An Examination of User Behaviour during Web Information Tasks

by

Melanie Kellar

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

at

Dalhousie University
Halifax, Nova Scotia
March 2007

Your file Votre référence
ISBN: 978-0-494-27180-3
Our file Notre référence
ISBN: 978-0-494-27180-3

DALHOUSIE UNIVERSITY

To comply with the Canadian Privacy Act the National Library of Canada has requested that the following pages be removed from this copy of the thesis:

Preliminary Pages
- Examiners Signature Page (pii)
- Dalhousie Library Copyright Agreement (piii)

Appendices
- Copyright Releases (if applicable)

Dedication

To my parents,

for preparing me for this journey

And to my husband,

for holding my hand during it

Table of Contents

List of Tables

List of Figures

Abstract

Since the inception of electronic environments, researchers have been interested in how to provide better support for the tasks users perform in these environments. The research presented in this thesis is the result of three successive studies conducted to examine user behaviour within the web browser in the context of task. An exploratory field study was first conducted to examine how users interact with their web browsers during information seeking tasks on the Web. Based on the study findings a characterization of web information tasks was developed, which includes: Fact Finding, Information Gathering, Browsing, Communications, Transactions, and Maintenance. The study also found significant differences in how users interacted with their web browsers to complete these tasks. The findings from the field study also highlighted the fact that little is known about the monitoring activities of web users, which occur when users return to previously visited web pages to view new or updated information.

As a next step, a series of semi-structured interviews were conducted to concentrate on the role of web-based monitoring in the context of web information tasks. The results from this study suggested that monitoring is an activity that occurs, to varying degrees, within all web information tasks. This implies that information monitoring activities require different types of web browser support, depending on the underlying web information task. Based on the study results, a series of recommendations for the design of task-specific tools to support web-based monitoring were developed. A laboratory study was then conducted to evaluate three task-specific web browser monitoring tools, which were developed based on the recommendations resulting from the semi-structured interviews. The results of this third study reinforced the notion that different monitoring activities require different types of support and also yielded several potential improvements to the tools.

The findings from these three studies provide new understanding of (1) the tasks users engage in on the Web; (2) how users interact with their web browsers to complete these tasks; and (3) how web browsers can better support users during these tasks.

List of Abbreviations and Symbols Used

ACM:	Association for Computing Machinery
Ajax:	Asynchronous JavaScript and XML
ANOVA:	Analysis of Variance
BHO:	Browser Helper Object
BR:	Browsing
CS:	Computer Science
CO:	Communications
DLL:	Dynamic Link Library
IE:	Internet Explorer
FF:	Fact Finding
IG:	Information Gathering
HCI:	Human-Computer Interaction
MA:	Maintenance
NTS:	New Task Session
OT:	Other
PIM:	Personal Information Management
RSS:	Real Simple Syndication
TR:	Transactions

SD: Standard Deviation

URL: Uniform Resource Locator

WTS: Within Task Session

WWW: World Wide Web

Acknowledgements

I would first like to acknowledge the funding sources that supported me during my Ph.D., which include the Natural Sciences and Engineering Council (NSERC) of Canada, Dalhousie University, CANARIE Inc., and the Social Sciences and Humanities Research Council (SSHRC) of Canada.

My supervisor, Dr. Carolyn Watters, and my committee members, Dr. Michael Shepherd and Dr. Kori Inkpen, have all had a profound impact on my graduate school experience - and I really believe there is no dollar amount I can attach to that ;) Together, they made sure I was always cared for intellectually, emotionally, and financially.

Dr. Watters has been a tremendous mentor throughout my Ph.D. and no matter how busy she was, she always made me a priority. During our innumerable discussions and meetings, she inspired and challenged me. She taught me how to step back and look at the big picture and also to never underestimate the value of a pie chart. I initially became interested in pursuing graduate studies during Dr. Shepherd's undergraduate Information Retrieval course and began my research under his direction. When my research interests began to diverge from his, Dr. Shepherd unselfishly encouraged me to follow my own research path and remained my number one cheerleader. Dr. Inkpen has been a wonderful mentor and friend and I would like to thank her for welcoming me into her research lab. Her background in HCI had a significant impact on my research I have truly valued her perspective and optimism. I will definitely miss our late night paper writing sessions! Finally, I would also like to thank my external examiner, Dr. Gary Marchionini, for his insightful comments and feedback on my work.

