these original visions could become commonplace in IDEs and could incorporate the lessons of social feedback.

A global learning tool in an IDE would allow a programmer to learn best practices from expert programmers. First, an appropriate metric would be investigated for identifying expert programmers, such as the number of lines written per day or the total amount of time spent in a particular language. Once experts are identified, they could be used as models to learn from. This would allow programmers to understand how their programming style and techniques could be improved, by being able to answer questions such as: Do my functions contain too many lines of code? Do I compile my code too often? Do I use infrequently-used functions? Am I a slow programmer? Should I spend more time reading code, before I start editing existing code? What features of the debugger do I use less often than an expert? Answering these questions would be extremely valuable, and help programmers improve their craft.

8.2.3.2 Support in General User Interfaces

The previous example of the use social feedback within a programming IDE also brings up the idea of further generalization to other user interfaces. In the same way that current web and programming tools do not support learning from others, many system interfaces have the same limitation. This is despite the fact that applications (like Microsoft Word, Excel, or PowerPoint) are extremely complex pieces of software, contain hundreds of commands, and are used by people every day. Again the same ideas can apply for global learning – How often do expert Excel users use shortcut keys? What shortcut keys do they make use of most frequently? – and local learning – How did Jim embed that video in his presentation? Has Sally reviewed my edits to her work yet? One might also consider how operating system-level social feedback may be offered to provide support through any application (recent work as shown that this may be technically feasible for a large subset of applications, e.g. [53]).

8.2.3.3 Ubiquitous Social Feedback

Research in ubiquitous computing investigates how useful system interactions can be moved away from traditional computers into the physical environment [154,155]. Social feedback need not be tied to a traditional desktop computer and so can be applied in the context of ubiquitous computing as well. Current mobile phones offer input capabilities through a number of sensors (to collect interaction history from the physical world), network capabilities (to send and receive descriptions of behaviour), and high-resolution screens (to display representations of behaviour). Because we have technology that allows us to make the core parts of social feedback systems mobile, learning from the behaviour of others can be extended into the environment around us.

For example, imagine a system to help someone improve their physical fitness through jogging. A phone application could be developed to track people’s activity as they jog (i.e., their interaction history), and it could then identify models (other system users) for them to learn from. Models may be ‘expert’ joggers, but they may also be people who a user can self-identify with (a role model). These role models could be people who were once at the same level of ability as the user but have succeeded in reaching a goal that the user would also like to reach (e.g., being able to run 10 km). Through tracking and comparing their jogging routine with their models, a user could receive appropriate feedback on their progress as compared to the past experiences of their role model. Throughout a user's progress toward their goal, they would be able to have clear and realistic targets based on how someone just like them had performed. This approach to providing a social learning experience may provide people with more relevant and meaningful support than typical approaches which simply display user's personal performance or provide a typical training programme.

In the future one could imagine many other ways that people could benefit from one another's experiences – for example, to learn time management skills, healthy eating, or to quit smoking. Through social feedback we have a general approach that can allow us to learn from one another and build on each other's experiences and successes.

REFERENCES

1. Adar, E., Teevan, J., and Dumais, S. T. Resonance on the web: web dynamics and revisitation patterns. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '09)*, 2009, 1381-1390.

2. Alexander, J., Cockburn, A., Fitchett, S., Gutwin, C., and Greenberg, S. 2009. Revisiting read wear: analysis, design, and evaluation of a footprints scrollbar. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '09)*, 2009, 1665-1674.

3. Altman, I. *The Environment and Social Behaviour*, Monterary, CA, USA. Brooks/Cole Publishing, 1975.

4. Amershi, S. and Morris, M.R. Co-located Collaborative Web Search: Understanding Status Quo Practices. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '09)*, 2009, 3937-3640.

5. Amershi, S. and Morris, M.R. CoSearch: A System for Collocated Collaborative Web Search. In . In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '08)*, 2008, 1647-1656.

6. Amershi, S., Fogarty, J., and Weld, D. Regroup: interactive machine learning for on-demand group creation in social networks. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '12),* 2012, 21-30.

7. Aula, A. & Nordhausen, K. Modelling successful performance in web searching. *JASIST*, *57,* 12, 2006, 1678-1693.

8. Aula, A., Jhaveri, N. & Kaki, M. Information search and re-access strategies of experienced Web users. *Proc. WWW*, 583-592.

9. Aula, A., Khan, R.M. & Guan, Z. How does search behaviour change as search becomes more difficult? In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '10)*, 2010, 35-44.

10. Baecker,R.M, Grudin, J., Buxton,W.A.S., and Greenberg, S., Eds. *Readings in Human-Computer Interaction: Toward the Year 2000*. Morgan Kaufmann, San Mateo, CA. 1995.

11. Bandura, A. Social cognitive theory of self-regulation. *Organizational Behaviour and Human Decision Processes*, 50, 2, December 1991, 248-287.

12. Bandura, A. *Social Foundations of Thought and Action: A Social Cognitive Theory*. Englewood Cliffs, NJ, Prentice-Hall. 1986.

13. Bateman, S., Brooks, C., McCalla, G., and Brusilovsky, P. Applying Collaborative Tagging to E-Learning. In the Proceedings of the *Workshop on Tagging and Metadata for Social Information Organization*, held in conjunction with the *Proceedings of the International Conference on the World Wide Web (WWW '07)*, 2007. 7 pages. Available at: http://www2007.org/workshops/paper_56.pdf

14. Bateman, S., Gutwin, C., Osgood, N., and McCalla, G. Interactive usability instrumentation. In the *Proceedings the ACM SIGCHI Symposium on Engineering Interactive Computing Systems (EICS '09*), 2009, 45-54.

15. Belkin, N. J. Information Concepts for Information Science. *Journal of Documentation*, 34, 10, 1978, 55-85.

16. Bellotti, V., and Bly, S., Walking Away from the Desktop Computer: Distributed Collaboration and Mobility in a Product Design Team, In *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '96)*, 1996, 209-218.

17. Bennett, P., Svore, K. & Dumais, S. Classification-enhanced ranking. In the *Proceedings of the International Conference on the World Wide Web (WWW '10)*, 2010, 111-120.

18. Benyon, D. The New HCI? Navigation of Information Space. *Knowledge-Based Systems*, 14, 8, December 2001, 425-430.

19. Borgatti, S.P. and Cross, R. A Relational View of Information Seeking and Learning in Social Networks. *Management Science*. 49, 4, 2003, 432-445.

20. Boud, D., Keogh, R. & Walker, D. *Reflection: Turning Experience into Learning.* Routledge, 1985.

21. Boyle, M. and Greenberg, S. The language of privacy: Learning from video media space analysis and design. ACM Transactions on Computer-Human Interaction. 12, 2, June 2005, 328-370.

22. Brand-Gruwel, S., Wopereis, I. & Vermetten, Y. Information problem solving by experts and novices: Analysis of a complex cognitive skill. *Computers in Human Behaviour*, *21*, 2005, 487-508.

23. Broder, A. A taxonomy of web search. *ACM SIGIR Forum*, 36,2, Fall 2002.

24. Brush, A.J.B., Meyers, B., Scott, J., and Venolia, G. Exploring awareness needs and information display preferences between coworkers. . In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '09)*, 2009, 2091-2094.

25. Brusilovsky, P. User Modelling and User Adapted Interaction, 11(1-2), 2001, 87-110.

26. Brusilovsky, P., Chavan, G., and Farzan, R. Social adaptive navigation support for open corpus electronic textbooks. In: P. De Bra and W. Nejdl (eds.) Proceedings of *The International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'2004)*, Eindhoven, the Netherlands, 2004, 24-33.

27. Buckland, M. K. Information as thing. *Journal of the American Society for Information Science (JASIST)*, 42, 1991, 351–360.

28. Bush, V. As we may think. *Atlantic Monthly*, 176, 1, 1945, 101–108.

29. Card, S. Information Visualization, in A. Sears and J.A. Jacko (eds.), *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies, and Emerging Applications*, Lawrence Erlbaum Assoc Inc, 2007.

30. Card, S. Pirolli, P., Van der Wege, M., Morrison, J. B., Reeder, R. W., Schraedley, P. K., and Boshart, J. Information scent as a driver of web behaviour graphs. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '01)*, 2001, 498–505.

31. Carmel, E., Crawford, S. and Chen, H. Browsing in Hypertext: a cognitive study. *IEEE Transactions on Systems, Man and Cybernetics*, 22, 1992, 865-884.

32. Carnegie Mellon University. Press Release: Picking the Brains of Strangers Improves Efforts To Make Sense of Online Information. *Press Release*. May 7th, 2012. Available at http://www.cmu.edu/news/stories/archives/2012/may/may7_distributedsensemaking.html. Accessed: June 7th 2012.

33. Case, D.O. Principle of least effort. In K. E. Fisher, S. Erdelez & L. McKechnie, *Theories of information behaviour*. ASIST Monograph Series, Medford, New Jersey. Information Today, Inc, 2005.

34. Chi, E.H. Information Seeking Can Be Social. *IEEE Computer*, 42, 3, March 2009, 42-46.

35. Choo, C.W., Detlor, B., and Turnbull, D. Information seeking on the Web: An integrated model of browsing and searching [Online Serial]. *First Monday*, 5, 2, 2000.

36. Claypool, M., Le, P., Wased, M., and Brown, D. Implicit interest indicators. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI '01)*, 2001, 33-40

37. Cockburn, A. Greenberg, S. McKenzie, B., Smith, J. M. Kaasten, S. Webview: a graphical aid for revisiting web pages, In the *Proceedings of the Computer Human Interaction Specialist Interest Group of the Ergonomics Society of Australia (OzCHI'99)*, 1999, 15- 22.

38. Collins, A., Brown, J.S. & Newman, S.E. Cognitive apprenticeship: Teaching the craft of reading, writing and mathematics. *Knowing, Learning and Instruction*, 8, 1, 1989, 453-494.

39. Consolvo, S., McDonald, D. & Landay, J. Theory-driven design strategies for technologies that support behaviour change in everyday life. In the . In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '09)*, 2009, 405-414.

40. Cosley, D., Lam, S. K., Albert, I., Konstan, J., and Riedl, J. Is Seeing Believing? How Recommender Systems Influence Users' Opinions. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '03)*, 2003, 585-592.

41. Davis, S., Gutwin, C. Using Relationship to Control Disclosure in Awareness Servers. In the *Proceedings of Graphics Interface (GI '05)*, 2005, 75-84.

42. Dennett, D. C. *Darwin's dangerous idea*. New York: Simon and Schuster, 1995.

43. Dervin, B. Sense-Making theory and practice: an overview of user interests in knowledge seeking and use. *Journal of Knowledge Management*, 2, 2, 1998, 36-46.

44. Dervin, B. Useful theory for librarianship: Communication, not information. *Drexel Library Quarterly,*13, 3, 16-32.

45. Dieberger, A., and Lonnqvist, P. Visualizing Interaction History on a Collaborative Web Server. In the *Proceedings of the Conference on Hypertext and Hypermedia (HT '00)*, 2000, 221-222.

46. Dieberger, A. Where Did All the People Go? A Collaborative Web Space with Social Navigation Information. Poster presented at the *Proceedings of the International Conference on the World Wide Web (WWW '00)*. Available online at: http://homepage.mac.com/juggle5/WORK/publications/SwikiWriteup.html. Accessed: March 3, 2011.

47. Dimitrova, V. StyLE-OLM: Interactive Open Learner Modelling. *International Journal of Artificial Intelligence in Education*, 13,1, 2003, 35-78.

48. Dou, Z., Song, R. & Wen, J. A large-scale evaluation and analysis of personalized search strategies. In the *Proceedings of the International Conference on the World Wide Web (WWW '07)*, 2007, 581-590.

49. Dourish, P. and Bly, S. Portholes: Supporting Awareness in a Distributed Work Group. In the *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '92)*, 1992, 541-547.

50. Dourish, P. and Chalmers, M. Running out of space: Models of information navigation. In the *Proceedings of the Conference on Human Computer Interaction* (*HCI '94*), 1994. 4 pages.

51. Dourish, P. Where the Footprints Lead: Tracking Down Other Roles for Social Navigation. In D. Benyon, A. Munro and K. Höök (Eds.) *Designing Information Spaces: The Social Navigation Approach*, New York, Springer, 2003, 83-104.

52. Dourish, P., and Bellotti, V. Awareness and coordination in shared workspaces. In *Proceedings of the 1992 ACM conference on Computer-Supported Cooperative Work (CSCW '92),* 1992, 107-114.

53. Eagan, J., Beaudouin-Lafon, M., and Mackay, W. Cracking the cocoa nut: user interface programming at runtime. In the *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '11)*, 2011, 225-234.

54. Evans, B.M., and Chi, E.H. An elaborated model of social search. *Information Processing and Management*,46, 6, November 2010, 656-678.

55. Farzan, R. *A Study of Social Navigation Support Under Different Situational and Personal Factors*. PhD Dissertaion. University of Pittsburgh. 2009.

56. Festinger, L. A theory of social comparison process. *Human Relations,* 7, 1954, 117-140.

57. Fisher, K., Counts, S., and Kittur, A. Distributed sensemaking: improving sensemaking by leveraging the efforts of previous users. In Proceedings of the *ACM Annual Conference on Human Factors in Computing Systems (CHI '12)*, 2012, 247-256.

58. Flatla, D.R., Gutwin, C., Nacke, L.E., Bateman, S., and Mandryk, R.L. Calibration Games: Making Calibration Tasks Enjoyable by Adding Motivating Game Elements. In the *Proceedings of the Symposium on User Interface Software and Technology (UIST '11)*, 2011, 403-412.

59. Fogg, B.J. Motivating, influencing, and persuading users. In Julie A. Jacko and Andrew Sears (Eds.). L., *The Human-Computer Interaction Handbook,* CRC Press, 2002, 358-370.

60. Fogg, B.J. *Persuasive Technology: Using Computers to Change What We Think and Do*. Morgan Kaufmann Publishers, San Franciso, 2003.

61. Froehlich, J., Findlater, L. & Landay, J. The design of eco-feedback technology. In the *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI '10)*, 2010, 1999-2008.

62. Furnas, G.W. Effective View Navigation. In the *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI '97)*, 1997, 367-374.

63. Goecks, J. *Understanding the Social Navigation User Experience*. PhD Dissertation. Georgia Institute of Technology. August 2009.

64. Goecks, J., and Mynatt, E.D. Social Approaches to End-User Privacy Management. In L. Cranor & S. Garfinkel (Eds.), *Security and Usability: Designing Secure Systems That People Can Use*, O'Reilly, 2005.

65. Golovchinsky, G., Pickens, J., and Back, M., A Taxonomy of Collaboration in Online Information Seeking, In the *Proceedings of the International Workshop on Collaborative Information Retrieval*, June 2008. 4 pages. Available at: http://arxiv.org/ftp/arxiv/papers/0908/0908.0704.pdf

66. Golovchinsky, G., Qvarfordt, P., and Pickens, J. Collaborative information seeking. *IEEE Computer*, 42,3, 2009, 47-51.

67. Gutwin, C., and Greenberg, S. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work,* 11, 3, November 2002, 411-446.

68. Gutwin, C., Greenberg, S., Blum, R., Dyck, J., Tee, K. and McEwan, G. Supporting Informal Collaboration in Shared-Workspace Groupware. *Journal of Universal Computing (JUCS)*, 14, 9, May 2008, 1411-1434.

69. Hawkey, K. and Inkpen, K.M. Privacy Gradients: Exploring Ways to Manage Incidental Information During Co-Located Collaboration. In the *Proceedings of the Extended Abstracts on Human Factors in Computing Systems (CHI EA '05),* 2005, 1431-1434.

70. Hill, W. C., Hollan, J. D., Wroblewski, D., and McCandless, T. Edit wear and read wear. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '92),* 1992, 3-9.

71. Hollan, J., Hutchins, E., and Kirsh, D. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transaction of Computer-Human Interaction.* 7, 2, June 2000, 174-196.

72. Hölscher, C. & Strube, G. Web search behaviour of Internet experts and newbies. *Computer Networks*, 33, 2000, 337-346.

73. Hook, K., Benyon, D., and Munro, A. Editor's Introduction: Footprints in the Snow. In: Hook,K., Benyon,D.,Munro,A.(Eds.), *Social Navigation of Information Space*. Springer-Verlag, London, 2003, 1-13.

74. Hudson, S E., and Smith, I. Techniques for Addressing Fundamental Privacy and Disruption Tradeoffs in Awareness Support Systems. In the *Proceedings of the Conference on Computer Supported Cooperative Work*, 1996, 248-257.

75. Indratmo, Vassileva, J. Social Interaction History: A Framework for Supporting Exploration of Social Information Spaces. In the *Proceedings of International Conference on Computational Science and Engineering, IEEE International Conference*, 2009, 538-545.

76. Kaasten, S. and S. Greenberg. Designing an integrated bookmark / history system for Web browsing. In the Proceedings of the *Western Computer Graphics Symposium*, 2000, 4pages. Available at: http://grouplab.cpsc.ucalgary.ca/grouplab/uploads/Publications/Publications/2000-Kaasten.Skigraph.pdf

77. Kärkkäinen, T., Vaittinen, T., and Väänänen-Vainio-Mattila, K. I don't mind being logged, but want to remain in control: a field study of mobile activity and context logging. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '10)*, 2010, 163-172.

78. Kawai, T., Bannai, Y. and Tamura, H. ARGUS: An Active Awareness System Using Computer-Controlled Multiple Cameras Video Program, In the *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '96),* 1996, 7.

79. Kay, J., and Kummerfield, B. Lifelong User Modelling Goals, Issues and Challenges. In the *Proceedings of Life Long User Modelling Workshop*, in conjunction with *User Modelling, Adaptation and Personalization*. June 22-26, 2009, Trento, Italy.

80. Kay, J., and Thomas, R. C. Studying Long-Term System Use. In the *Communications of the ACM*, 38, 7, July 1995, 61-69.

81. Kellar, M., Watters, C., and Sheppard, M. A field study characterizing Web-based information-seeking tasks. *Journal of the American Society for Information Science and Technology (JASIST)*, 58, 7, May 2007, 999-1018.

82. Kelly, D. and Teevan, J. Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum*, 37,2, September 2003, 18-28.

83. Khan, K. & Locatis, C. Searching through the cyberspace: the effects of link display and link density on information retrieval from hypertext on the World Wide Web. *Journal of the American Society for Information Science and Technology (JASIST*), *49*, 2, 1998, 176-182.

84. Klein, G., Moon, B., and Hoffman, R. F. Making sense of sensemaking I: Alternative perspectives. *IEEE Intelligent Systems*, 21, 4, 2006, 70–73.

85. Kleinberg, J. M. Authoritative sources in a hyperlinked environment. *Journal of the ACM*, 46, 5, 1999, 604-632.

86. Kobsa, A. Generic user modelling systems. In P. Brusilovsky, A. Kobsa, and W. Nejdl, (Eds.) the *Adaptive Web: Methods and Strategies of Web Personalization*, Lecture Notes In Computer Science, vol. 4321. Springer-Verlag, Berlin, Heidelberg, 2007, 136-154.

87. Kodagoda, N. and Wong, B.L.W. Effects of low and high literacy on user performance in information search and retrieval. In the *Proceedings of the Conference on Human-Computer Interaction*, 2008, 173-181.

88. Konstan, J.A., and Riedl, J. Collaborative Filtering: Supporting Social Navigation In Large, Crowded Infospaces, In *Designing Information Spaces: the Social Navigation Approach*, Springer-Verlag, London, 2003, 43-82.

89. Kraut, R., Fish, R., Root, R., & Chalfonte, B. Informal Communication in Organizations: Form, Function, and Technology, in R. Baecker (Eds.), *Readings in Groupware and Computer Supported Cooperative Work*, Morgan Kaufmann, 1993.

90. Lazonder, A.W., Biemans, H.J.A. and Worpeis, I.G.J.H. Differences between novice and experienced users in searching information on the World Wide Web. *Journal of the American Society for Information Science and Technology (JASIST)*, *51*, 6, 2000, 576-581.

91. Li, I., Forlizzi, J. & Dey, A. Know thyself: monitoring and reflecting on facets of one's life. . In the *Proceedings of the Extended Abstracts on Human Factors in Computing Systems (CHI EA '10),* 2010, 4489-4492.

92. MacKay, B., Kellar, M., and Watters, C. An evaluation of landmarks for re-finding information on the web. In the *Proceedings of the Extended Abstracts on Human Factors in Computing Systems (CHI EA '05),* 2005, 1609-1612.

93. Malone, T. W., and Lepper, M. R. Making learning fun: A taxonomy of intrinsic motivations for learning. *Aptitude Learning and Instruction,* 3, 1987, 223-253.

94. Marchionini, G. Exploratory search: from finding to understanding. *Communications of the ACM*, 49, 4, April 2006, 41-46.

95. Marchionini, G. *Information Seeking in Electronic Environments*. Cambridge University Press, 1995.

96. McCalla, G. The ecological approach to the design of E-learning environments: Purpose-based capture and use of information about learners. *Journal of Interactive Media Education*, 7, 2004.

97. Meiss, M.R., Menczer, F., Fortunato, S., Flammini, A., & Vespignani, A. Ranking web sites with real user traffic. In the *Proceedings of the International Conference on Web Search and Data Mining (WSDM '08),* 2008, 1067-75.

98. Mertens, R., Farzan, R., and Brusilovsky, P. 2006. Social navigation in web lectures. In the *Proceedings of the Conference on Hypertext and Hypermedia (Hypertext '06)*, 2006, 41-44.

99. Milic-Frayling, N., Jones, R., Rodden, K., Smyth, G., Blackwell, A., and Sommerer, R. Smartback: supporting users in back navigation. In *Proceedings of the International Conference on World Wide Web (WWW '04)*. 2004, 63-71.

100. Millen, D., Feinberg, J., and Kerr, B. Dogear: Social bookmarking in the enterprise. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '06)*, 2006, 111-120.

101. Moraveji, N., Morris, M.R., Morris, D., Czerwinski, M. and Riche, N. ClassSearch: Facilitating the development of web search skills through social learning. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '06)*, 2006, 1797-1806.

102. Moraveji, N., Russell, D., Bien, J. & Mease, D. Measuring improvement in user search performance resulting from optimal search tips. In the *Proceedings of International Conference on Research and Development in Information Retrieval (SIGIR '11)*, 2011, 355-363.

103. Morris, D., Morris, M.R., and Venolia, G. SearchBar: a search-centric web history for task resumption and information re-finding. In *Proceeding of the Conference on Human Factors in Computing Systems* (*CHI '08)*, 2008, 1207-1216.

104. Morris, M. R., and Teevan, J. *Collaborative Web Search: Who, What, Where, When, and Why*. Morgan and Claypool Publishers, San Rafael, CA, USA, 2010.

105. Morris, M.R. A survey of collaborative web search practices. *In Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, 2008, 1657-1660.

106. Morris, M.R. and Horvitz, E. S3: Storable, Shareable Search. In the *Proceedings of Interact*, 2007, 120–123.

107. Morris, M.R., and Horvitz, E. SearchTogether: an interface for collaborative web search. In the *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '07)*, 2007, 3-12.

108. Morrison, J.B., Pirolli, P., and Card, S.K. A taxonomic analysis of what World Wide Web activities significantly impact people's decisions and actions. In the *Proceedings of the Extended Abstracts on Human Factors in Computing Systems (CHI EA '01)*, 2001, 163–164.

109. Nicol, D., Smeaton, C., and Slater, A. F. Footsteps: Trailblazing the web. Computer Networks and ISDN Systems, 27 (1995), 879–885.

110. Nielsen, J. Incompetent research skills curb users' problem solving. *Alertbox.* April 11, 2011. Available at: http://www.useit.com/alertbox/search-skills.html

111. Noelle, D.C. Learning from Advice. In the *Encyclopedia of Cognitive Science.* Wiley Online, 2006.

112. Norman, D.A. Some observations on mental models. In D. Gentner & A. Stevens (Eds.), *Mental Models,* pp. 7-14. Hillsdale, NJ: Erlbaum.

113. Norman, D.A. *The Design of Everyday Things*. Basic Books, 1989.

114. Oard, D. W., and Kim, J. Modelling information content using observable behaviour. In the *Proceedings of the Annual Meeting of the American Society for Information Science and Technology*, USA, 2001, 38-45.

115. Olson, J., Grudin, J., and Horvitz, E. Toward Understanding Preferences for Sharing and Privacy. *Microsoft Technical Report (MSR-TR-2004-138)*, 2004, 10 pages. Available at: http://research.microsoft.com/pubs/70123/tr-2004-138.pdf

116. Olston, C. and Chi, E.H. ScentTrails: Integrating browsing and searching on the Web. *ACM Transactions on Computer-Human Interaction (TOCHI),* 10, 3, September 2003, 177-197.

117. Page, L., Brin, S., Motwani, R., and Winograd, T. The PageRank Citation Ranking: Bringing order to the Web. *Stanford University Technical Report*. 1998. Available at: http://ilpubs.stanford.edu:8090/422

118. Palen, L., and Dourish, P. Unpacking "Privacy" for a Networked World, In the *Proceedings of the ACM Conference on Human-Factors in Computing Systems (CHI '03)*, 2003, 129-136.

119. Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P. and Back, M. Algorithmic Mediation for Collaborative Exploratory Search. In the *Proceedings of International Conference on Research and Development in Information Retrieval (SIGIR '08)*, 2008, 315-322.

120. Pinelle, D., and Gutwin, C. Designing for loose coupling in mobile groups. In *Proceedings of the International Conference on Supporting Group Work (GROUP '03)*, 2003, 75-84.

121. Pirolli, P., and Card, S. Information foraging in information access environments. In the *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '95)*, 1995, 51–58.

122. Pirolli, P., and Card, S. Information foraging. *Psychology Review*, 106, 4, 1999, 643–675.

123. Pirolli, P., and Card, S. K. The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis. In the *Proceedings of the International Conference on Intelligence Analysis*, 2005. 6 pages. Available at: http://vadl.cc.gatech.edu/documents/2__card-sensemaking.pdf

124. Power, C., McQuillan, I., Petrie, H., Kennaugh, P., Daley, M., and Wozniak, G. No Going Back: An Interactive Visualization Application for Trailblazing on the Web. In the *Proceedings of the International Conference Information Visualisation*, 2008, 133-142.

125. Purcell, K., Brenner, J., and Rainie, L. *Search Engine Use 2012*. Technical Report, Pew Internet and American Life Project. March 2012. Available at http://www.pewinternet.org/Reports/2012/Search-Engine-Use-2012.aspx

126. Resnick, P., Neophytos, I., Suchak, M., Bergstrom, P., and Riedl, J. GroupLens: an Open Architecture for Collaborative Filtering of Netnews. In the Proceedings of the *ACM Conference on Computer-Supported Cooperative Work (CSCW '94)*, 1994, 175-186.

127. Rich, E. User Modelling via Stereotypes. *Cognitive Science*, 3, 1979, 335-366.

128. Rozin, P., Royzman, E.B. Negativity Bias, Negativity Dominance, and Contagion. *Personality and Social Psychology Review,* 5, 296–320.

129. Russell, D. M., Stefik, M. J., Pirolli, P. L., and Card, S. K. The Cost Structure of Sensemaking. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '93)*, 1993, 269–276.

130. Saito, H. and Miwa, K. A cognitive study of information seeking processes in the WWW: Effects of searcher's knowledge and experience. In the *Proceedings of the International Conference on Web Information Systems Engineering (WISE'01)*, 2001, 321-333.

131. Schein, A.I., Popescul, A., Ungar, L.H., and Pennock, D.M. Methods and metrics for cold-start recommendations. In *Proceedings of the Conference on Research and Development in Information Retrieval (SIGIR '02)*, 2002, 253-260.

132. Schwarz, J. and Morris, M. Augmenting web pages and search results to support credibility assessment. In the *Proceedings of the Annual Conference on Human Factors in Computing Systems (CHI '11),* 2011. 1245-1254.

133. Sellen, A.J., Murphy, R., & Shaw, K.L. How knowledge workers use the Web. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '02)*, 2002, 227–234.

134. Shah, C. and Marchionini, G. Awareness in collaborative information seeking. *Journal of the American Society for Information Science and Technology (JASIST)*, 61, 2010, 1970–1986.

135. Shneiderman, B. The eyes have it: A task by data type taxonomy for information visualizations. In the *Proceedings of IEEE Symposium on Visual Languages*, 1996, 336–343.

136. Smeaton, A.F., Lee, H., Foley, C., and McGivney, S. Collaborative Video Searching on a Tabletop. *Multimedia Systems Journal*, 1, 4, 2006, 375-391.

137. Smyth, B. A Community-Based Approach to Personalizing Web Search. *IEEE Computer*, 40, 8, 2007, 42-50.

138. Spärck-Jones, K. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, *28,* 1, 1972, 11-21.

139. Spence, R. A framework for navigation, *International Journal of Human-Computer Studies (IJHCS),* 51, 1999, 919–945.

140. Spence, R. *Information Visualization: Design for Interaction (Second Edition)*. Pearson Education Limited, 2007.

141. Svensson, M., Höök, K., & Cöster, R. Designing and evaluating kalas: A social navigation system for food recipes. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 12, 3, 2005, 374-400.

142. Svensson, M., Höök, K., Laaksolahti, J., & Waern, A. Social Navigation of Food Recipes, In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '01)*, 2001, 341-348.

143. Svensson, M., Laaksolahti, J., Höök, K., Waern, A.: A recipe based on-line foodstore. In the *Proceedings of the International Conference on Intelligent User Interfaces (IUI '00)*, 2000, 260-263.

144. Takano, H. and Winograd, T. Dynamic bookmarks for the WWW. In the *Proceedings of the Conference on Hypertext and Hypermedia (Hypertext '98)*, 1998, 297-298.

145. Tam, J., and Greenberg, S. A framework for asynchronous change awareness in collaborative documents and workspaces. *International Journal of Human-Computer Studies (IJHCS),* 64, 7, July 2006, 583-598.

146. Teevan, J., Alvardo, C. Ackerman, M., and Karger, D. The Perfect Search Engine is not Enough: A Study of Orienteering Behaviour in Directed Search. In the *Proceedings of the Conference on Human Factors in Computing Systems (CHI '01)*, 415-422.

147. Teevan, J., Dumais, S.T., and Horvitz, E. Personalizing Search via Automated Analysis of Interests and Activities. In *Proceedings of the 28th Annual ACM Conference on Research and Development in Information Retrieval (SIGIR '05),* 2005, 449-456.

148. Teevan, J., Dumais, S.T., and Liebling, D.J. A Longitudinal Study of How Highlighting Web Content Change Affects People's Web Interactions. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '10)*, April 2010, 1353-1356.

149. Thaler, R.H. & Sunstein, C.R. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. Yale, 2008.

150. Tversky, B. Cognitive Maps, Cognitive Collages and Spatial Mental Models. In *Proceedings of European Conference on Spatial Theory – A Theoretical Basis for GIS (COSIT '93)*, 1993, 14-24.

151. Van Kleek, M., Moore, B., Xu, C., & Karger, D.R. Eyebrowse: Real-Time Web Activity Sharing and Visualization. In the *Proceedings of the Extended Abstracts on Human Factors in Computing Systems (CHI EA ' 10)*, 2010, 3643-3648.

152. Van Kleek, M., Smith, D.A., Shadbolt, N. and schraefel, m.c. A decentralized architecture for consolidating personal information ecosystems: The WebBox. *In the Proceedings of the International Workshop on Personal Information Management (PIM 2012)*, 2012, 4 pages. Available at http://pimworkshop.org/2012/pdf/kleek_2012_decentralized.pdf

153. Wang, Y., Aroyo, L., Stash, N., and Rutledge, L. Interactive User Modelling for Personalized Access to Museum Collections: The Rijksmuseum Case Study. In Proceedings of the *International Conference on User Modelling (UM '07)*, 2007, 385-389.

154. Weiser, M. Some Computer Science Issues in Ubiquitous Computing. *Communications of the ACM (CACM)*, 36, 7, July 1993, 74-83.

155. Weiser, M. The Computer of the 21st Century. *Scientific American*, 265, 3, September, 1991, 94-104.

156. Wexelblat, A. and Maes, P. Footprints: history-rich tools for information foraging. In Proceedings of *the Conference on Human Factors in Computing System (CHI '99),* 1999, 270-277.

157. Wexelblat, A. Communities through Time: Using History for Social Navigation. In *Community Computing and Support Systems, Social Interaction in Networked Communities*, 1519, 1998, 281-298.

158. White, R.W. and Morris, D. Investigating the querying and browsing behaviour of advanced search engine users. In the *Proceedings of the Conference on Research and Development in Information Retrieval (SIGIR '07)*, 2007, 255–262.

159. White, R.W. and Roth, R. *Exploratory Search: Beyond the Query-Response Paradigm.* San Rafael, CA, Morgan and Claypool. 2009.

160. White, R.W., Bilenko, M., and Cucerzan, S. Studying the use of popular destinations to enhance web search Interaction. In the *Proceedings of the Conference on Research and Development in Information Retrieval (SIGIR '07)*, 2007, 159–166.

161. White, R.W., Dumais, S., and Teevan, J. Characterizing the influence of domain expertise on web search behaviour. In the *Proceedings of the International Conference on Web Search and Data Mining (WSDM '09)*, 2009, 132–141.

162. White, R.W., Ruthven, I., Jose, J.M. A study of factors affecting the utility of implicit relevance feedback. In the *Proceedings of the Conference on Research and Development in Information Retrieval (SIGIR '05)*, 2005, 35-42.

163. Whittaker, S., Frohlich, D., and Daly-Jones, O. Informal Workplace Communication: What Is It Like and How Might We Support It? In the *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '94)*, 1994, 131-137.

164. Wikipedia contributors. Computer supported cooperative work. *Wikipedia, The Free Encyclopedia,* March 7th, 2011, 08:48 UTC. Accessed March 11th, 2011. Available from: http://en.wikipedia.org/w/index.php?title=Computer_supported_cooperative_work&oldid=417574299.

165. Wikipedia contributors. Learning.. In *Wikipedia, The Free Encyclopedia*, November 30th, 2011, UTC 17:47. Accessed December 2nd, 2011. Available at: http://en.wikipedia.org/w/index.php?title=Learning&oldid=463312340.

166. Wileman, R. E. *Visual communicating*. Englewood Cliffs, N.J.: Educational Technology Publications, 1993.

167. Willett, W., Heer, J., and Agrawala, M. Scented Widgets: Improving Navigation Cues with Embedded Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13, 6, November 2007, 1129-1136.

168. Wilson, P. *Computer Supported Cooperative Work: An Introduction*. Kluwer Academic Pub, 1991.

169. Wobbrock, J.O., Findlater, L., Gergle, D. and Higgins, J.J. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only ANOVA Procedures. In the *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '11),* 2011, 143-146.

170. Wright, A. *Glut: Mastering information through the ages*. Washington DC, Joseph Henry Press, 2007.

171. Zipf, G. K. *Human Behavior and the Principle of Least Effort*, Addison-Wesley, 1949.

APPENDIX A: STUDY MATERIALS

University Ethical Approval for the Dissertation Research

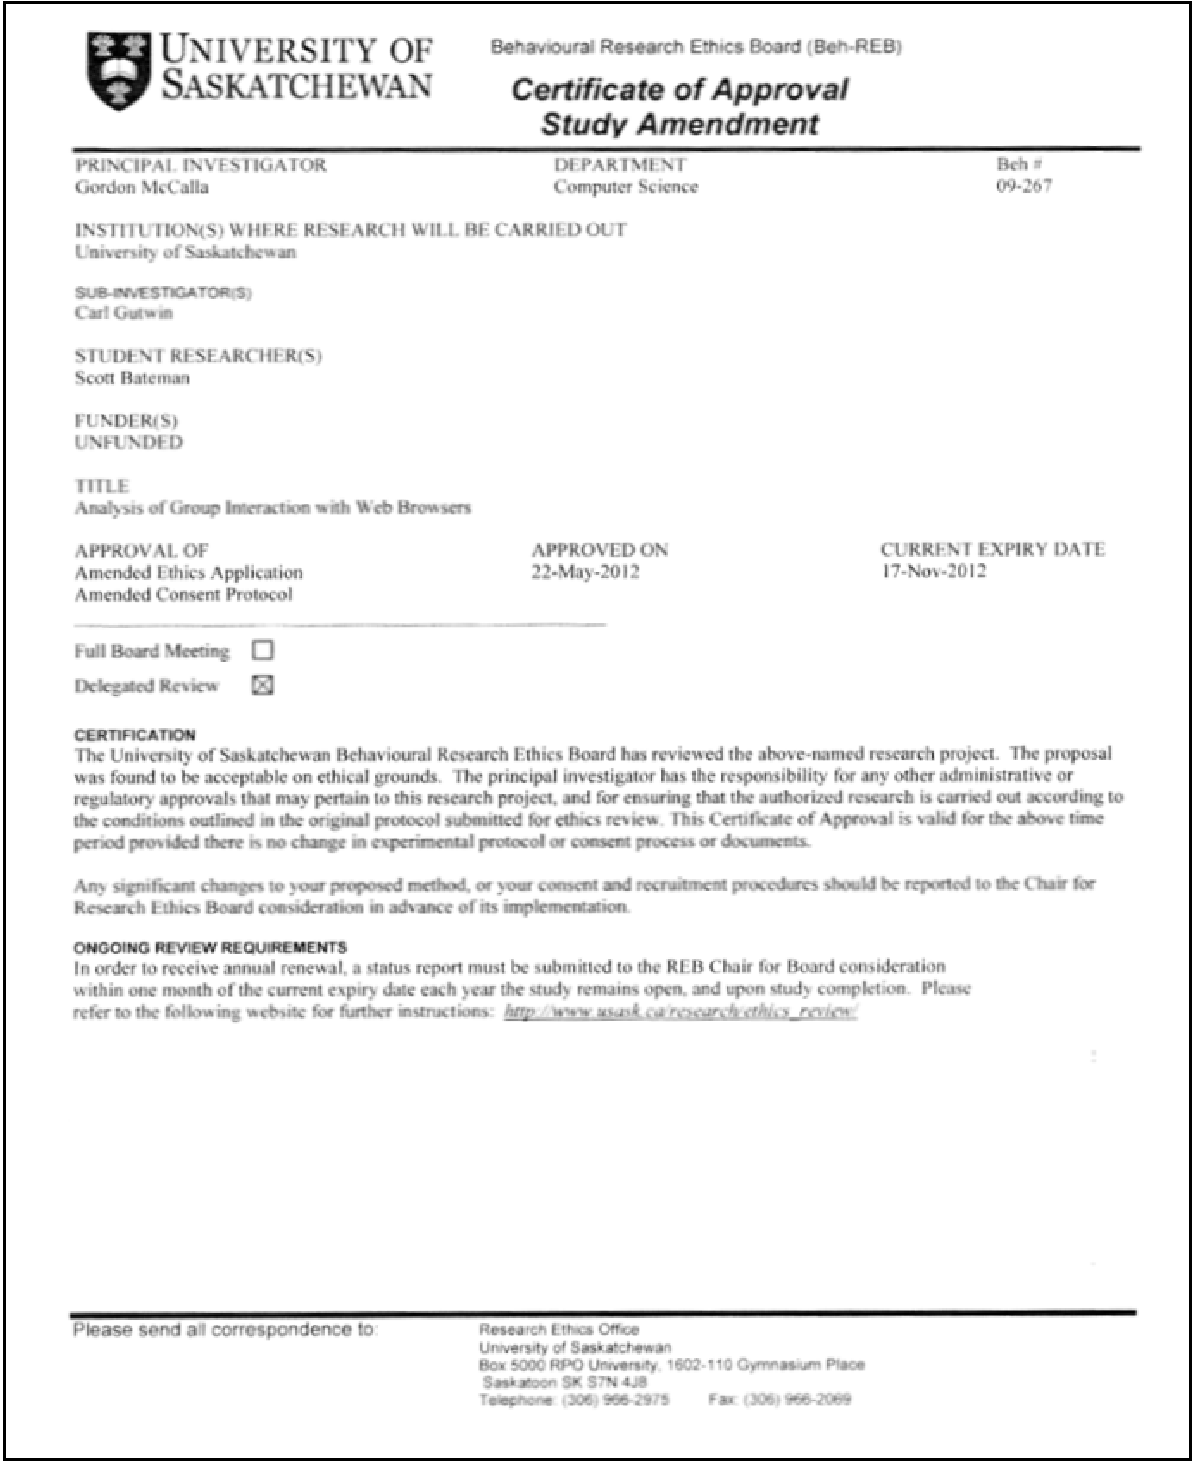

UNIVERSITY OF SASKATCHEWAN

Behavioural Research Ethics Board (Beh-REB)

Certificate of Approval
Study Amendment

PRINCIPAL INVESTIGATOR	DEPARTMENT	Beh #
Gordon McCalla	Computer Science	09-267

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT
University of Saskatchewan

SUB-INVESTIGATOR(S)
Carl Gutwin

STUDENT RESEARCHER(S)
Scott Bateman

FUNDER(S)
UNFUNDED

TITLE
Analysis of Group Interaction with Web Browsers

APPROVAL OF	APPROVED ON	CURRENT EXPIRY DATE
Amended Ethics Application Amended Consent Protocol	22-May-2012	17-Nov-2012

Full Board Meeting ☐

Delegated Review ☒

CERTIFICATION
The University of Saskatchewan Behavioural Research Ethics Board has reviewed the above-named research project. The proposal was found to be acceptable on ethical grounds. The principal investigator has the responsibility for any other administrative or regulatory approvals that may pertain to this research project, and for ensuring that the authorized research is carried out according to the conditions outlined in the original protocol submitted for ethics review. This Certificate of Approval is valid for the above time period provided there is no change in experimental protocol or consent process or documents.

Any significant changes to your proposed method, or your consent and recruitment procedures should be reported to the Chair for Research Ethics Board consideration in advance of its implementation.

ONGOING REVIEW REQUIREMENTS
In order to receive annual renewal, a status report must be submitted to the REB Chair for Board consideration within one month of the current expiry date each year the study remains open, and upon study completion. Please refer to the following website for further instructions: http://www.usask.ca/research/ethics_review/

Please send all correspondence to:

Research Ethics Office
University of Saskatchewan
Box 5000 RPO University, 1602-110 Gymnasium Place
Saskatoon SK S7N 4J8
Telephone: (306) 966-2975 Fax: (306) 966-2069

WebWear – Ethical Consent Form

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
INFORMED CONSENT FORM

Research Project: **WebWear – new browser functionality**
(Ethics Approval #: BEH 09-267)

Investigators: Dr. Carl Gutwin, Department of Computer Science (966-4888)
Dr. Gordon McCalla, Department of Computer Science (966-4888)
Scott Bateman, Department of Computer Science

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.

This study is concerned with detecting preferences and performance benefits from sharing contextual web histories between contacts.

The goal of the research is to evaluate an interface of contextual web histories within a web browser.

The session will require no more than 90 minutes, during which you will be asked complete a number of hypothetical web search tasks using three different interfaces.

At the end of the session, you will be given more information about the purpose and goals of the study, and there will be time for you to ask questions about the research.

The data collected from this study will be used in articles for publication in journals and conference proceedings.

As one way of thanking you for your time, we will be pleased to make available to you a summary of the results of this study once they have been compiled (usually within two months). This summary will outline the research and discuss our findings and recommendations. If you would like to receive a copy of this summary, please write down your email address here.

Contact email address:____

All personal and identifying data will be kept confidential. If explicit consent has been given, textual excerpts, photographs, or video recordings may be used in the dissemination of research results in scholarly journals or at scholarly conferences. Anonymity will be preserved by using pseudonyms in any presentation of textual data in journals or at conferences. The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. Do you have any questions about this aspect of the study?

You are free to withdraw from the study at any time without penalty and without losing any advertised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. If you withdraw, your data will be deleted from the study and destroyed.

Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-4888, gutwin@cs.usask.ca

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact:

- Dr. Carl Gutwin, Professor, Dept. of Computer Science, (306) 966-4888, gutwin@cs.usask.ca
- Office of Research Services, University of Saskatchewan, (306) 966-4053

Participant's signature:____

Date:_________________

Investigator's signature:____

Date:_________________

A copy of this consent form has been given to you to keep for your records and reference. This research has the ethical approval of the Office of Research Services at the University of Saskatchewan.

WebWear – Textual Excerpt Consent Form

UNIVERSITY OF
SASKATCHEWAN

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
TRANSCRIPT / TEXTUAL EXCERPT CONSENT FORM

Research Project:	**WebWear – new browser functionality (Ethics Approval #: BEH 09-267)**
Investigators:	Dr. Carl Gutwin, Department of Computer Science (966-4888)
	Dr. Gordon McCalla, Department of Computer Science (966-4888)
	Scott Bateman, Department of Computer Science

TRANSCRIPTS

"I, ___________________________________, agree to allow transcripts of my conversations with the investigator or other participants during the experiment to be used for public presentation of the research results in the manner described in the consent form. However, I understand that I will be given the opportunity to read any transcript excerpts that are intended for public participation and to withdraw consent for them to be reported, if so desired. I also understand that I will receive a copy of any transcripts presented publically for my records. I understand that all identifying information will be removed from the transcripts and names will be changed prior to publication. "

Participant		**Investigator**	
Name:	________________________________	Name:	________________________________
Signature:	________________________________	Signature:	________________________________
Date:	________________________________	Date:	________________________________

TEXTUAL EXCERPTS

"I, ___________________________________, agree to allow excerpts of text that I wrote to be used for public presentation of the research results in the manner described in the consent form. However, I understand that I will be given the opportunity to read any excerpts that are intended for public participation and to withdraw consent for them to be reported, if so desired. I also understand that I will receive a copy of any textual excerpts presented publically for my records. I understand that all identifying information will be removed from the excerpts and names will be changed prior to publication."

Participant		**Investigator**	
Name:	________________________________	Name:	________________________________
Signature:	________________________________	Signature:	________________________________
Date:	________________________________	Date:	________________________________

WebWear – Study Tasks

Division Task 1:
Your friend from work, Mary, has a history of breast cancer in her family. Chatting about it over lunch last week, you decided to research breast cancer symptoms together. You are trying to get the widest coverage of different information sources, so you are both trying to collect the URLs (webpage addresses) for 3 introductory articles that mention common breast cancer symptoms that you can look at together over lunch next week. Of course, it would be best to find different articles than Mary. Copy-and-paste the URLs for each of the webpages you find below.

Division Task 2:
You, Mary, and John are planning a retreat for executives at your company. Typically the retreats (really business meetings held over a few days) have been held in scenic locations. Since the reservations need to be made quickly, the three of you have decided to narrow your search down to locations in the Florida keys (where it will be nice at the time of the meetings in November). The executives are typically happy with well-known hotel chains (like the Ramada or Marriott) with meeting facilities. They are less concerned about price, but want a really nice setting. You have all decided to identify a candidate location, and look at the websites together tomorrow. Of course, it would be best to find a hotel that the others are not considering, so that you have more options to consider. Copy-and-paste the URL for an appropriate hotel into the field below.

Division Task 3:
Lynn called you earlier tonight in a panic. The office's cleaning service has gone out of business, and she needs to get a new one in place tomorrow morning. On top of this her laptop has died, right when she was in the middle of collecting some contact information for other cleaning services in Saskatoon. She needs you to get a couple (two) more phone numbers for cleaning services in Saskatoon. Find two phone numbers that Lynn probably doesn't already have, and copy-and-paste the URLs where you found them in the space below (i.e., don't take them from the Google results page).

Follow Task 1:
John knows you are interested in getting a new smart phone, he tells you that he saw an article that was linked to on the website Slashdot.org a few weeks ago that suggests that the Apple iPhone may not be the greatest, and that you should think about a Windows phone because some people think it is more beautiful; and, one of those people was a co-founder of Apple. You decide that you would like to find the article, but John can't remember the name of it. Copy-and-paste the URL for the article into the field below.

Follow Task 2:
Your boss Lynn always makes great appetizers for office get-togethers. She said that she always gets her appetizer recipes from a particular website, but you can't remember which one it is. It's late on Thursday and you would like to make an appetizer for some visitors coming over on Friday evening after work. You particularly liked the 'Pea Hummus' that Lynn made a few days ago. You decide to try and find the website that Lynn uses. Find the URL for a page on the website that Lynn always uses, and copy-and-paste it into the field below.