While there are too many to mention personally, I would like to thank everyone who over the course of my Ph.D. provided me with comments and criticisms of my research. Whether it was through the review of one of my papers, or during a brief conversation over a beer at a conference, your valuable feedback helped shaped my research.

I cannot imagine how different my Ph.D. experience would have been without the support and friendship I received from the members of the EDGE and WIFL research labs at

Dalhousie University. Thank you for piloting, brainstorming, editing, and generally putting up with me.

While there have been several teachers who have had a profound impact on my education, it was under the guidance of my high school math teacher, Mr. Roddie Duguay that I realized I "could do math". Without his encouragement and talent as a math teacher, I would never have pursued an undergraduate degree in Math and Computer Science.

I would like to thank my friends and family for not giving up on me when I failed to return their phone calls and for understanding why it sometimes took me a week to catch up on the latest episode of the Amazing Race. I am grateful to my sister, Natalie, for her constant encouragement and for being there when I became a little unhinged (I have also forgiven her for finishing graduate school before me!).

My husband, Donnie Kellar, was an endless source of support over the course of my Ph.D. Thank you for tolerating my late nights and long absences, especially towards the end when I was unable to engage in a conversation that did not revolve around my thesis.

Most of my success in life can be attributed to my parents, Anne and Joe Fitzgerald. Thank you for nurturing me and encouraging me. From an early age you instilled in me the value of determination and hard work, without which I am sure I would not have been able to do this.

Finally, I would like to thank all of the participants who took part in my studies.

Chapter 1

Introduction

Since the early 1990's, the World Wide Web (WWW) has drastically changed how people acquire and circulate information. Now, people often conduct their research using search engines and online library portals; read the news and their favourite comics online; and communicate with friends and family through email, blogs, and social networks.

Several recent surveys illustrate the wide range of activities currently supported by the Web. The Web serves as an important source of news and up-to-date information for millions of people, with an estimated 60% of web users reading news and sports information online (Statistics Canada, 2006). The Web is also an important source of up-to-the- minute information during world events, such natural disasters. For example, in the aftermath of Hurricanes Katrina and Rita in 2005, it was estimated that 50% of all web users in the United States went online to find information or news about the hurricanes and the aftermath (Horrigan & Morris, 2005). The Web also serves as a rich information source for people conducting research on a variety of topics. For instance, it was estimated that 87% of web users have used the Web to conduct research on a scientific topic (Horrigan, 2006) and 39% have used the Web to look for housing information, such as real-estate listings (Fallows, 2006). It was also estimated that over half of web users have searched the Web for health related information (Statistics Canada, 2006).

The Web is also proving to be an important medium for communications and transactions. Email has been reported to be the most common web activity and it was estimated that 91%

of web users use the Web for email (Statistics Canada, 2006). Blogs and social networking sites are also becoming a popular means of communication. In 2006, it was estimated that 55% of American teenagers who use the Web access social networking web sites, such as MySpace and Facebook (Lenhart & Madden, 2007). The Web also supports an increasing number of financial transactions, such as banking, bill payments, and shopping. Statistics Canada (2006) estimated that 55% of web users use the Web to pay their bills, 58% use the Web for online banking, and 43% have purchased a product or service online.

With the wide range of tasks supported by the Web, it is therefore natural that researchers would want to study and categorize the behaviour of people who use the Web. Byström and Hansen (2005) state that the concept of task is an important area of study for two reasons. First, it "is important in gaining an understanding of why people seek information, the type of information they seek, the methods they chose to acquire it, and the use they make of it" and second, it "provides a framework for analyzing and developing information access in general and for analyzing and designing information retrieval systems in particular".