Follow Task 3:
John tells you that he saw a great book with bird photographs last week on Amazon.com that the office should get together and buy for Mary for her birthday (John says "she is mildly obsessed with birds"). Mary has a lot of books about birds, but John says he was able to figure out that she doesn't have any coffee table style books like this one, and it only costs about $38. You are in charge of buying the present, so you decide to try and find the book John had seen. Copy-and-paste the URL of the Amazon.com page into the space below.

WebWear – Study Handouts (Used to explain the task scenarios and study systems)

Task Scenarios

The Setting

You like to get a few tasks of your list before bed each night. It’s Thursday and you were busy tonight, so it’s quite late and you have just sat down to your task list.

Because it’s late and you have a busy Friday ahead of you tomorrow, you would like to be as efficient as possible in getting your tasks done. While you want to get the tasks done quickly, you also want to make sure you find what you are looking for as best you can.

The People Involved in Your Tasks

John

John is a computer programmer in your work group. He is a really friendly guy who seems to know everything related to technology. John’s work is slow and methodical, you can always count on him to do consider lots of options and do a thorough job.

Mary

Mary is your closest friend at work. You have worked with her for a long time. She is a very efficient worker, and likes to make decisions quickly.

Lynn

Lynn is your manager at work, she is like a “Super Woman”: extremely good at her job, manages a young family, and she is very well read. You look up to Lynn and really appreciate her advice.

The Everyone System

The Everyone System shows how much everyone in the world visits different pages on the Web.

How Many People Visited the Page I am Visiting?

The number people who visited the page you are currently visiting is shown by a bar at the top of the page you are visiting. There are four different levels shown by the bar.

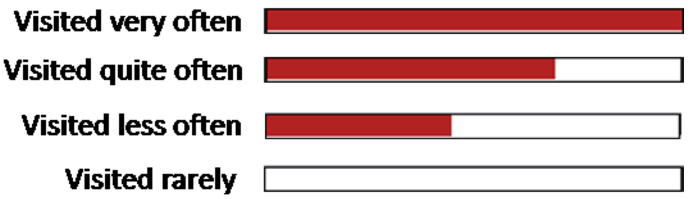

How Many People Visited that Link?

When you see a link that you can click to go to another webpage, you will see bars that give you the same information as above.

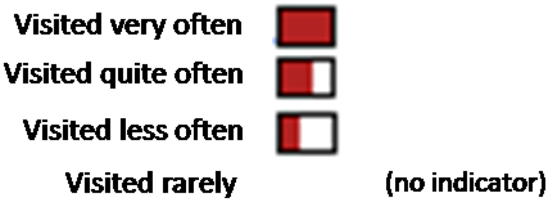

Where are the Link Indicators on the Page?

When a page is long, you might not know where the links with indicators are on the page. When a page is long, a grey bar appears with the location of the link indicators down the page. You can click on an indicator to jump the scrollbar to that part of the page.

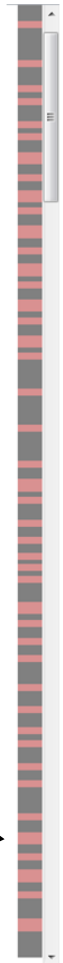

Link indicator,
for link near the bottom of
the page

The Friends System

The Friends System shows you the activities of your ***contacts***, including your friends and colleagues. Your contacts have decided to share certain parts of the web browsing history with you in hopes that it might be useful.

How Many of My Contacts Visited the Page and the Website I am Visiting?

The number people who visited the exact page and the website (where the page is) you are currently visiting is shown by two bars at the top of the page. The page information is shown in the purple color, while the website is shown by green. There are four different levels shown by the bars.

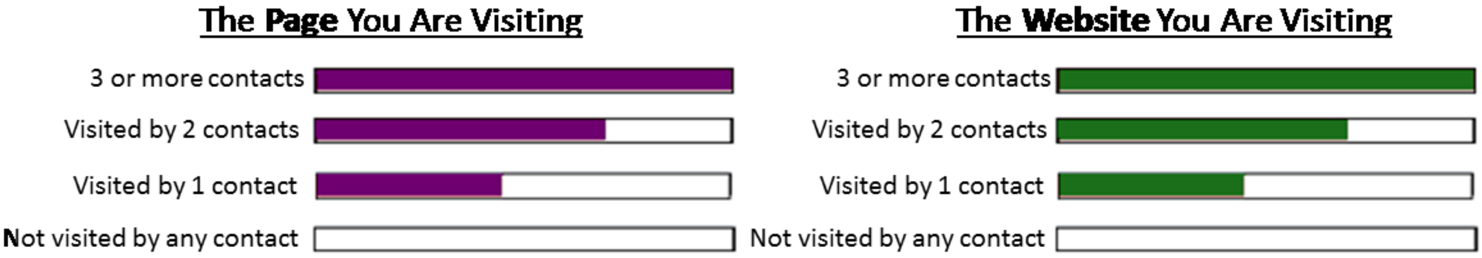

How Many People Visited that Link?

When you see a link that you can click to go to another webpage, you will see bars that give you the same information as above.

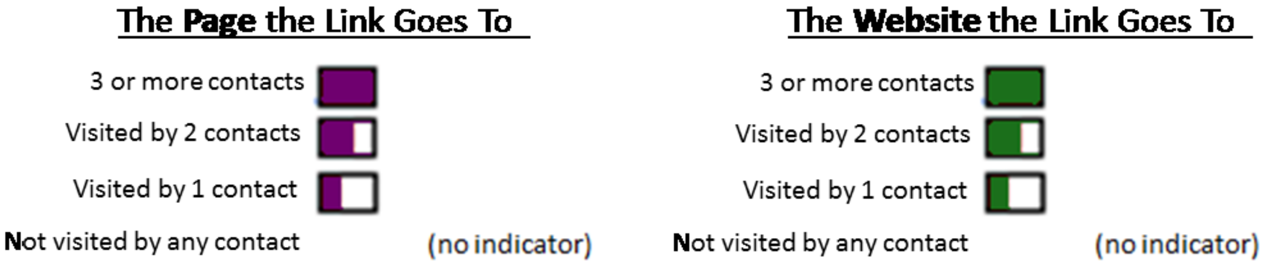

Getting More Information

Because you are seeing information from your contacts (people you know), it might be helpful to see who visited the particular page or website, when they visited, for how long, and what other pages they visited on the Website. To see this information simply put your mouse over any of the little bars on the page (at the top or beside the links).

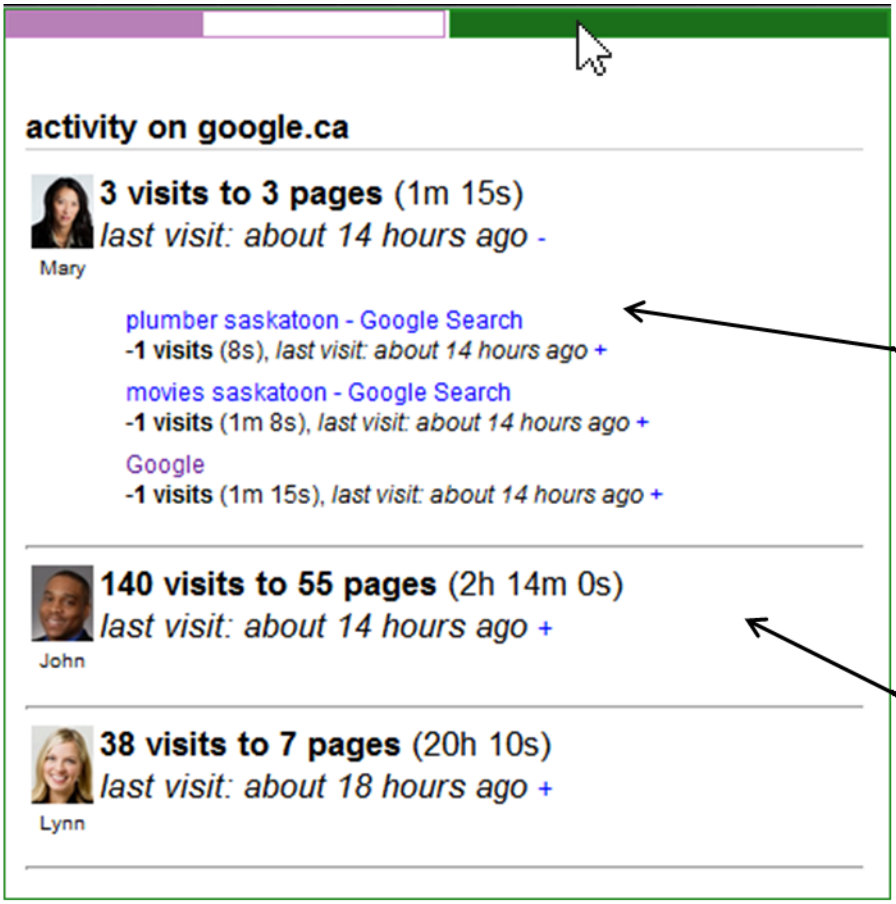

activity on google.ca

3 visits to 3 pages (1m 15s)
last visit: about 14 hours ago -
Mary

plumber saskatoon - Google Search
-**1 visits** (8s), *last visit: about 14 hours ago* +

movies saskatoon - Google Search
-**1 visits** (1m 8s), *last visit: about 14 hours ago* +

Google
-**1 visits** (1m 15s), *last visit: about 14 hours ago* +

140 visits to 55 pages (2h 14m 0s)
last visit: about 14 hours ago +
John

38 visits to 7 pages (20h 10s)
last visit: about 18 hours ago +
Lynn

More details about each page visited,
Summary of each page visited for the website. The most recent pages appear at the top. Clicking on the page title will take you the page.

Summary by contact,
The number of visits, how many different pages, how much time was spent, and when the last visit occurred.

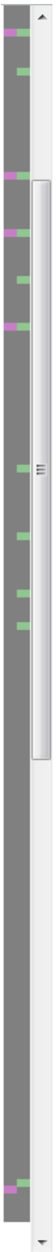

Where are the Link Indicators on the Page?

When a page is long, you might not know where the links with indicators are on the page. When a page is long, a grey bar appears with the location of the link indicators down the page. You can click on an indicator to jump the scrollbar to that part of the page.

Link indicators,
for link near the top of the page
(purple for a page indicator, green
for a website indicator)

What Did my Contacts Search For?

You can also see searches that your contacts have shared with you. When you use a search engine, like Google, these will appear automatically. As you type your search only your contacts searches that match yours will be show. E.g., typing ‘w’ only shows the most recent searches that contain a ‘w’.

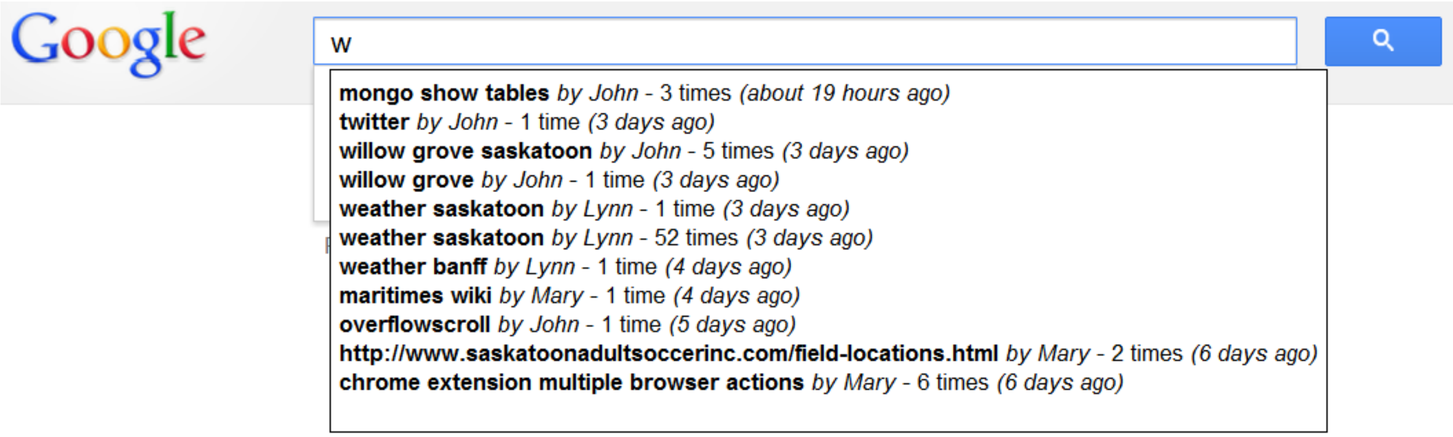

w

mongo show tables *by John* - 3 times *(about 19 hours ago)*
twitter *by John* - 1 time *(3 days ago)*
willow grove saskatoon *by John* - 5 times *(3 days ago)*
willow grove *by John* - 1 time *(3 days ago)*
weather saskatoon *by Lynn* - 1 time *(3 days ago)*
weather saskatoon *by Lynn* - 52 times *(3 days ago)*
weather banff *by Lynn* - 1 time *(4 days ago)*
maritimes wiki *by Mary* - 1 time *(4 days ago)*
overflowscroll *by John* - 1 time *(5 days ago)*
http://www.saskatoonadultsoccerinc.com/field-locations.html *by Mary* - 2 times *(6 days ago)*
chrome extension multiple browser actions *by Mary* - 6 times *(6 days ago)*

WebWear – Post Task Questionnaire

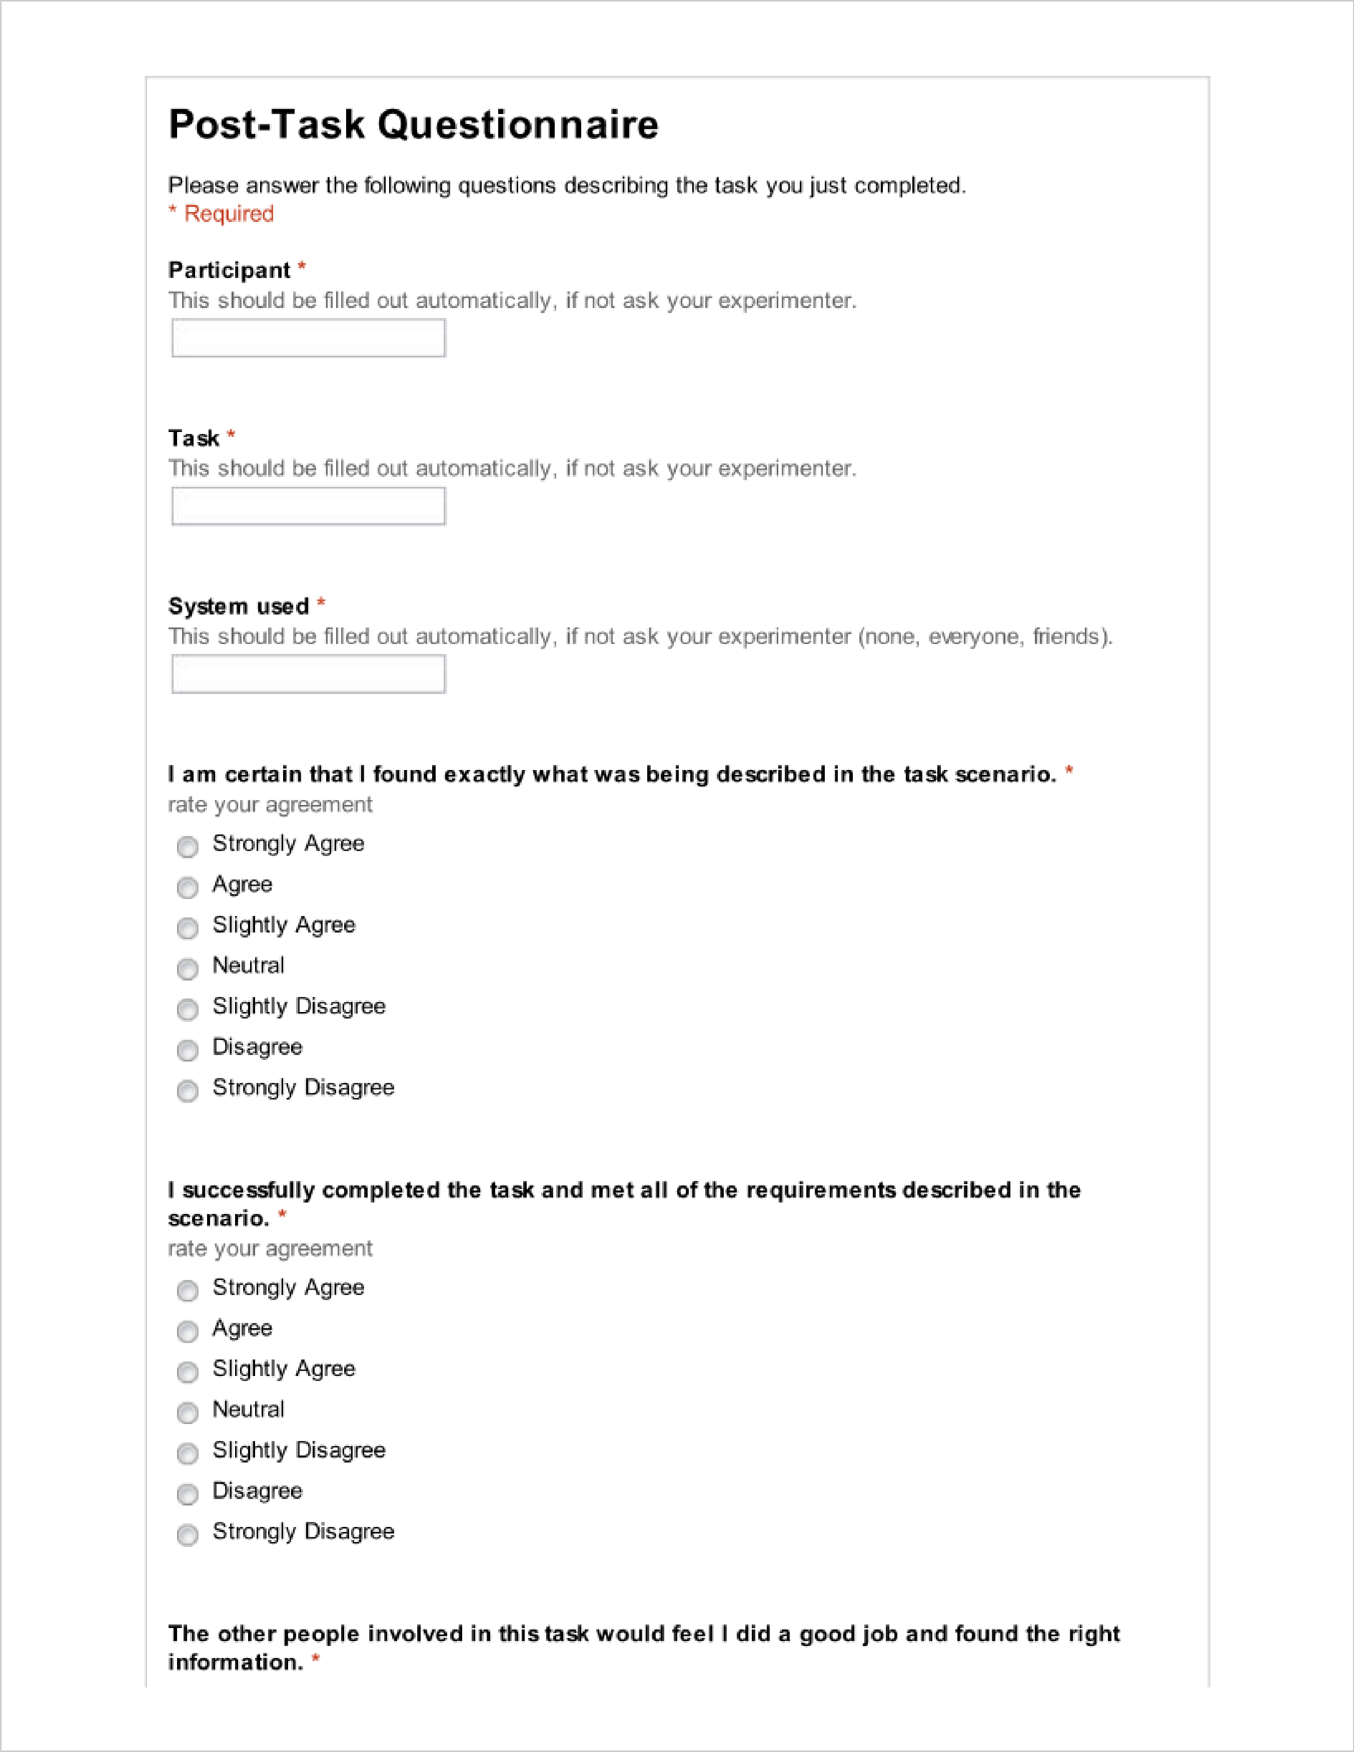

Post-Task Questionnaire

Please answer the following questions describing the task you just completed.
* Required

Participant *
This should be filled out automatically, if not ask your experimenter.

Task *
This should be filled out automatically, if not ask your experimenter.

System used *
This should be filled out automatically, if not ask your experimenter (none, everyone, friends).

I am certain that I found exactly what was being described in the task scenario. *
rate your agreement

- ○ Strongly Agree
- ○ Agree
- ○ Slightly Agree
- ○ Neutral
- ○ Slightly Disagree
- ○ Disagree
- ○ Strongly Disagree

I successfully completed the task and met all of the requirements described in the scenario. *
rate your agreement

- ○ Strongly Agree
- ○ Agree
- ○ Slightly Agree
- ○ Neutral
- ○ Slightly Disagree
- ○ Disagree
- ○ Strongly Disagree

The other people involved in this task would feel I did a good job and found the right information. *

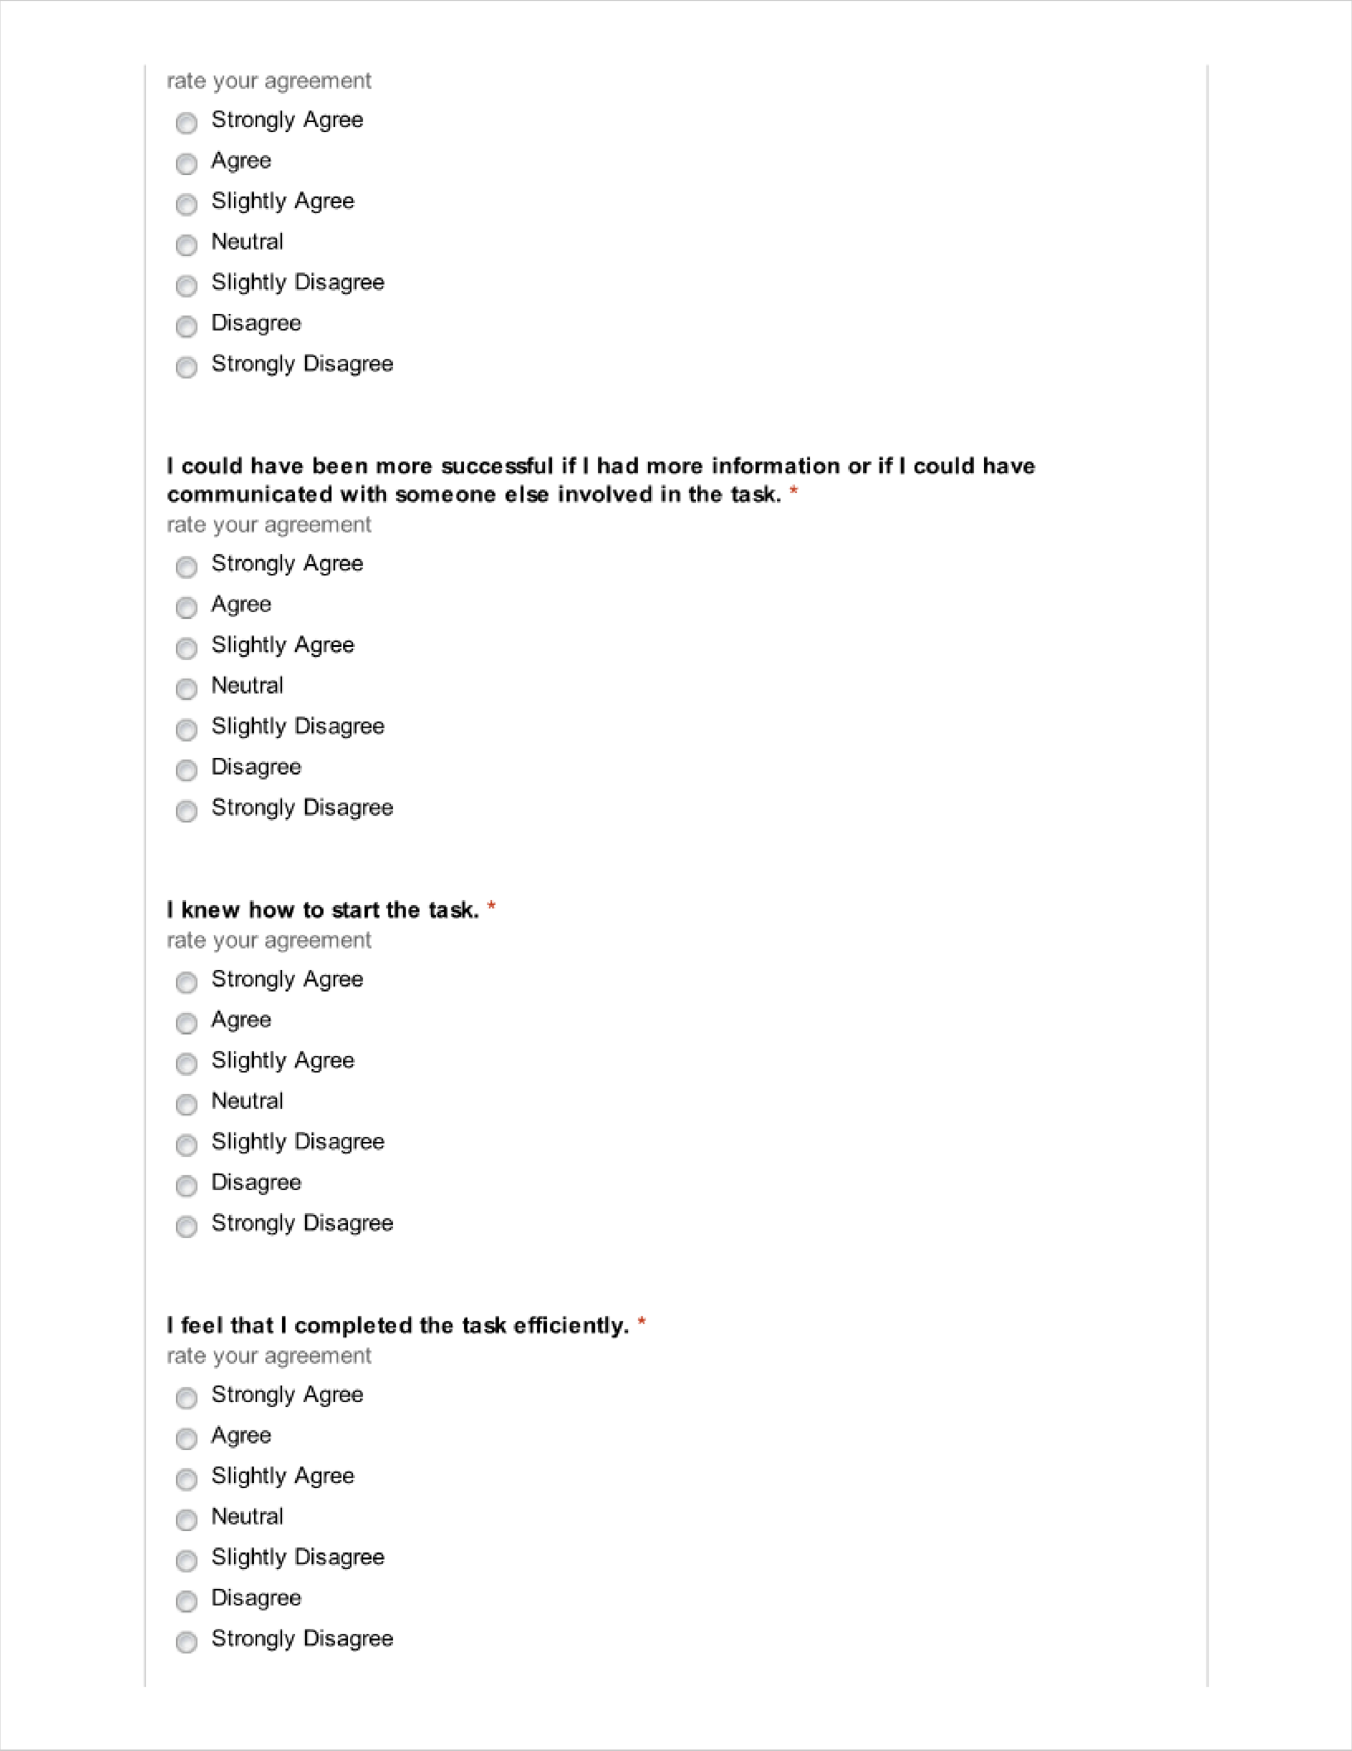

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I could have been more successful if I had more information or if I could have communicated with someone else involved in the task. *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I knew how to start the task. *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I feel that I completed the task efficiently. *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

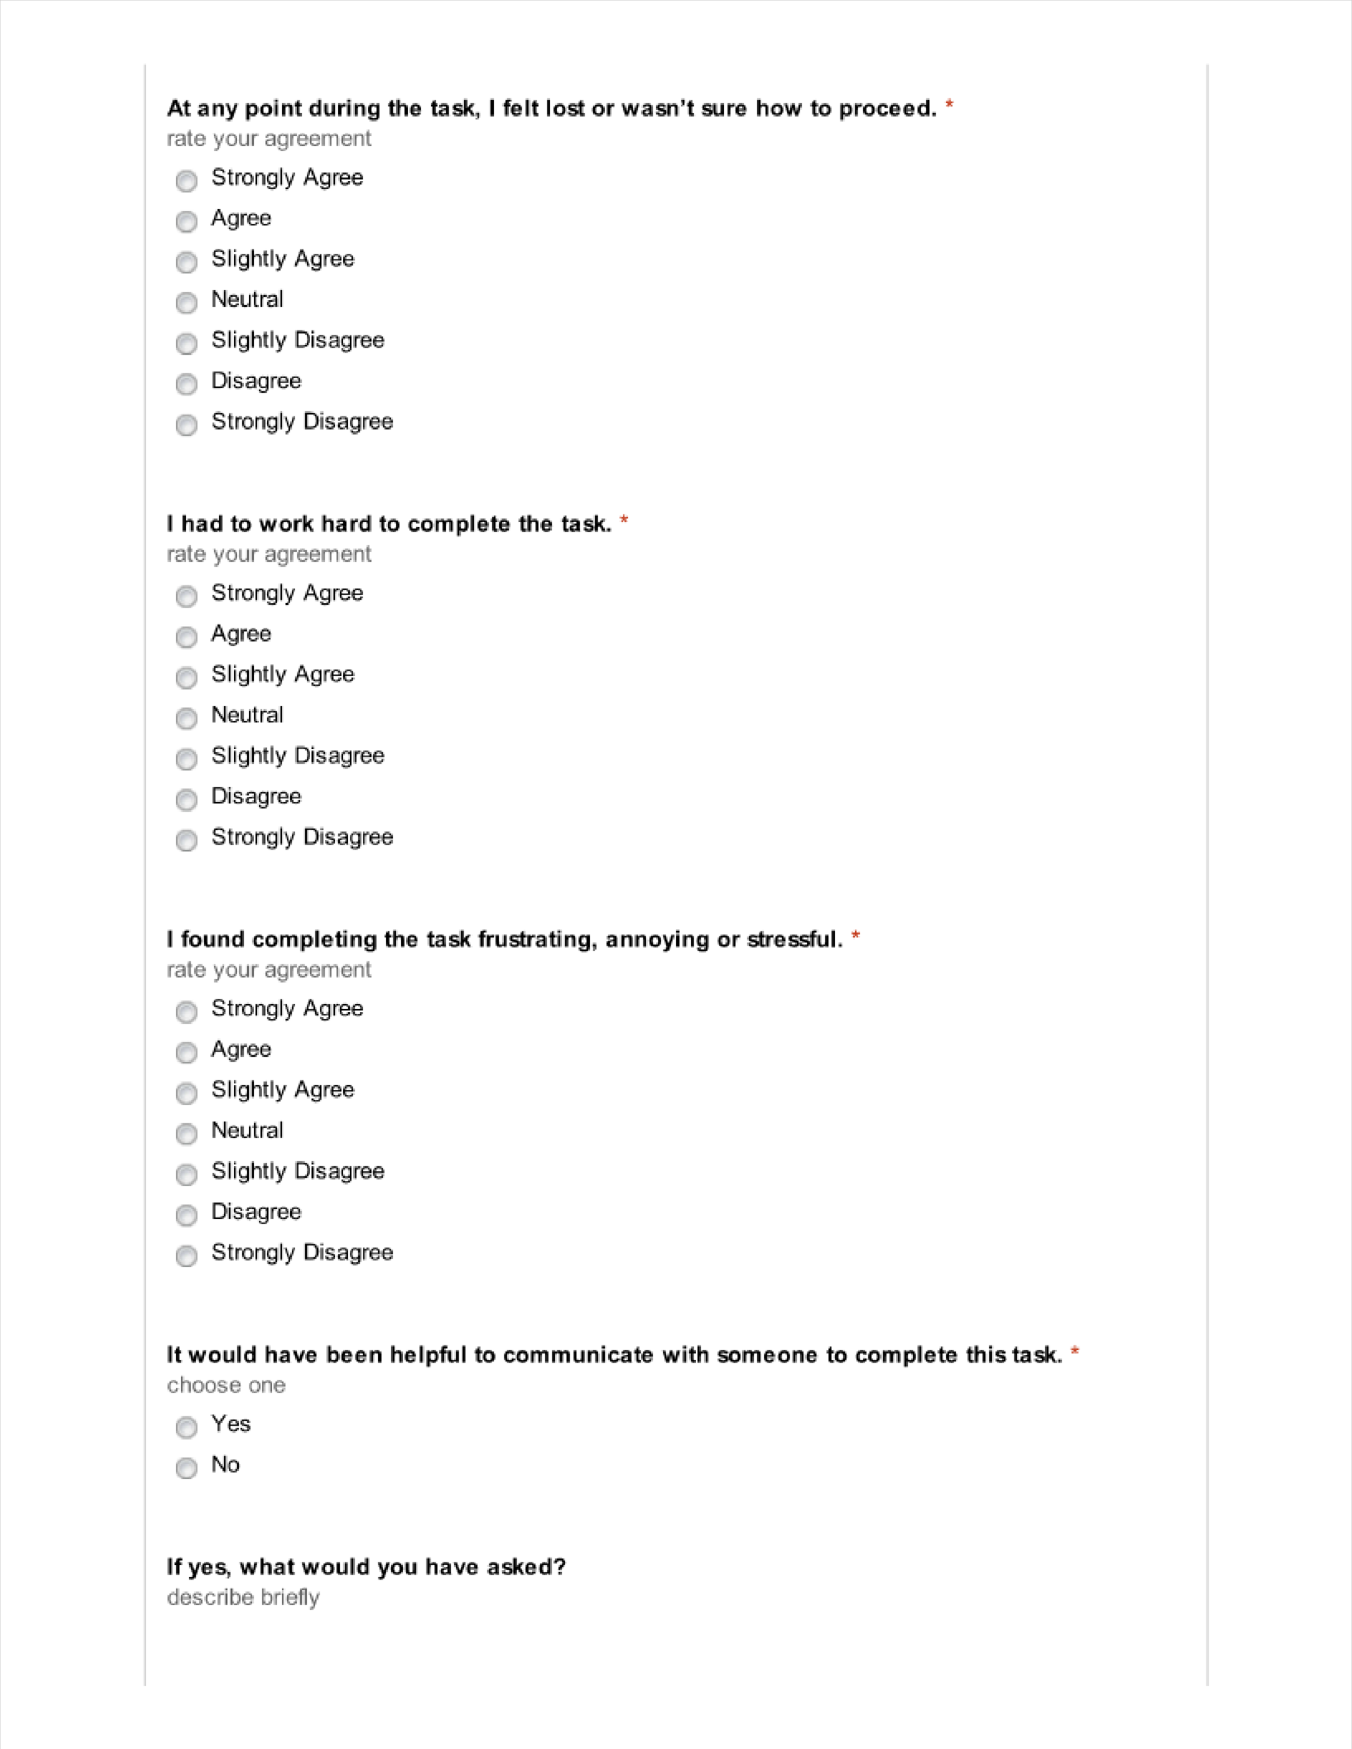

At any point during the task, I felt lost or wasn't sure how to proceed. *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I had to work hard to complete the task. *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I found completing the task frustrating, annoying or stressful. *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

It would have been helpful to communicate with someone to complete this task. *

choose one

○ Yes
○ No

If yes, what would you have asked?

describe briefly

Do you have any comments you would like to share?

WebWear – Post System Questionnaire

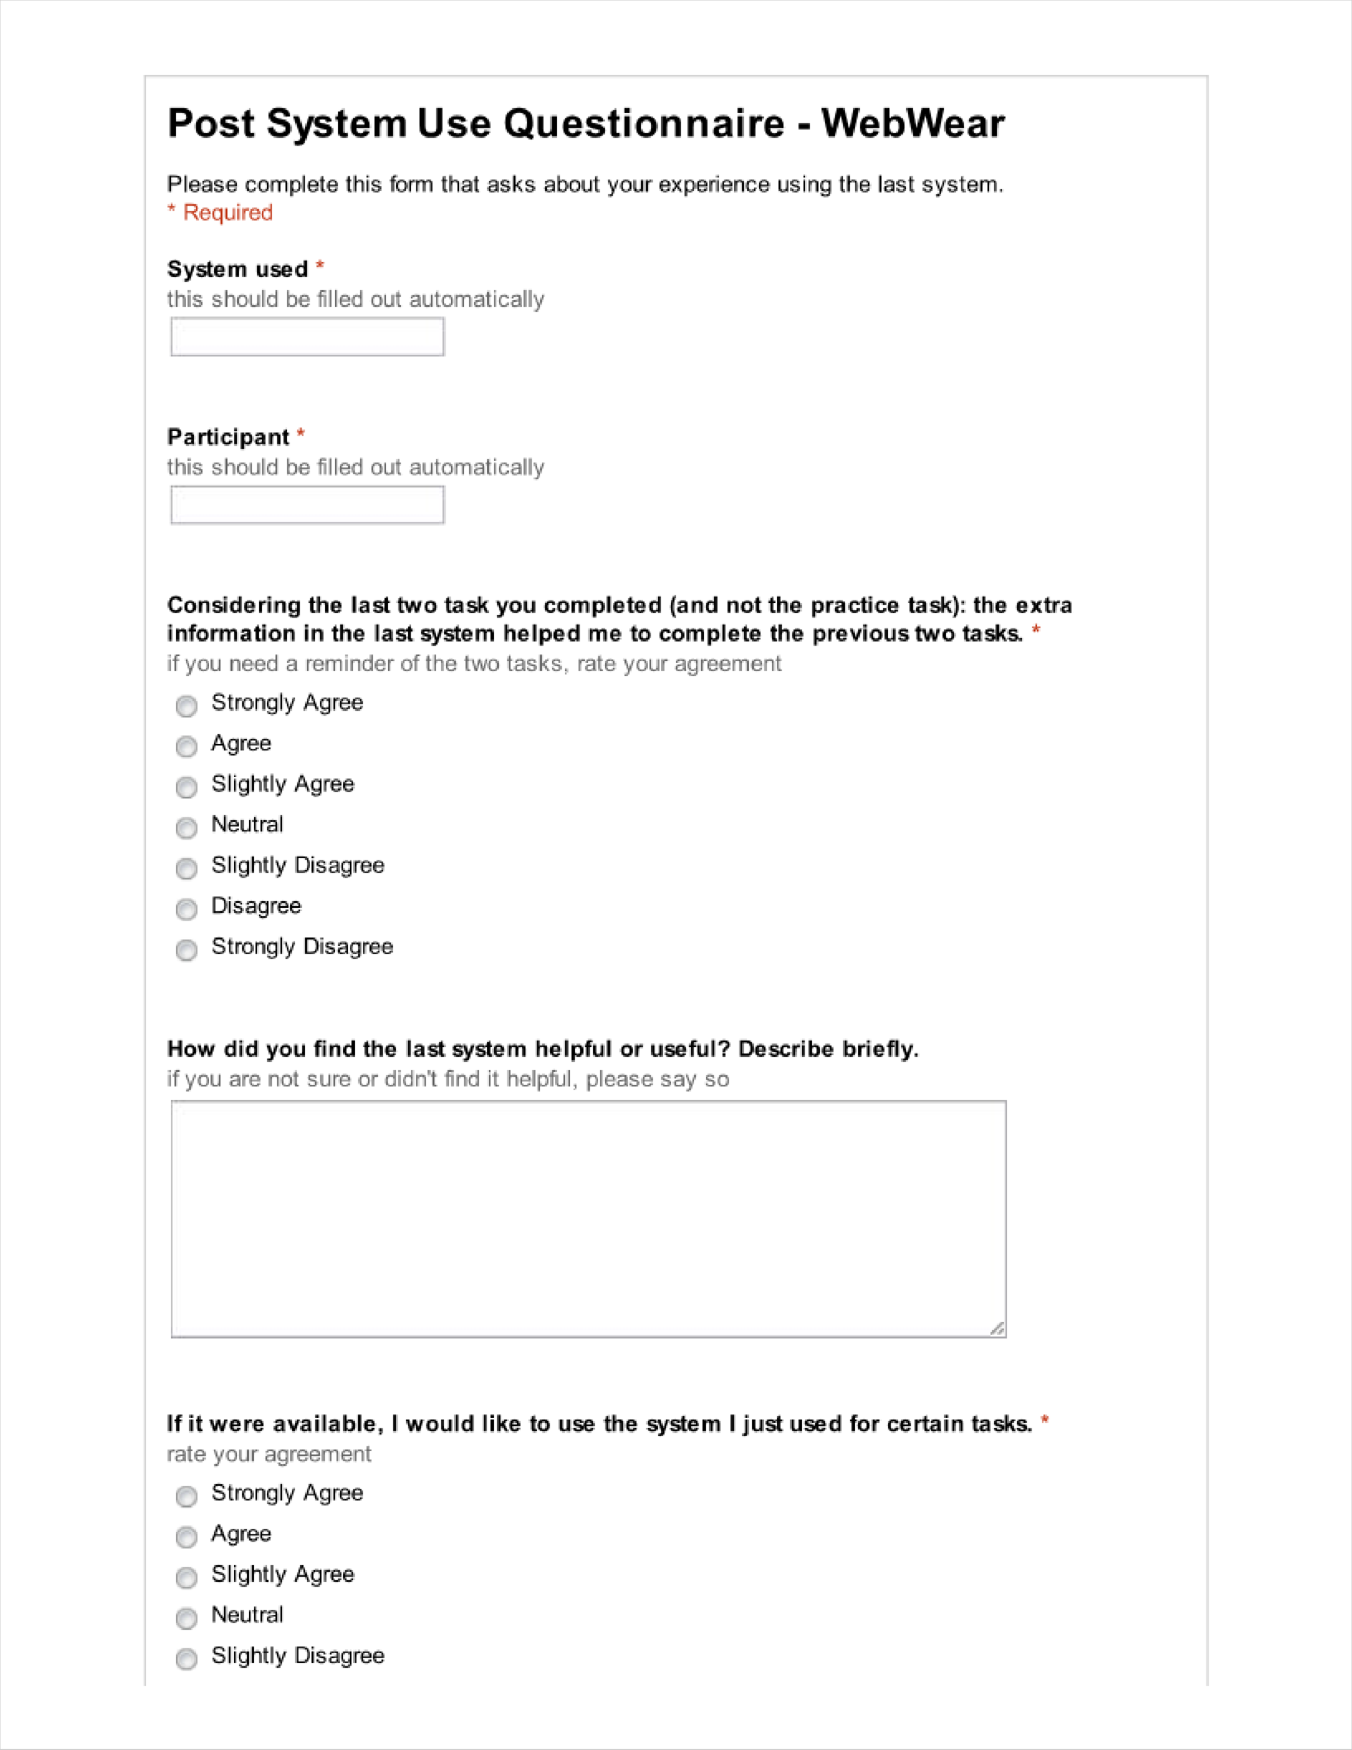

Post System Use Questionnaire - WebWear

Please complete this form that asks about your experience using the last system.
* Required

System used *
this should be filled out automatically

Participant *
this should be filled out automatically

Considering the last two task you completed (and not the practice task): the extra information in the last system helped me to complete the previous two tasks. *
if you need a reminder of the two tasks, rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

How did you find the last system helpful or useful? Describe briefly.
if you are not sure or didn't find it helpful, please say so

If it were available, I would like to use the system I just used for certain tasks. *
rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree

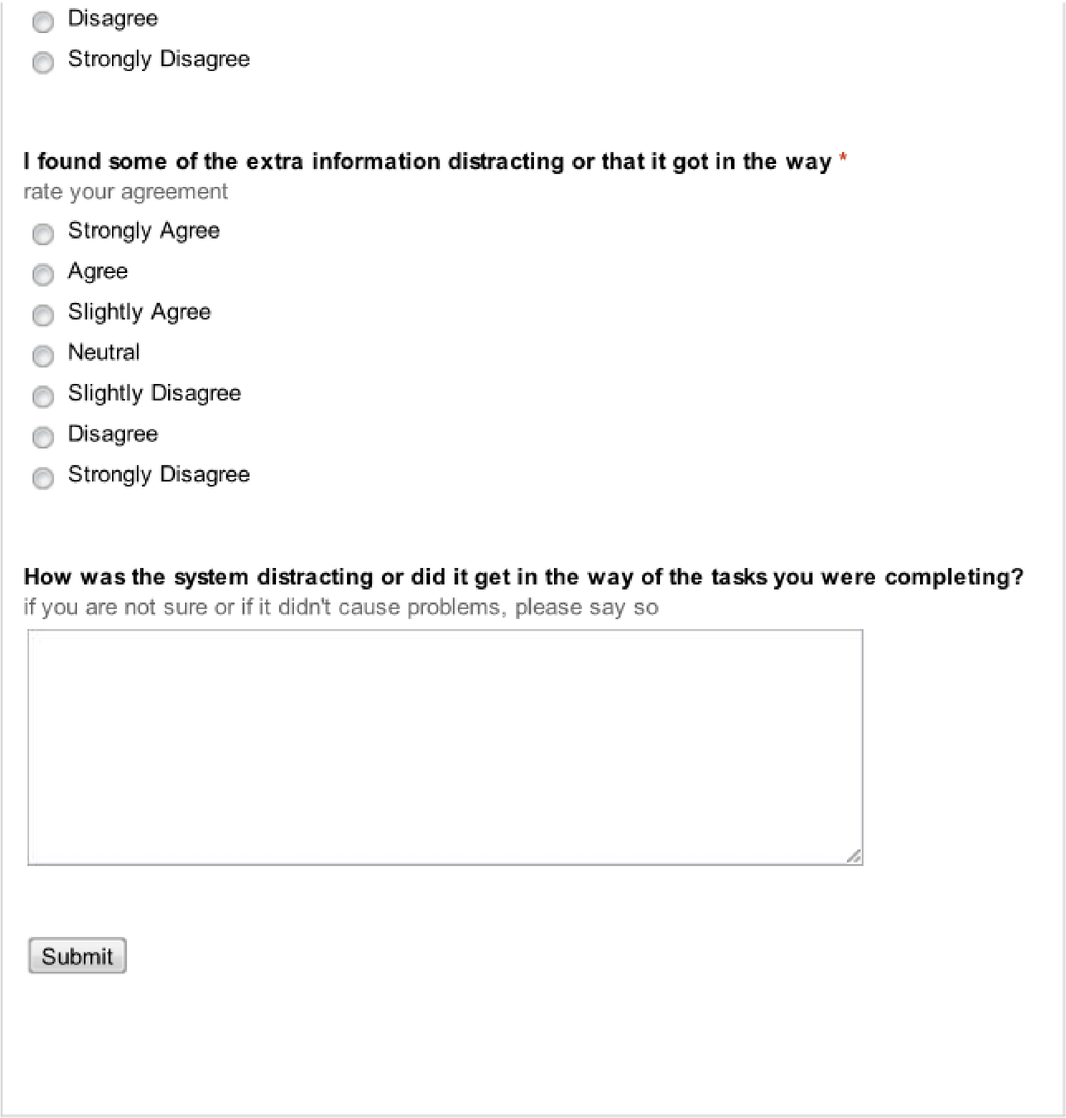

○ Disagree
○ Strongly Disagree

I found some of the extra information distracting or that it got in the way *

rate your agreement

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

How was the system distracting or did it get in the way of the tasks you were completing?