One gap that we identified in the literature is a lack of research categorizing the high level tasks and activities of web users. While previous research has examined the tasks in which users engage on the Web (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002), not all task models are in agreement and some have only studied a subset of the tasks on the Web. Additionally, the last study of this kind of was published in 2002 and the Web has changed substantially since this time. Blogs, wikis, and social networking sites have exploded in popularity; search engine use has increased significantly; and new web technologies such as Ajax[1] are enabling a new generation of web applications (e.g., Google Maps, Flickr).

There is also a lack of research examining how users interact with their Web browsers within the *context* of task. For example, there is a large body of research examining navigation patterns on the Web (Catledge & Pitkow, 1995; Tauscher & Greenberg, 1997; Weinreich, Obendorf, Herder, & Mayer, 2006) and the use of web browser navigation mechanisms (Cockburn, McKenzie, & JasonSmith, 2002; Milic-Frayling, Sommerer, & Rodden, 2003).

---

[1] Asynchronous JavaScript and XML

However, these studies are typically conducted in a field setting without any context of task, or in a laboratory setting for a focused set of tasks. While revisitation patterns on the Web have been studied extensively in field settings (Cockburn & McKenzie, 2001; Herder, 2005; Tauscher & Greenberg, 1997), researchers have not studied how task may impact revisitation. The steps a user takes when revisiting a web site may be influenced by their task and intentions. For instance, a user who revisits a web page in order to re-find a previously found fact may require different navigational support than a user who is revisiting a web page in order to monitor new information.

There are several factors that have contributed to this lack of research, including the dynamic nature of the Web, methodological challenges in studying user behaviour on the Web, and the characteristics of web users. First, the Web and its use are moving targets that are continually evolving. The dynamic nature of the Web means that users are engaging in a wide variety of ever changing tasks and activities. Therefore, the research literature must be continually updated to reflect these changes.

Second, collecting rich and detailed user data on the Web can be very difficult from a methodological standpoint. While research conducted in the field provides a more realistic picture of users' natural behaviour on the Web, it is difficult to carry out because there is a lack of appropriate tools for collecting both contextual information (e.g., task) and detailed web browser interactions (Fenstermacher & Ginsburg, 2003; Hawkey & Inkpen, 2005b). While laboratory research is somewhat easier to conduct from a methodological standpoint, there is a lack of realism that can impede users' natural behaviour as the tasks are often contrived and users do not typically have access to their usual web browsers and tools.

Third, user behaviour on the Web can be influenced by individual user characteristics. Cognitive differences (Ford, Wilson, Foster, Ellis, & Spink, 2002), domain knowledge, and web experience (Hölscher & Strube, 2000) have all been found to impact information seeking behaviour. There is also evidence to suggest that users' search behaviour may differ between home and work environments (Rieh, 2003). These factors make it difficult to compare new results to previous studies and to generalize results to communities of web users.

1.1 Research Summary

The goals of this thesis were to develop new understandings of (1) the types of tasks users engage on the Web; (2) how users interact with their web browsers to complete these tasks; and (3) how web browsers can better support users during these tasks. A series of three studies, which are summarized below, were designed to meet these goals.

1.1.1 A Field Study Exploring Information Seeking Tasks on the Web

As a first step, a week long field study was conducted to examine the characteristics of information seeking tasks as well as the differences in how users interact with their web browsers to complete these tasks. Over the course of the study, participants were asked to annotate their web usage with task information and to use a custom web browser that logged most of their interactions within the browser. Analysis of the study results found several distinguishing characteristics within each task type and also provided recommendations for how web browsers may better support these tasks. The results of the field study were also used to develop a classification of the information goals and tasks in which users engage on the Web. Web information tasks consist of Fact Finding, Information Gathering, Browsing, Transactions, Communications, and Maintenance. Finally, the study highlighted the fact that little is known about the types of monitoring activities users engage in on the Web, their behaviour during these activities, and whether the design of current tools effectively supports web-based monitoring.

1.1.2 Semi-structured Interviews Examining Web-based Monitoring

The findings from the field study indicated that web-based monitoring warranted further study. As a next step, a series of semi-structured interviews were conducted to concentrate on the role of monitoring in the context of the web information tasks identified in the field study. The results from the semi-structured interviews indicated that monitoring is an activity that occurs within different web information tasks and that different monitoring activities require different types of support. The findings from the interviews were used to develop a series of general and task-specific design recommendations to better support web-based monitoring.