if you are not sure or if it didn't cause problems, please say so

WebWear – Post Experiment Questionnaire

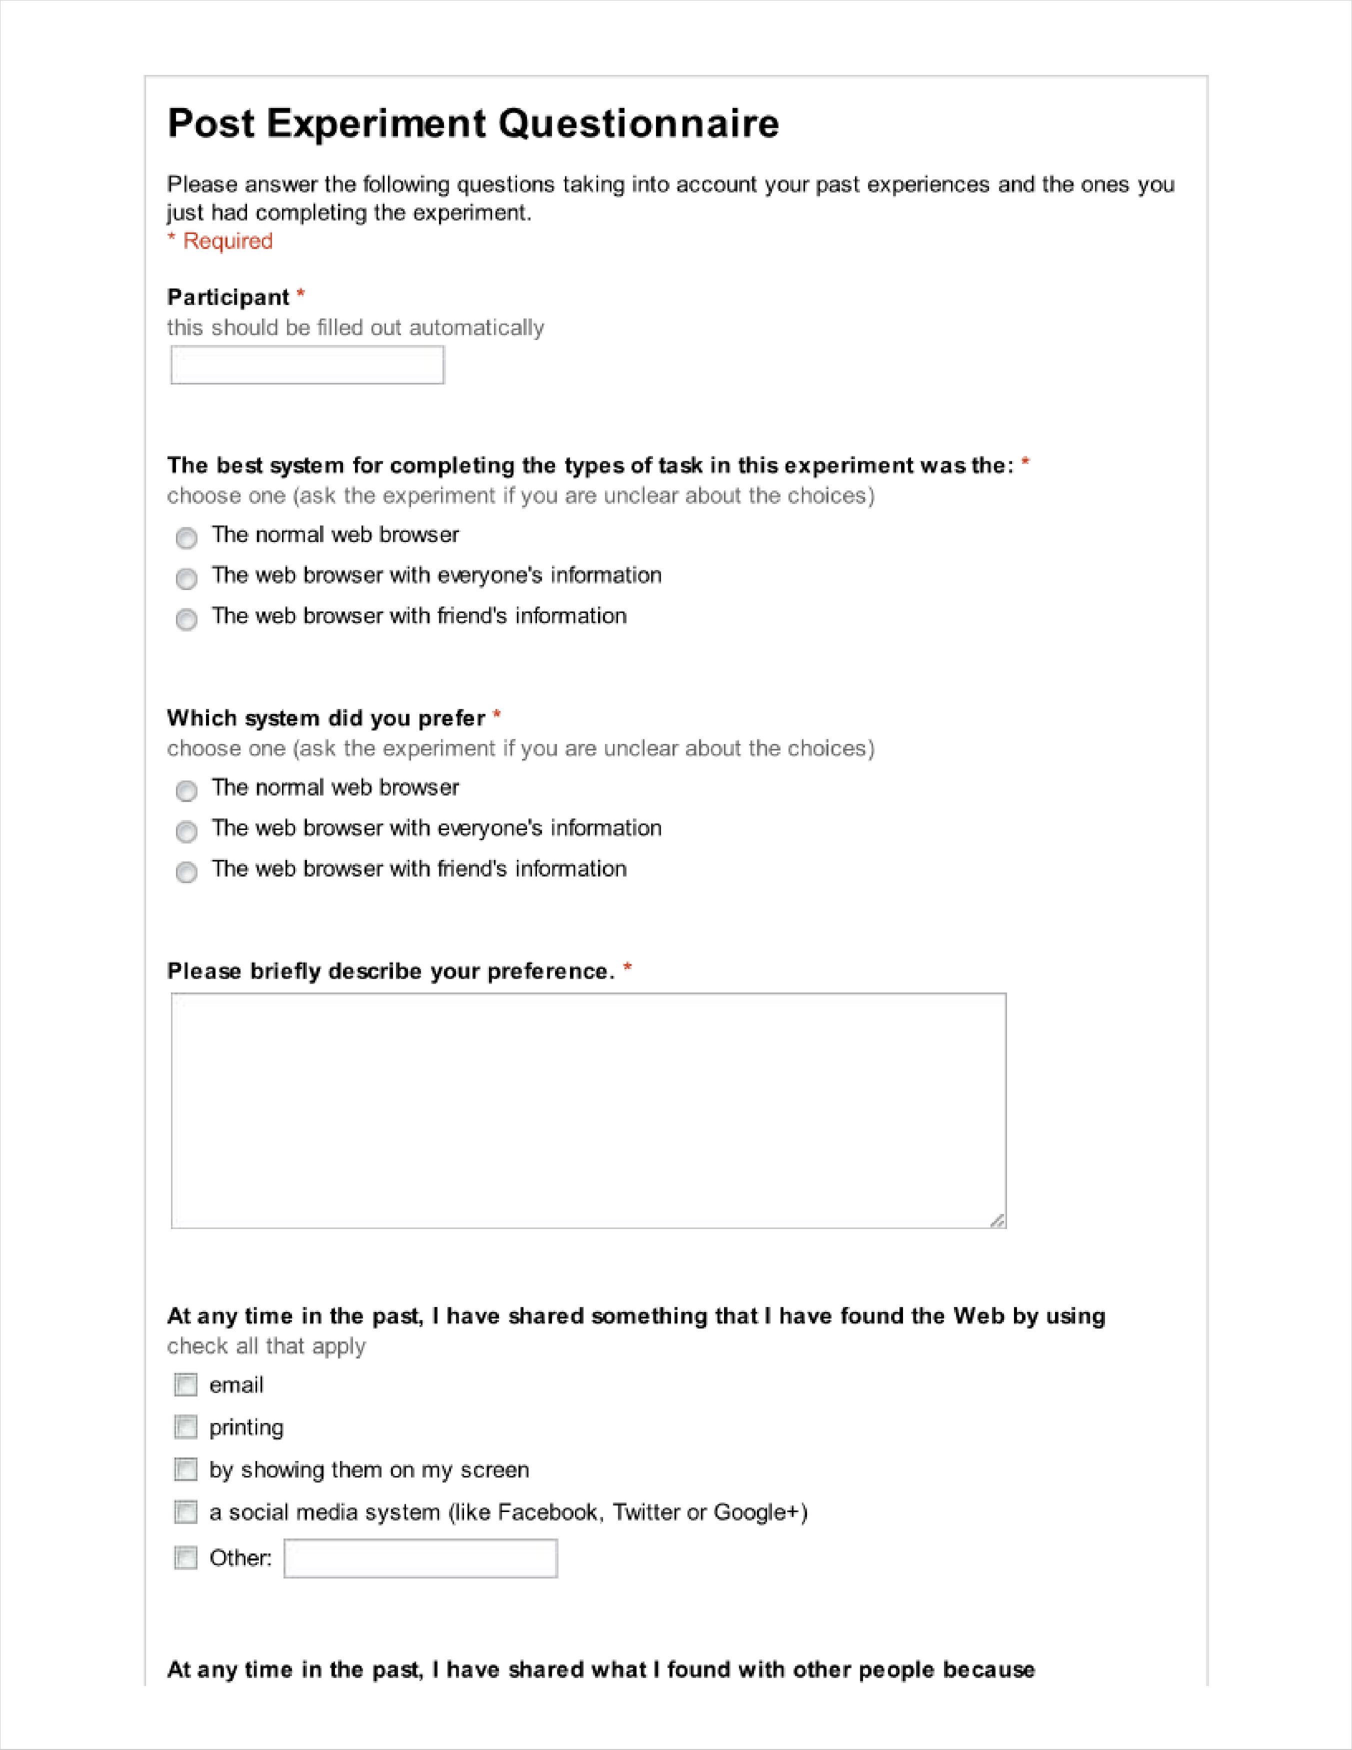

Post Experiment Questionnaire

Please answer the following questions taking into account your past experiences and the ones you just had completing the experiment.
* Required

Participant *
this should be filled out automatically

The best system for completing the types of task in this experiment was the: *
choose one (ask the experiment if you are unclear about the choices)

○ The normal web browser
○ The web browser with everyone's information
○ The web browser with friend's information

Which system did you prefer *
choose one (ask the experiment if you are unclear about the choices)

○ The normal web browser
○ The web browser with everyone's information
○ The web browser with friend's information

Please briefly describe your preference. *

At any time in the past, I have shared something that I have found the Web by using
check all that apply

☐ email
☐ printing
☐ by showing them on my screen
☐ a social media system (like Facebook, Twitter or Google+)
☐ Other:

At any time in the past, I have shared what I found with other people because

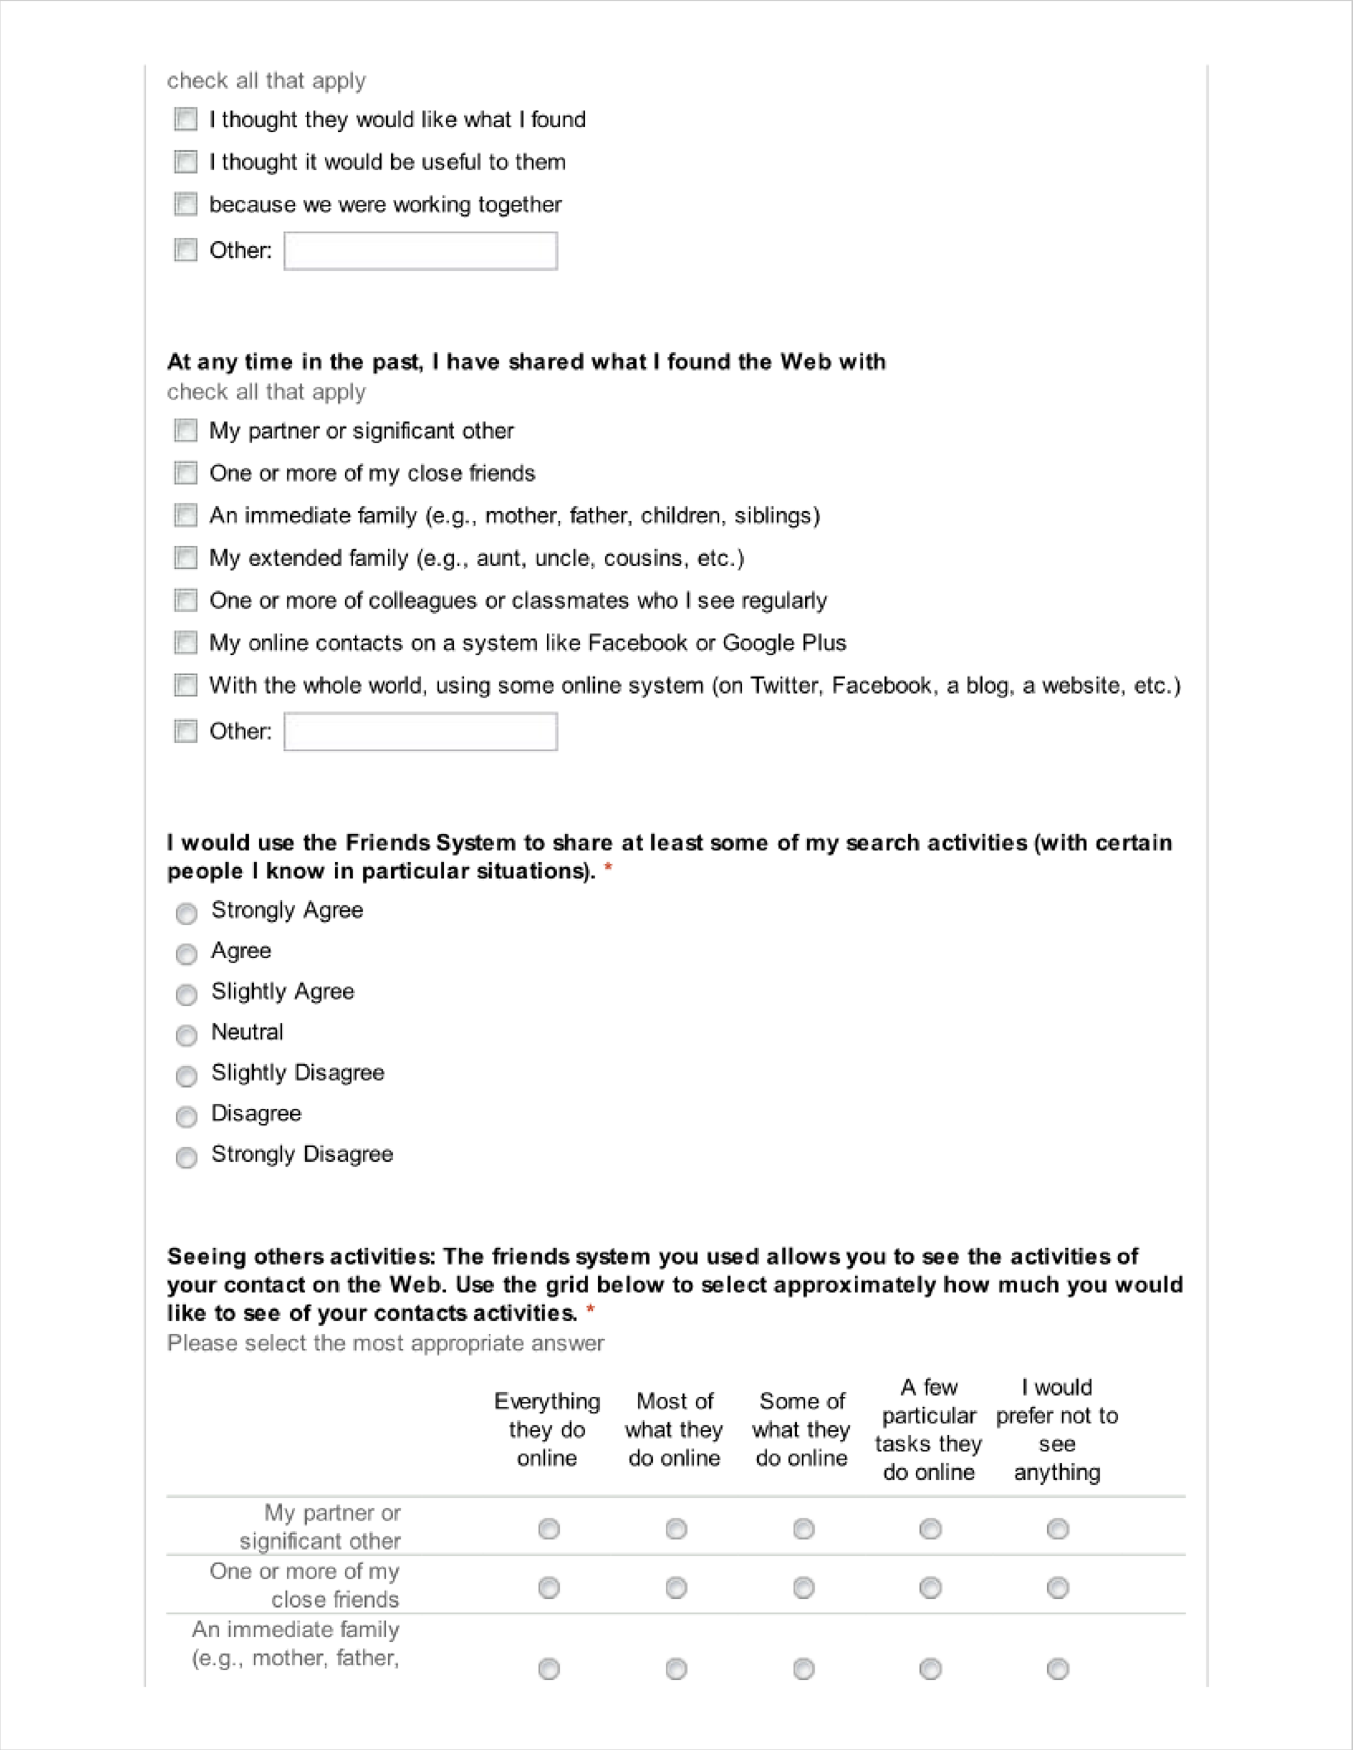

check all that apply

☐ I thought they would like what I found

☐ I thought it would be useful to them

☐ because we were working together

☐ Other:

At any time in the past, I have shared what I found the Web with

check all that apply

☐ My partner or significant other

☐ One or more of my close friends

☐ An immediate family (e.g., mother, father, children, siblings)

☐ My extended family (e.g., aunt, uncle, cousins, etc.)

☐ One or more of colleagues or classmates who I see regularly

☐ My online contacts on a system like Facebook or Google Plus

☐ With the whole world, using some online system (on Twitter, Facebook, a blog, a website, etc.)

☐ Other:

I would use the Friends System to share at least some of my search activities (with certain people I know in particular situations). *

○ Strongly Agree

○ Agree

○ Slightly Agree

○ Neutral

○ Slightly Disagree

○ Disagree

○ Strongly Disagree

Seeing others activities: The friends system you used allows you to see the activities of your contact on the Web. Use the grid below to select approximately how much you would like to see of your contacts activities. *

Please select the most appropriate answer

	Everything they do online	Most of what they do online	Some of what they do online	A few particular tasks they do online	I would prefer not to see anything
My partner or significant other	○	○	○	○	○
One or more of my close friends	○	○	○	○	○
An immediate family (e.g., mother, father,	○	○	○	○	○

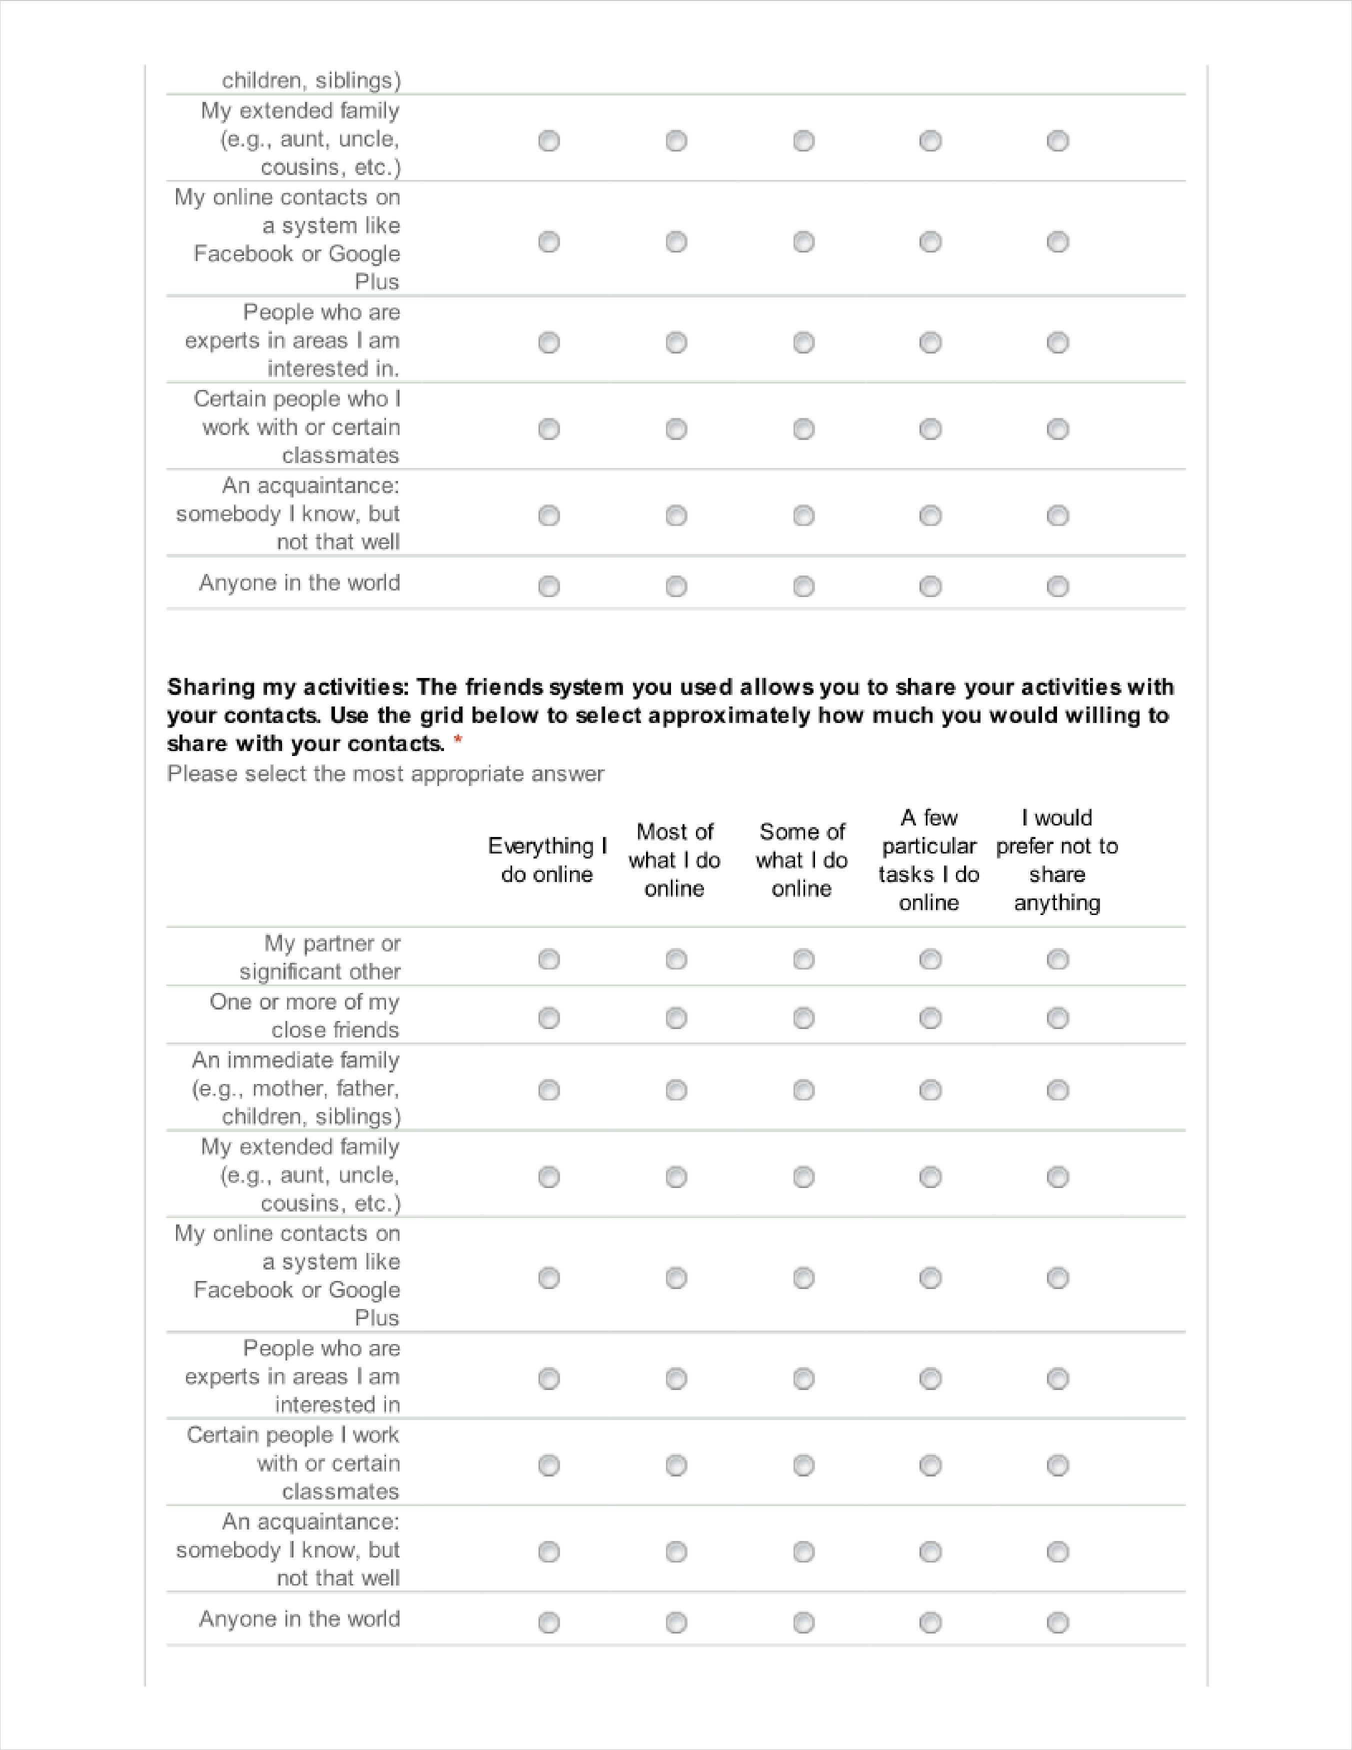

children, siblings)					
My extended family (e.g., aunt, uncle, cousins, etc.)	○	○	○	○	○
My online contacts on a system like Facebook or Google Plus	○	○	○	○	○
People who are experts in areas I am interested in.	○	○	○	○	○
Certain people who I work with or certain classmates	○	○	○	○	○
An acquaintance: somebody I know, but not that well	○	○	○	○	○
Anyone in the world	○	○	○	○	○