1.1.3 A Laboratory Evaluation of Task-specific Monitoring Tools

As a final step, three prototype task-specific tools were developed to support Browsing, Fact Finding, and Maintenance monitoring activities. The design of these tools was guided by the findings from the semi-structured interviews. A laboratory study was conducted to evaluate the appropriateness of these tools for specific tasks and the overall usability and utility of the tools. The results of this study reinforced the notion that different monitoring activities require different support and also yielded several potential improvements to the tools.

1.2 Organizational Overview

Chapter 2 provides an overview of the literature relevant to the research presented in this thesis. The chapter begins with a summary of user centred models of information seeking. This is followed by an overview of web-based information seeking tasks and strategies, as well as tools that support information seeking on the Web. The chapter then presents research examining revisitation on the Web, followed by a summary of web-based monitoring and monitoring tools research.

Chapter 3 provides an overview of the methodological challenges involved in studying user behaviour on the Web. The chapter provides an overview of the research strategies and data collection tools that have previously been used to study different aspects of user behaviour on the Web. This chapter also discusses many of the tradeoffs inherent to each research strategy.

Chapter 4 presents an overview of the methodology used to conduct the field study examining information seeking tasks on the Web, which forms the basis of this thesis research. This chapter presents the research questions motivating the field study, the methodologies used, and the data collection tools employed. Results from the field study are presented in the two subsequent chapters.

Chapter 5 describes the first set of analyses conducted using data collected during the field study. These analyses explore the characteristics of web information tasks, as well as how user interactions with the web browser differed across tasks. The results of these analyses were used to develop the Web Information Classification, which describes the high level

tasks and goals in which users engage on the Web. As well, a series of recommendations for improving the support of information seeking on the Web are presented.

Chapter 6 describes the second set of analyses conducted using the data collected during the field study. These analyses provide a novel exploration of the impact of task session, task type, and individual differences on the use of web browser navigation mechanisms. The results of these analyses were used to develop recommendations for the future design and evaluation of web browser navigation mechanisms.

Chapter 7 introduces the second study contributing to the thesis research, which explored web-based monitoring activities. This chapter describes a series of semi-structured interviews, which were designed to better understand the role of monitoring within web information tasks. The results of the interviews are used to develop both general and task-specific recommendations for the design of tools to support web-based monitoring.

Chapter 8 describes the development and evaluation of three task specific monitoring tools. These tools were developed based on the recommendations provided in Chapter 7.

Chapter 9 offers a summary of the research presented in this thesis, followed by a description of the major research contributions. The chapter concludes with an overview of the planned future work.

1.3 Previously Published Work

Much of the research presented in this thesis has been published in part in peer reviewed conference proceedings and journals. Many of the methodological choices and experiences, discussed in Chapter 3, Chapter 4 and Chapter 5, appear in (Kellar, Hawkey, Inkpen, & Watters, In Press). In Chapter 5, the characterization of web information tasks appears in (Kellar, Watters, & Shepherd, In Press) and the development of the Web Information Classification is described in (Kellar, Watters, & Shepherd, 2006a). In Chapter 6, the examination of task and the use of web browser navigation mechanisms appears in (Kellar, Watters, & Shepherd, 2006b). In Chapter 7, the examination of the role of task in web monitoring appears in (Kellar, Watters, & Inkpen, 2007). Finally, the evaluation of three task-specific monitoring tools in Chapter 8 is currently in preparation for journal submission.

Chapter 2

Related Work

This chapter describes the literature relevant to the research presented in this thesis. The first half of this chapter provides an overview of information seeking on the Web while the second half focuses on the revisitation and monitoring behaviours of web users.

2.1 Information Seeking

In Chapter 4 to Chapter 6, we examine information seeking tasks on the Web. In this section, we provide an overview of information seeking models and information seeking tasks, followed by the strategies and tools used in information seeking on the Web.