Sharing my activities: The friends system you used allows you to share your activities with your contacts. Use the grid below to select approximately how much you would willing to share with your contacts. *

Please select the most appropriate answer

	Everything I do online	Most of what I do online	Some of what I do online	A few particular tasks I do online	I would prefer not to share anything
My partner or significant other	○	○	○	○	○
One or more of my close friends	○	○	○	○	○
An immediate family (e.g., mother, father, children, siblings)	○	○	○	○	○
My extended family (e.g., aunt, uncle, cousins, etc.)	○	○	○	○	○
My online contacts on a system like Facebook or Google Plus	○	○	○	○	○
People who are experts in areas I am interested in	○	○	○	○	○
Certain people I work with or certain classmates	○	○	○	○	○
An acquaintance: somebody I know, but not that well	○	○	○	○	○
Anyone in the world	○	○	○	○	○

Do you have any comments about the last task you would like to share?

describe briefly

Diary Study – Informed Consent Form (webpage for participant records, was also available in the study system)

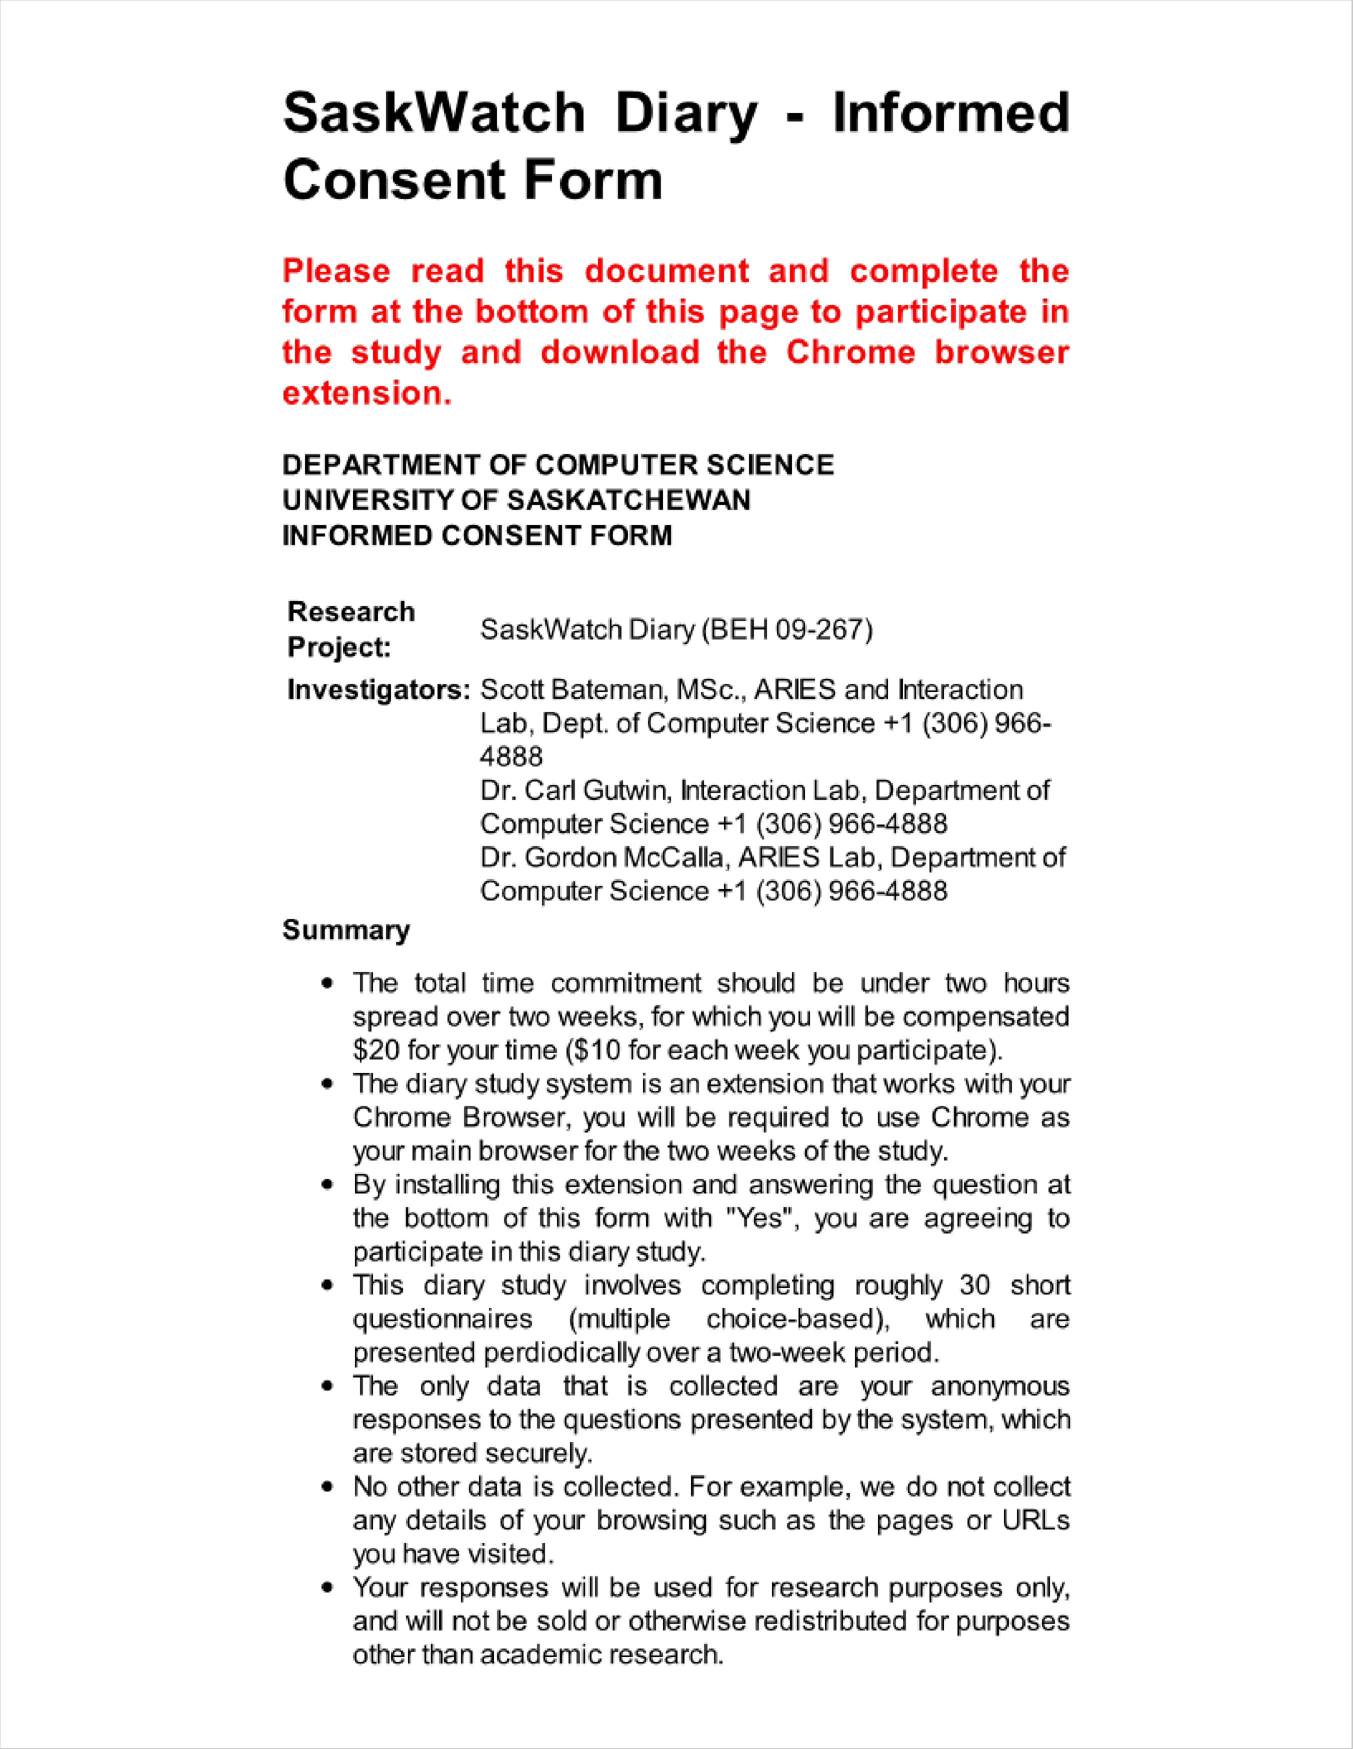

SaskWatch Diary - Informed Consent Form

Please read this document and complete the form at the bottom of this page to participate in the study and download the Chrome browser extension.

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF SASKATCHEWAN
INFORMED CONSENT FORM

Research Project:	SaskWatch Diary (BEH 09-267)
Investigators:	Scott Bateman, MSc., ARIES and Interaction Lab, Dept. of Computer Science +1 (306) 966-4888 Dr. Carl Gutwin, Interaction Lab, Department of Computer Science +1 (306) 966-4888 Dr. Gordon McCalla, ARIES Lab, Department of Computer Science +1 (306) 966-4888

Summary

- The total time commitment should be under two hours spread over two weeks, for which you will be compensated $20 for your time ($10 for each week you participate).
- The diary study system is an extension that works with your Chrome Browser, you will be required to use Chrome as your main browser for the two weeks of the study.
- By installing this extension and answering the question at the bottom of this form with "Yes", you are agreeing to participate in this diary study.
- This diary study involves completing roughly 30 short questionnaires (multiple choice-based), which are presented perdiodically over a two-week period.
- The only data that is collected are your anonymous responses to the questions presented by the system, which are stored securely.
- No other data is collected. For example, we do not collect any details of your browsing such as the pages or URLs you have visited.
- Your responses will be used for research purposes only, and will not be sold or otherwise redistributed for purposes other than academic research.

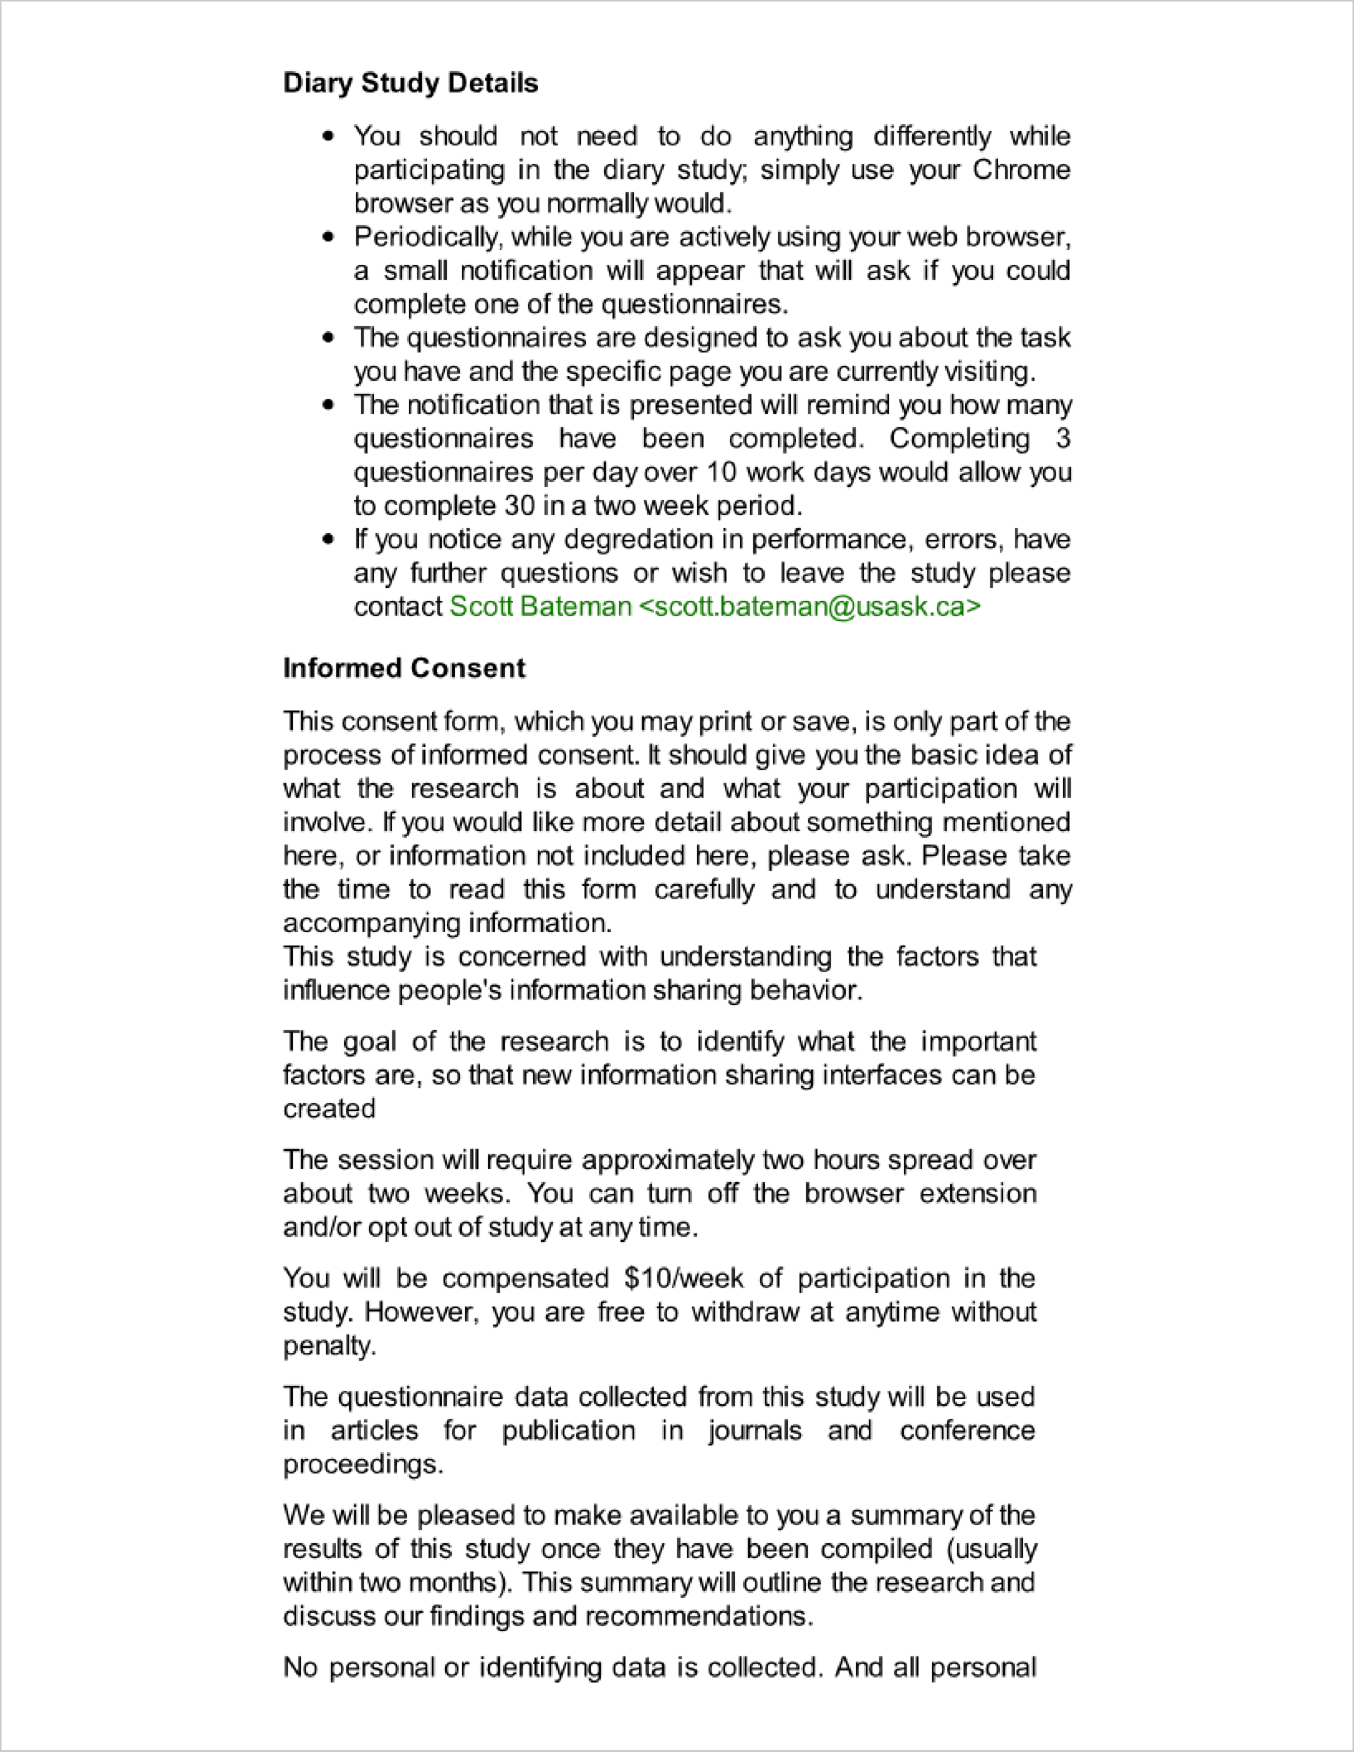

Diary Study Details

- You should not need to do anything differently while participating in the diary study; simply use your Chrome browser as you normally would.
- Periodically, while you are actively using your web browser, a small notification will appear that will ask if you could complete one of the questionnaires.
- The questionnaires are designed to ask you about the task you have and the specific page you are currently visiting.
- The notification that is presented will remind you how many questionnaires have been completed. Completing 3 questionnaires per day over 10 work days would allow you to complete 30 in a two week period.
- If you notice any degredation in performance, errors, have any further questions or wish to leave the study please contact Scott Bateman <scott.bateman@usask.ca>

Informed Consent

This consent form, which you may print or save, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail about something mentioned here, or information not included here, please ask. Please take the time to read this form carefully and to understand any accompanying information.
This study is concerned with understanding the factors that influence people's information sharing behavior.

The goal of the research is to identify what the important factors are, so that new information sharing interfaces can be created

The session will require approximately two hours spread over about two weeks. You can turn off the browser extension and/or opt out of study at any time.

You will be compensated $10/week of participation in the study. However, you are free to withdraw at anytime without penalty.

The questionnaire data collected from this study will be used in articles for publication in journals and conference proceedings.

We will be pleased to make available to you a summary of the results of this study once they have been compiled (usually within two months). This summary will outline the research and discuss our findings and recommendations.

No personal or identifying data is collected. And all personal

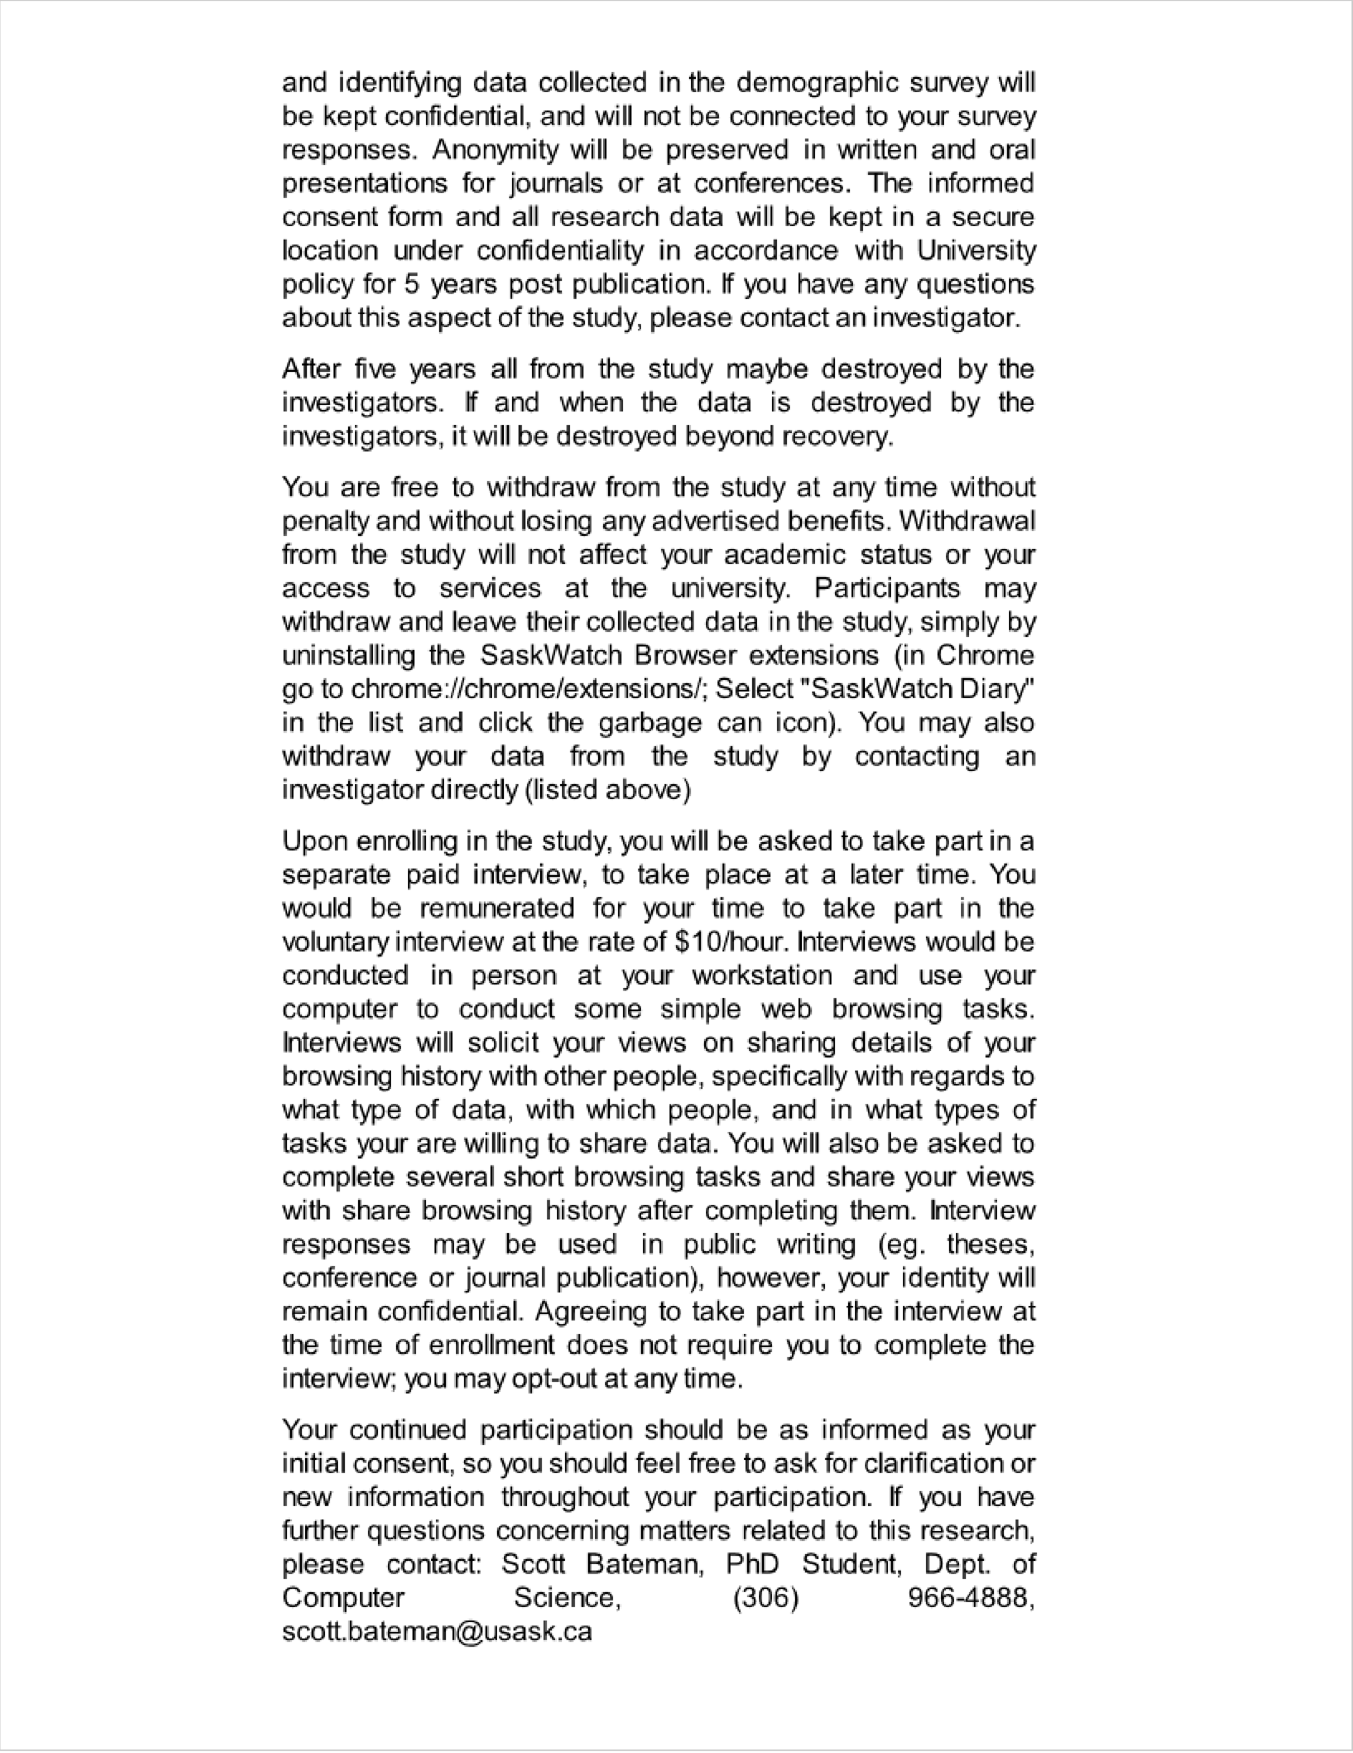

and identifying data collected in the demographic survey will be kept confidential, and will not be connected to your survey responses. Anonymity will be preserved in written and oral presentations for journals or at conferences. The informed consent form and all research data will be kept in a secure location under confidentiality in accordance with University policy for 5 years post publication. If you have any questions about this aspect of the study, please contact an investigator.

After five years all from the study maybe destroyed by the investigators. If and when the data is destroyed by the investigators, it will be destroyed beyond recovery.

You are free to withdraw from the study at any time without penalty and without losing any advertised benefits. Withdrawal from the study will not affect your academic status or your access to services at the university. Participants may withdraw and leave their collected data in the study, simply by uninstalling the SaskWatch Browser extensions (in Chrome go to chrome://chrome/extensions/; Select "SaskWatch Diary" in the list and click the garbage can icon). You may also withdraw your data from the study by contacting an investigator directly (listed above)

Upon enrolling in the study, you will be asked to take part in a separate paid interview, to take place at a later time. You would be remunerated for your time to take part in the voluntary interview at the rate of $10/hour. Interviews would be conducted in person at your workstation and use your computer to conduct some simple web browsing tasks. Interviews will solicit your views on sharing details of your browsing history with other people, specifically with regards to what type of data, with which people, and in what types of tasks your are willing to share data. You will also be asked to complete several short browsing tasks and share your views with share browsing history after completing them. Interview responses may be used in public writing (eg. theses, conference or journal publication), however, your identity will remain confidential. Agreeing to take part in the interview at the time of enrollment does not require you to complete the interview; you may opt-out at any time.

Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation. If you have further questions concerning matters related to this research, please contact: Scott Bateman, PhD Student, Dept. of Computer Science, (306) 966-4888, scott.bateman@usask.ca

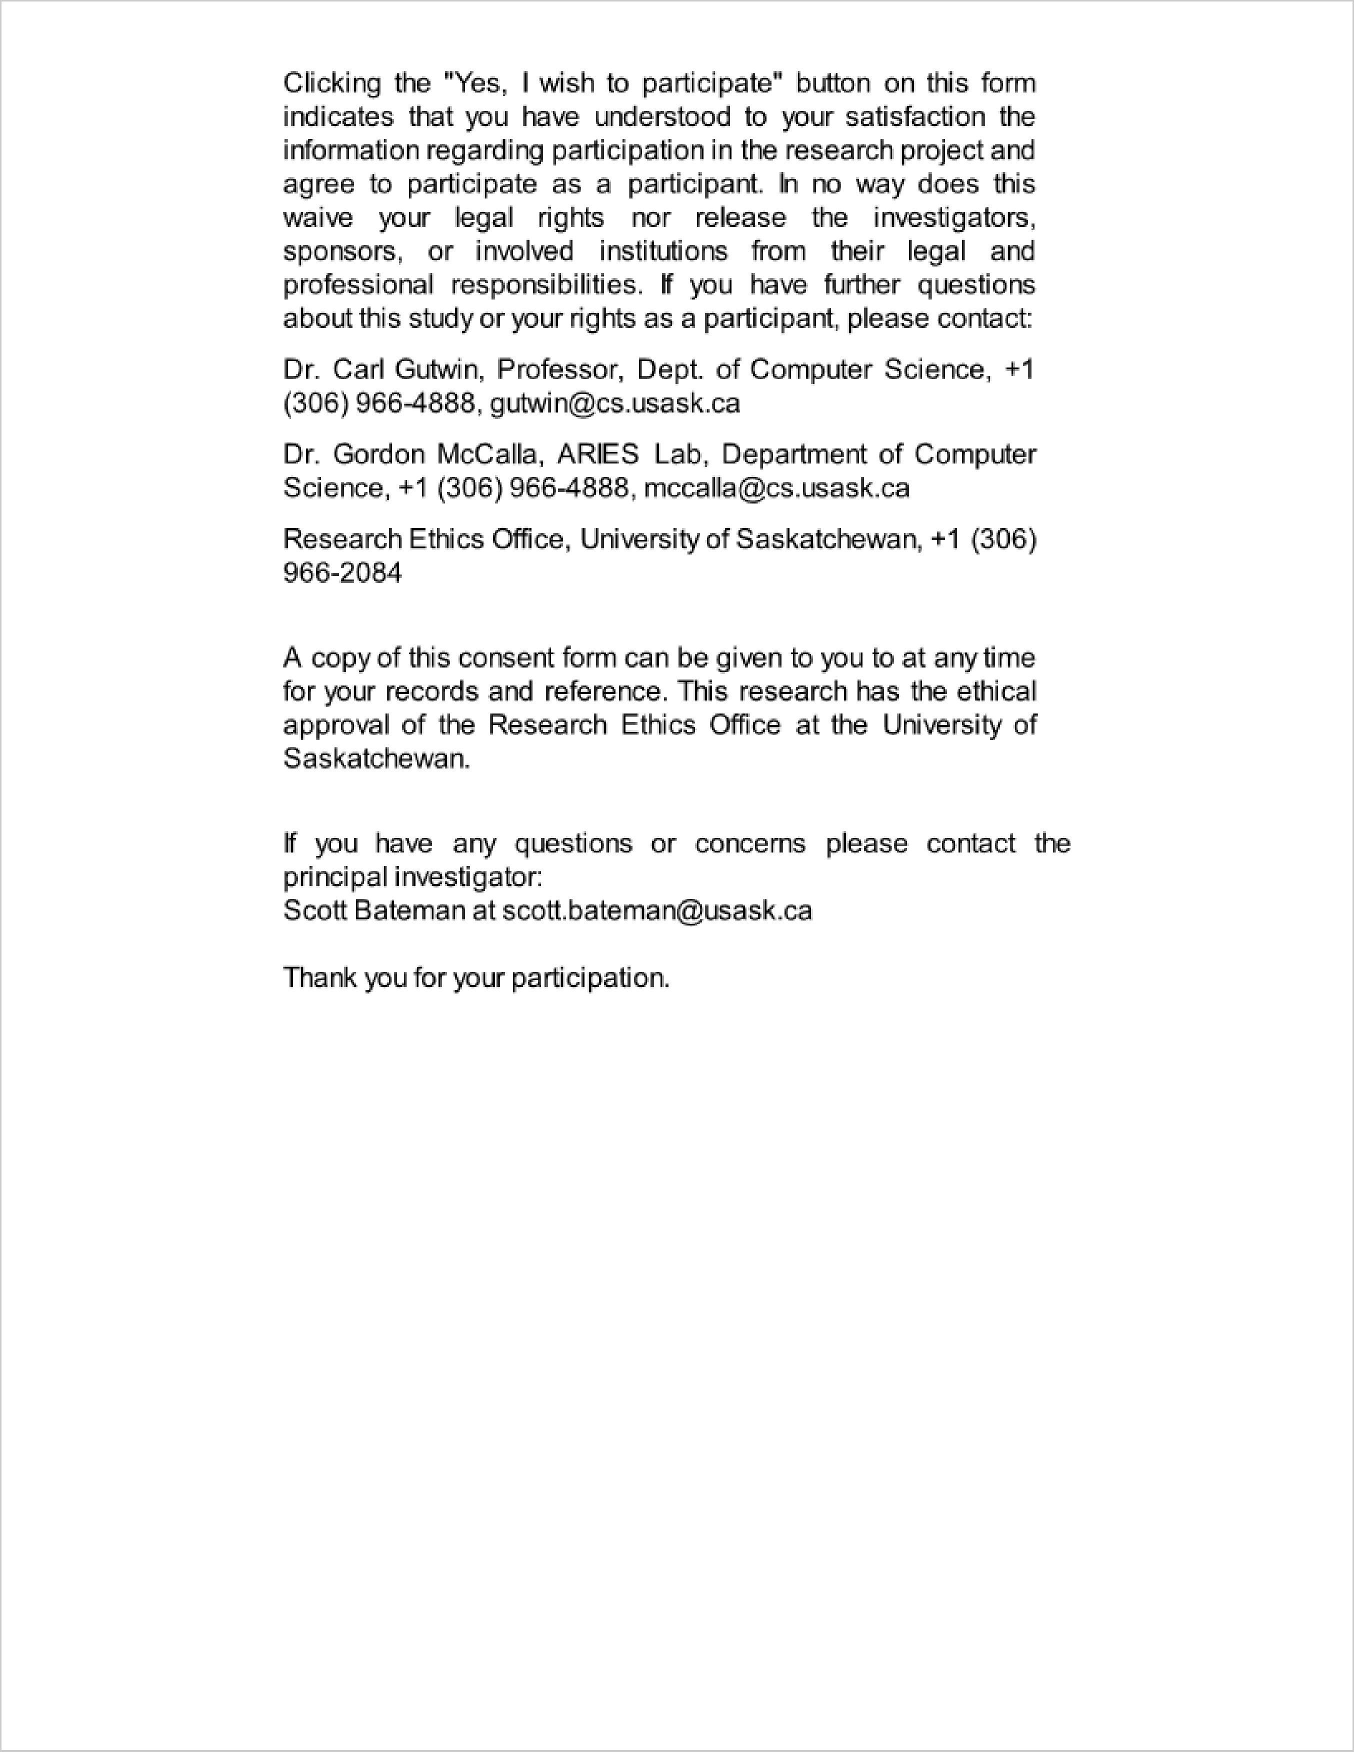

Clicking the "Yes, I wish to participate" button on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a participant. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. If you have further questions about this study or your rights as a participant, please contact:

Dr. Carl Gutwin, Professor, Dept. of Computer Science, +1 (306) 966-4888, gutwin@cs.usask.ca

Dr. Gordon McCalla, ARIES Lab, Department of Computer Science, +1 (306) 966-4888, mccalla@cs.usask.ca

Research Ethics Office, University of Saskatchewan, +1 (306) 966-2084

A copy of this consent form can be given to you to at any time for your records and reference. This research has the ethical approval of the Research Ethics Office at the University of Saskatchewan.

If you have any questions or concerns please contact the principal investigator:
Scott Bateman at scott.bateman@usask.ca

Thank you for your participation.

Diary Study – Survey

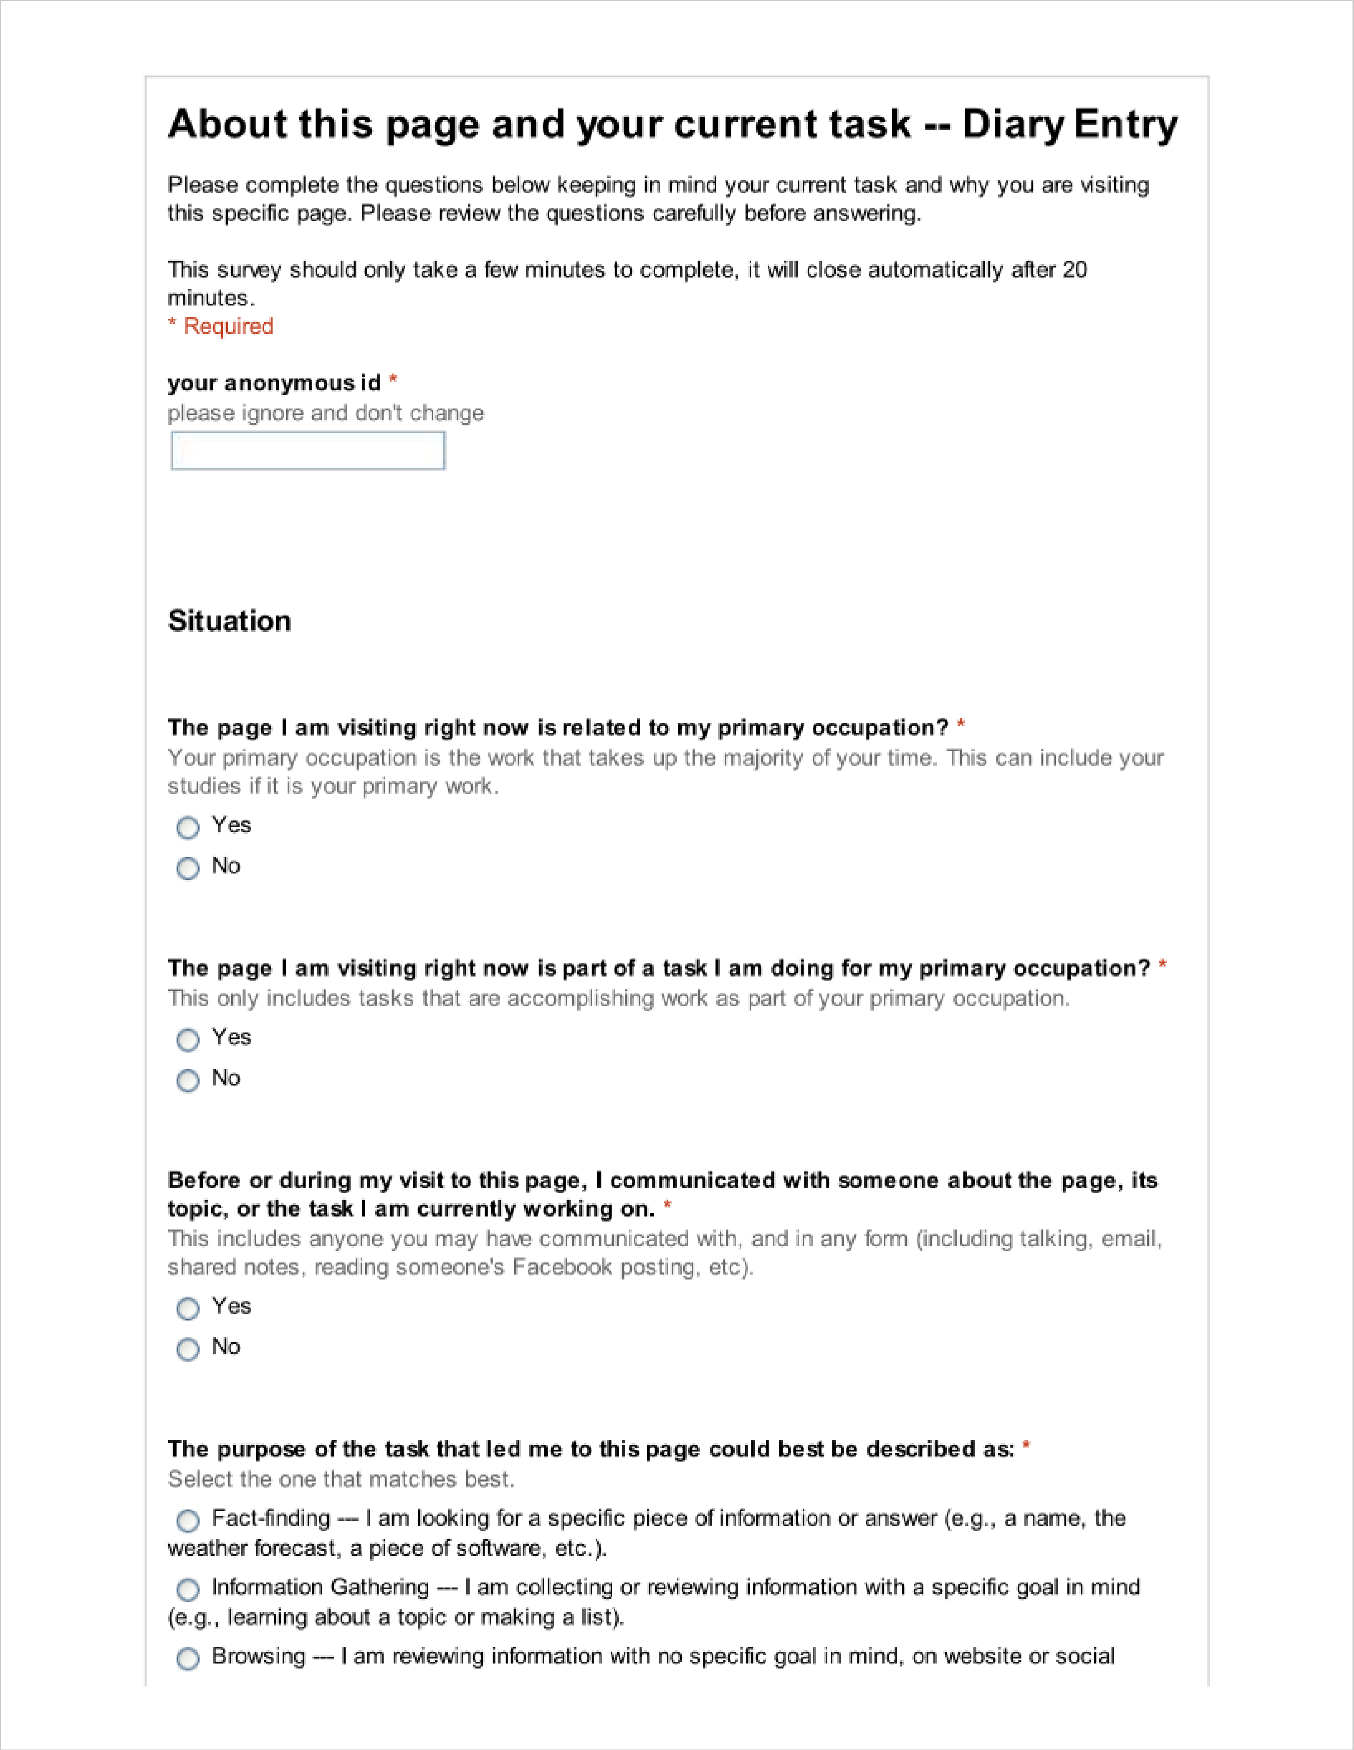

About this page and your current task -- Diary Entry

Please complete the questions below keeping in mind your current task and why you are visiting this specific page. Please review the questions carefully before answering.

This survey should only take a few minutes to complete, it will close automatically after 20 minutes.
* Required

your anonymous id *
please ignore and don't change

Situation

The page I am visiting right now is related to my primary occupation? *
Your primary occupation is the work that takes up the majority of your time. This can include your studies if it is your primary work.

○ Yes
○ No

The page I am visiting right now is part of a task I am doing for my primary occupation? *
This only includes tasks that are accomplishing work as part of your primary occupation.

○ Yes
○ No

Before or during my visit to this page, I communicated with someone about the page, its topic, or the task I am currently working on. *
This includes anyone you may have communicated with, and in any form (including talking, email, shared notes, reading someone's Facebook posting, etc).

○ Yes
○ No

The purpose of the task that led me to this page could best be described as: *
Select the one that matches best.