2.1.1 Models of Information Seeking

Marchionini (1995) describes information seeking as "a process in which humans purposefully engage to change their state of knowledge". A large number of information seeking models and theories exist that attempt to understand how users satisfy their information needs, in both electronic and non-electronic environments. Much of users' activity on the Web is related to a need for information; therefore, models of information seeking behaviour have been used to classify users' web-based tasks. We present a subset of the common user centred models of information seeking published in the Information Science literature.

Ellis (1989) developed a behavioural model of information seeking through the study of the information seeking habits of social scientists. The model consists of six activities: starting, chaining, browsing, differentiating, monitoring, and extracting. Two more activities, verifying and ending, were later added by Ellis, Cox, and Hall (1993). Meho and Tibbo (2003) further extended Ellis' model to support web-based information seeking through three additional activities: accessing, networking, and information managing.

Kuhlthau's (1991) model of the information search process is in some ways similar to Ellis' (1989) model of information seeking, but also incorporates feelings, thoughts, and actions. Based on a series of five studies conducted with actual library users, Kuhlthau's model consists of six stages of information seeking: initiation, selection, exploration, formulation, collection, and presentation.

Marchionini's (1995) model of information seeking describes a series of subprocesses and was developed in the context of electronic environments. The subprocesses consist of recognizing the information problem, understanding the problem, choosing a search system, formulating a query, executing the search, examining the results, extracting the relevant information, and the decision to stop/reflect/iterate upon the search process. Unlike Ellis' and Kuhlthau's models of information seeking, Marchionini's model represents a process instead of a set of independent behaviours (Kalbach, 2000).

Wilson and Walsh's (2002) model of information behaviour differs from many of the previous models by suggesting more high-level information seeking search processes: passive attention, passive search, active search, and ongoing search. Passive attention occurs when information is obtained without being actively sought such as while listening to the radio or television. Passive search is the serendipitous acquisition of information through search. Active search occurs when information is actively being sought through explicit searches. Ongoing search occurs when occasional searching is carried out to expand or update previously found information.

While the previously presented models of information seeking have been very useful in understanding information seeking behaviour, they cannot be used to characterize *all* tasks in which users engage on the Web. Information seeking on the Web is a newer branch of

research and differs from library based information seeking in the complexity of the resources and the tools used. Citing Marchionini (1995), Cothey (2002) notes that, "There is little underlying theory of web information searching as distinct from information search theory more generally and especially information searching in electronic environments."

2.1.2 Information Seeking Tasks

The first study of user behaviour on the Web was conducted by Catledge and Pitkow (1995). A three week long study was conducted where participants' behaviour was logged using a modified version of XMosaic. The browsing strategies of participants were classified into three categories: serendipitous, general purpose, and searcher.

Another study examining general user behaviour on the Web was conducted by Pitkow and Kehoe (1996) . They reported five main uses of the Web from the Fourth GVU WWW Survey (2001): browsing, entertainment, work, shopping, and other uses. They noted that the activities had remained fairly consistent since the Second GVU WWW Survey conducted in 1994.

While these studies provide insight into the behaviours of web users, they do not address the types of tasks in which users engage on the Web. Only a few in-depth studies have examined overall information seeking behaviour on the Web in relation to the user's intent or task. One of the most comprehensive studies was conducted by Choo, Detlor, and Turnbull (2000). They studied critical incidents[2] of information seeking on the Web among 34 knowledge workers. Using interviews, questionnaires, and data logging over a two week period, significant episodes of information seeking were coded into one of four modes of information seeking: undirected viewing, conditioned viewing, informal search, and formal search.

Morrison, Pirolli, and Card (2001) used a modified version of Choo et al.'s (2000) critical incident approach by studying significant web activities through 2188 responses to the Tenth

---

[2] Critical incidents are defined as situations or incidents, that are recent and complete and where the effect and consequences of the outcome are known. Flanagan, J. C. (1954). The Critical Incident Technique. *Psychological Bulletin*, 51(4): 327-358.

GVU WWW Survey (2001). Participants were asked to describe a recent episode in which they found information on the Web that led to a significant decision or action. The participants reported four main goals: collect, find, explore, and monitor.