○ Fact-finding --- I am looking for a specific piece of information or answer (e.g., a name, the weather forecast, a piece of software, etc.).
○ Information Gathering --- I am collecting or reviewing information with a specific goal in mind (e.g., learning about a topic or making a list).
○ Browsing --- I am reviewing information with no specific goal in mind, on website or social

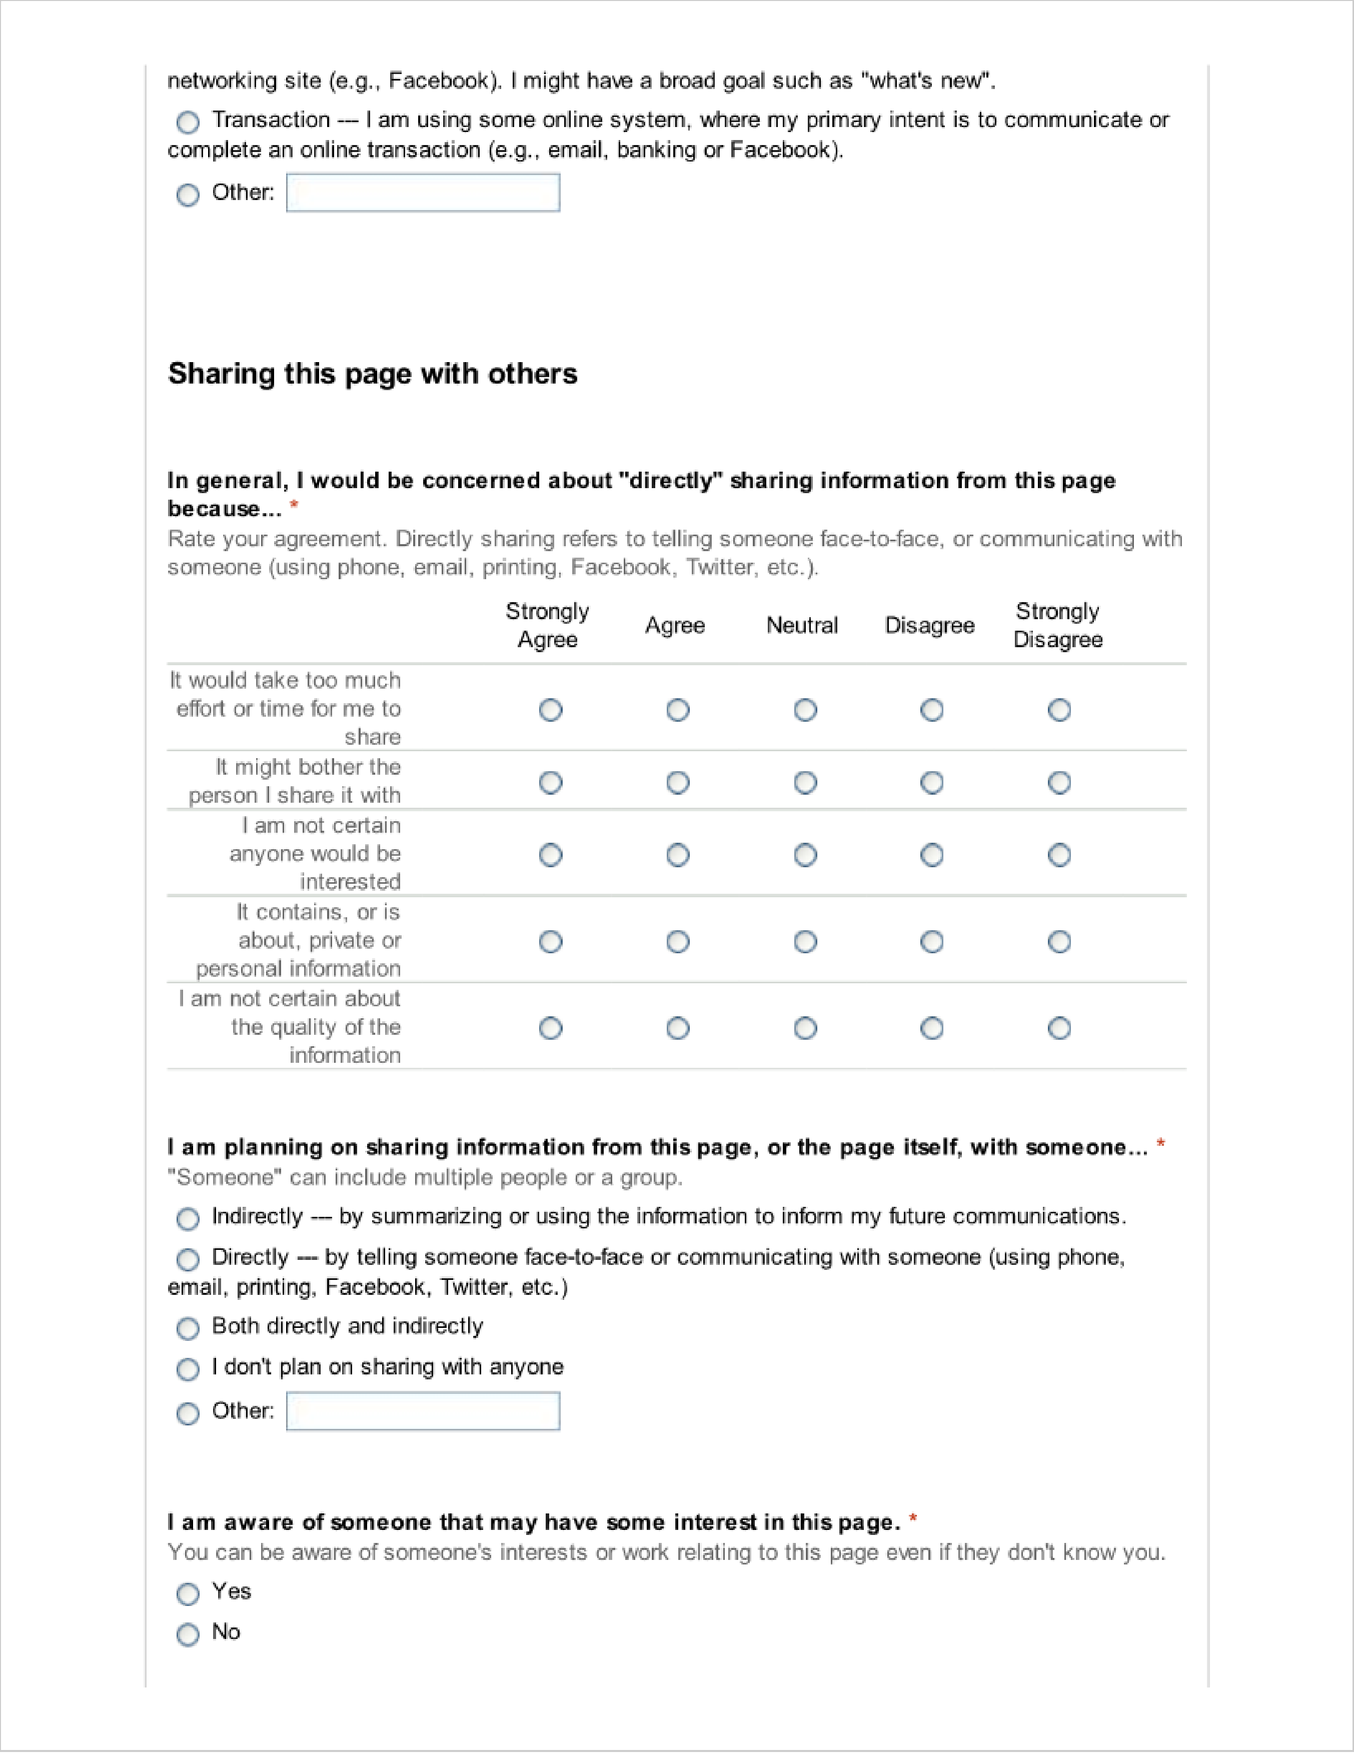

networking site (e.g., Facebook). I might have a broad goal such as "what's new".

○ Transaction --- I am using some online system, where my primary intent is to communicate or complete an online transaction (e.g., email, banking or Facebook).

○ Other: ☐

Sharing this page with others

In general, I would be concerned about "directly" sharing information from this page because... *

Rate your agreement. Directly sharing refers to telling someone face-to-face, or communicating with someone (using phone, email, printing, Facebook, Twitter, etc.).

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
It would take too much effort or time for me to share	○	○	○	○	○
It might bother the person I share it with	○	○	○	○	○
I am not certain anyone would be interested	○	○	○	○	○
It contains, or is about, private or personal information	○	○	○	○	○
I am not certain about the quality of the information	○	○	○	○	○

I am planning on sharing information from this page, or the page itself, with someone... *

"Someone" can include multiple people or a group.

○ Indirectly --- by summarizing or using the information to inform my future communications.

○ Directly --- by telling someone face-to-face or communicating with someone (using phone, email, printing, Facebook, Twitter, etc.)

○ Both directly and indirectly

○ I don't plan on sharing with anyone

○ Other: ☐

I am aware of someone that may have some interest in this page. *

You can be aware of someone's interests or work relating to this page even if they don't know you.

○ Yes

○ No

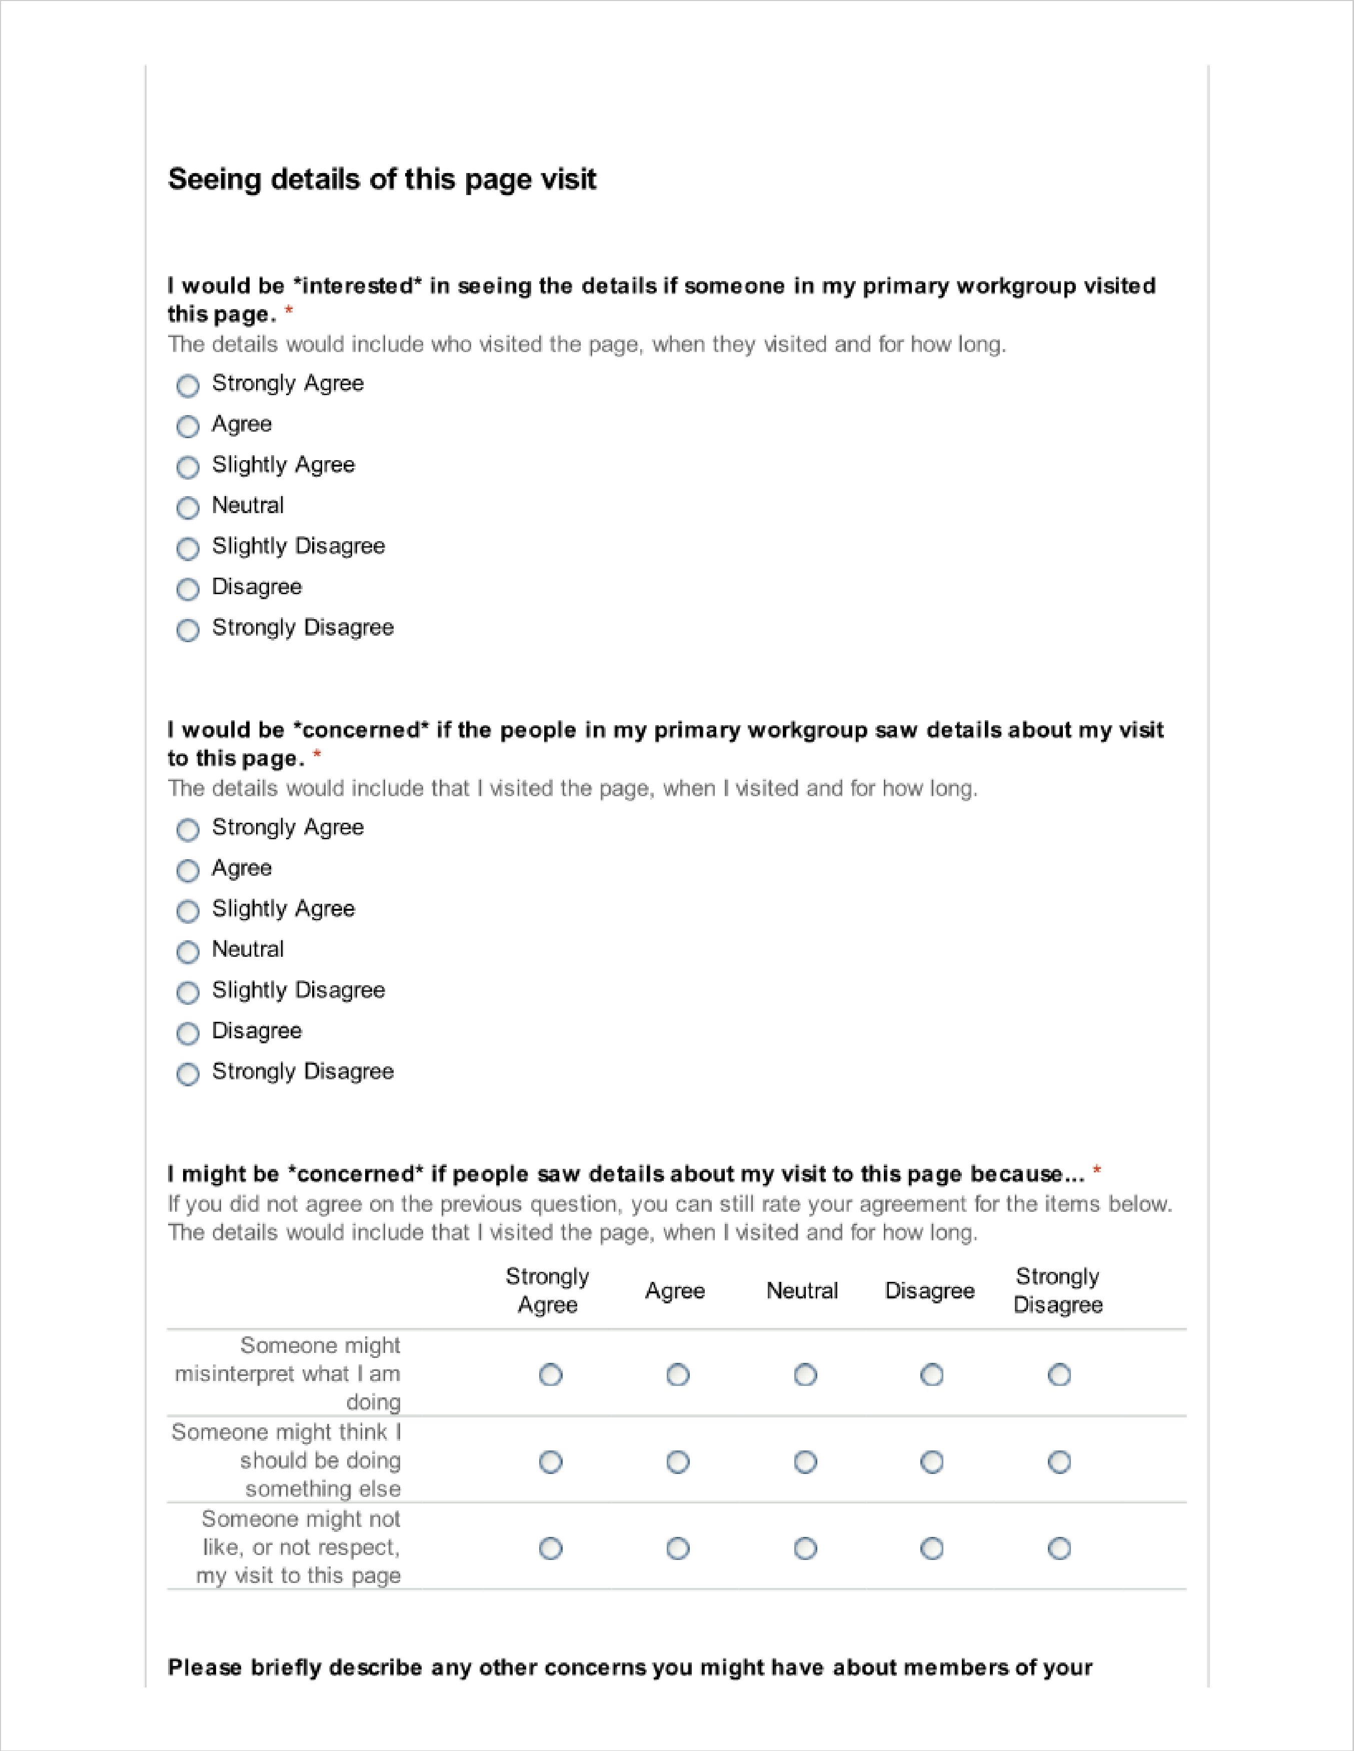

Seeing details of this page visit

I would be *interested* in seeing the details if someone in my primary workgroup visited this page. *

The details would include who visited the page, when they visited and for how long.

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I would be *concerned* if the people in my primary workgroup saw details about my visit to this page. *

The details would include that I visited the page, when I visited and for how long.

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

I might be *concerned* if people saw details about my visit to this page because... *

If you did not agree on the previous question, you can still rate your agreement for the items below. The details would include that I visited the page, when I visited and for how long.

	Strongly Agree	Agree	Neutral	Disagree	Strongly Disagree
Someone might misinterpret what I am doing	○	○	○	○	○
Someone might think I should be doing something else	○	○	○	○	○
Someone might not like, or not respect, my visit to this page	○	○	○	○	○

Please briefly describe any other concerns you might have about members of your

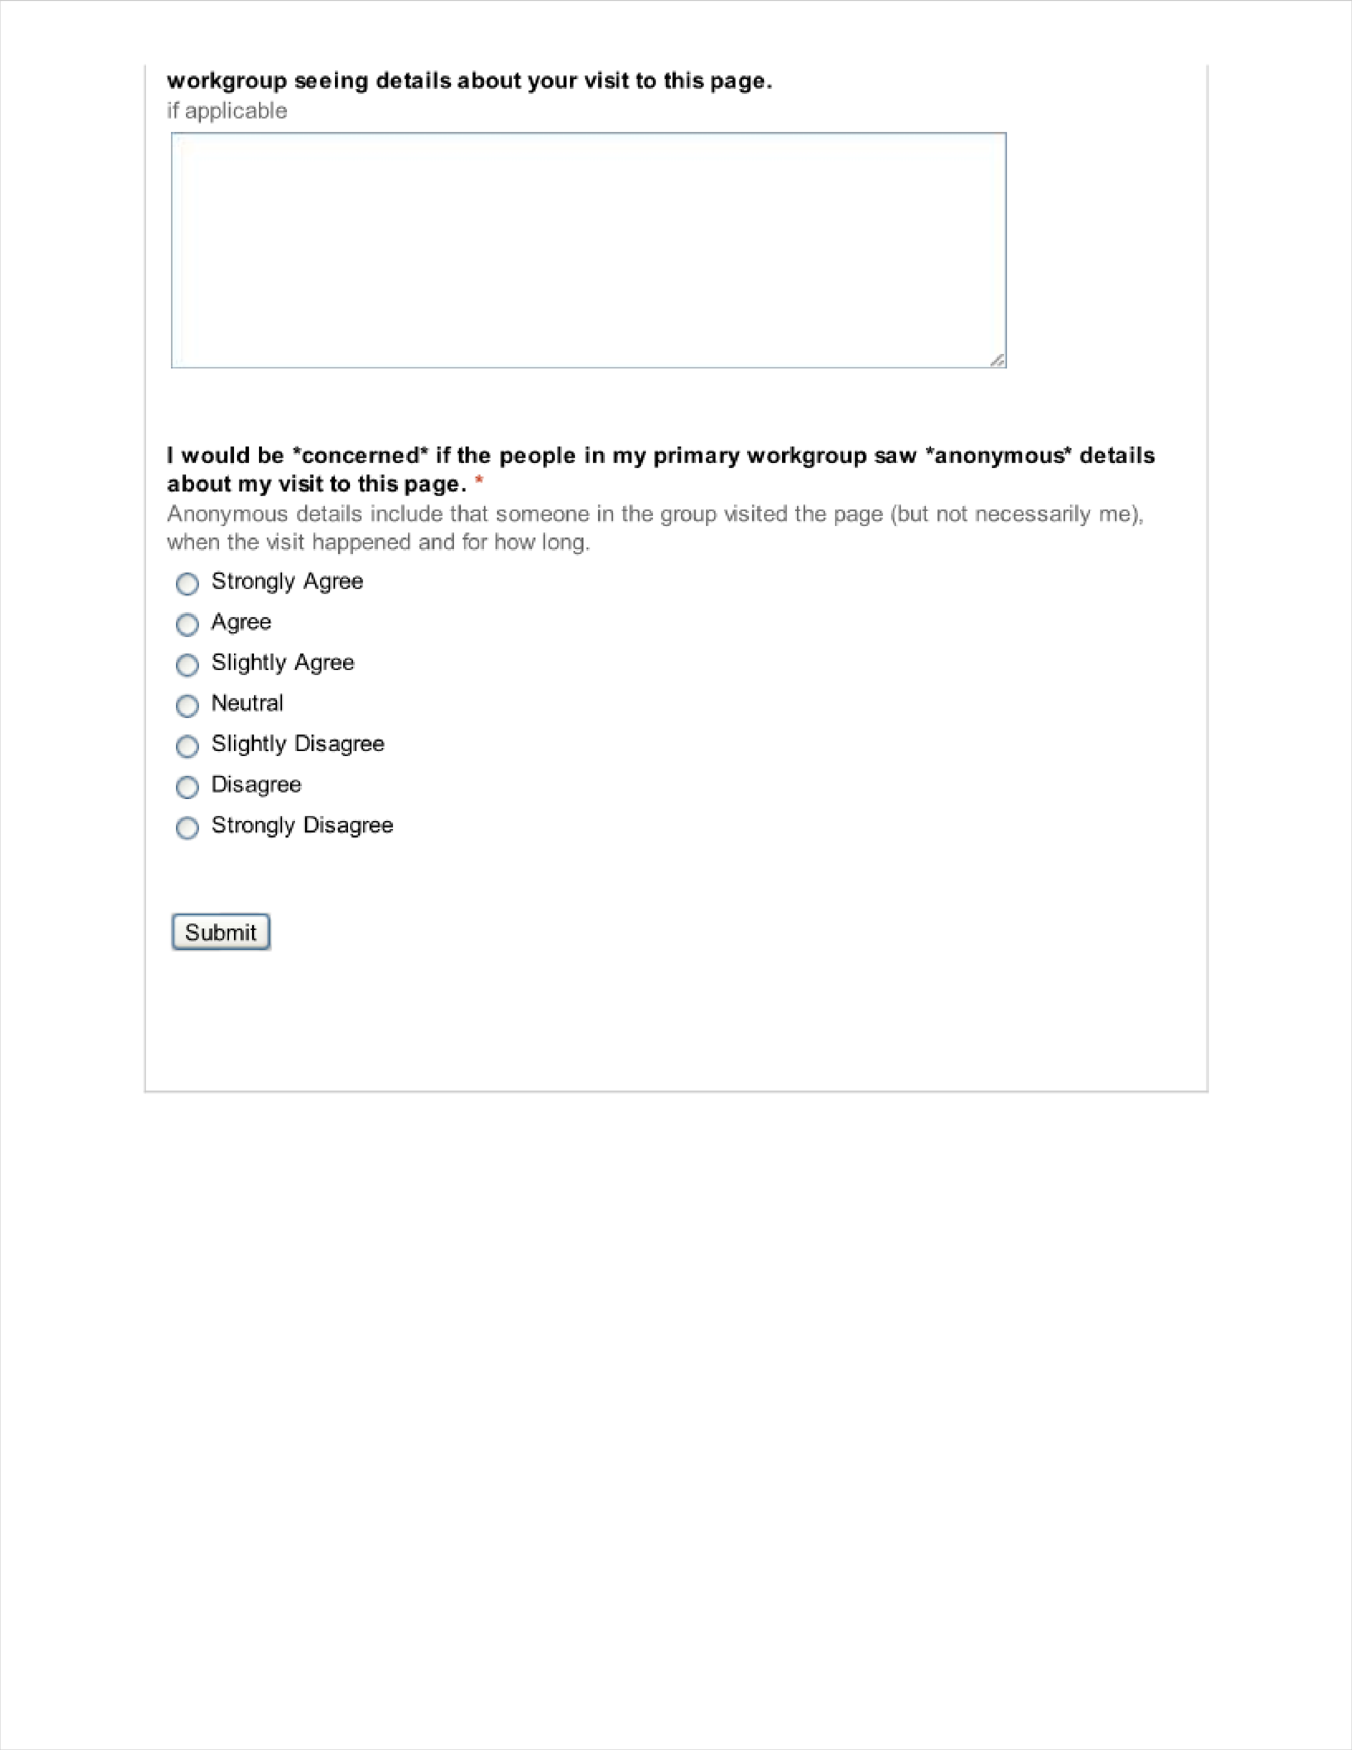

workgroup seeing details about your visit to this page.

if applicable

I would be *concerned* if the people in my primary workgroup saw *anonymous* details about my visit to this page. *

Anonymous details include that someone in the group visited the page (but not necessarily me), when the visit happened and for how long.

○ Strongly Agree
○ Agree
○ Slightly Agree
○ Neutral
○ Slightly Disagree
○ Disagree
○ Strongly Disagree

APPENDIX B: IMPLEMENTATION DETAILS

The Search Dashboard

The Search Dashboard (described in Chapter 4) was implemented as a .NET web application. The server-side code was written in C#, and the client-side application was implemented in HTML, CSS, and JavaScript. The client-side made heavy use of the jQuery JavaScript library, and the jQuery plugin jQuery UI. The system ran on Windows Server 2008, and served the application using the Microsoft Internet Information Service (IIS) webserver. All visual displays were custom built or made use of the templates provided from the above systems, with the exception of the gauge charts, which were created using the Google Chart Tools API.

Data for the application was stored in a Microsoft SQL Server (MSSQL) database. The database stored the created profiles of experts and typical users, and the logged profiles of system users. Profiles were created from opt-in logs web logs at Microsoft, and were processed on a virtual cluster using a proprietary map-reduce data processing language. Usage of the Search Dashboard was logged through interface events. The system used an intranet security token to authenticate and identify individual study participants.

WebWear

WebWear (described in Chapters 5 and 6) was built as an extension for the Google Chrome web browser. The extension was implemented in HTML, CSS, and JavaScript. The system interface made heavy use of the jQuery JavaScript library, and the plugin jQuery UI.

The extension settings and preferences were stored locally in the browser using local storage. All visitation history was stored in a MongoDB database. The client Chrome extension interacted with the database through a RESTful Python-based server, SleepyMongoose. This approach allowed no middleware software to be created; rather, the extension sent and received visitation data using AJAX request to the REST interface.

The server was an Ubuntu 10.04 virtual machine, with 1024 MBs of RAM.

SaskWatch Diary System

The SashWatch Diary system (described in Chapter 6) was built as an extension for the Google Chrome web browser. The extension made was implemented in HTML, CSS, and

JavaScript. The client-side made heavy use of the jQuery JavaScript library, and the plugin jQuery UI.

The Diary system requested the participant to complete a survey every 15 minutes, but only when a participant's web browser had focus. If the browser did not have focus, the system would check for focus two minutes later. After completing a survey, the system would wait 25 minutes before attempting to make another survey request.