Sellen, Murphy, and Shaw (2002) studied the web activities of 24 knowledge workers over two days. Participants were interviewed in front of their of web history at the end of the second day and described the different activities in which they engaged. Activities were classified into six main categories: finding, information gathering, browsing, transacting, communicating, and housekeeping.

Finally, Rozanski, Bollman, and Lipman (2001) analyzed the clickstream data of 2,466 users and reported seven main web usage occasions: quickies, just the facts, single mission, do it again, loitering, information please, and surfing. This work was conducted from a commercial standpoint since the focus of their work was for marketing purposes.

Table 1. Common categories of user behaviour found in previous research.

	Choo et al. (2000)	**Morrison et al. (2001)**	**Sellen et al. (2002)**	**Rozanski et al. (2002)**
1	Informal search	Find	Finding	Just The Facts/ Quickies
2	Formal Search	Collect	Information Gathering	Information Please/ Single Mission
3	Undirected Viewing	Explore	Browsing	Surfing/ Loitering
4	Conditioned Viewing	Monitoring	N/A	Do It Again
5	N/A	N/A	Transacting/ Communicating/ Housekeeping	N/A

Although these studies differed in methodology and research goals, there are similarities among the resultant categorizations, shown in Table 1. The first is the short answer or informal search, including fact finding and simple lookup. In this category the goal of the

user is to retrieve some short, specific information, possibly on one page. The second category, the formal search, is the more traditional bibliographic search in which the user's goal is to collect information on a topic. This may require multiple pages and overlapping data for confirmation or alternate views on the topic. The third category is the ludic notion of browsing, where the user is engaged in spontaneous information seeking. The fourth category is monitoring, which includes repeated visits to one or more web pages to monitor or check for dynamic information. As can be seen in Table 1 monitoring is not always included as a distinct information seeking task. The final category consists of the remaining web tasks studied by Sellen et al. (2002) which consist of non information seeking tasks such as transacting (e.g., online transactions), communicating (e.g., chat rooms and discussion boards), and housekeeping (e.g., maintaining web pages).

2.1.3 Information Seeking Goals and Strategies

There is a large body of research exploring more focused aspects of information seeking, such as categorizations of search engine queries and the search strategies employed by users on the Web. This area of research provides a better understanding of search behaviour on the Web and provides some insight into improving support for users engaging in web-based information seeking tasks.

Broder's (2002) web search taxonomy categorizes search strategies into three categories: navigational, where the user's goal is to reach a specific web site, informational, where the user's goal is to find information thought to exist on some web page, and transactional, where the user's goal is to perform a web-based activity. Broder concludes that although each type of strategy is motivated by different goals, search engines must be able to support all strategies.

Rose and Levinson (2004) extended Broder's taxonomy to create a search goal hierarchy, which was used to manually classify a set of AltaVista queries. The hierarchy consists of three top level search goals: navigational, informational, and resource. A navigational search goal is to reach a specific web site. An informational search goal is to learn something about a particular topic on the Web, and includes five search sub-goals: directed, undirected, advice, locate, and list. Directed search goals can be either open or closed, where open

searches occur when the answer is open-ended and closed searches occur when there is a single, unambiguous answer. A resource search goal is to obtain a particular resource on the Web and includes three sub-goals: download, entertainment, and interact. Rose and Levinson reported that only 35% of all queries appeared to be of the type traditionally supported by search engines (e.g., directed and undirected search) while over 40% of the queries were non-informational.

Rose (2006) describes three characteristics of users' information seeking behaviour that could be better supported through improved search interfaces. First, search interfaces should support the different search goals (navigational, informational, and resource) undertaken by web users. Second, search queries may have different meanings depending on the cultural or situational context and search interfaces should provide support for these differences. Finally, search interfaces should support the iterative nature of the search process, which could potentially be done through support for query exploration and refinement.

Lee, Liu, and Cho (2005) have further extended Rose and Levinson's work to automatically classify web search goals. They report that using user click behaviour (frequency of user clicks on query results) and anchor-link distribution (relationship between the query text and anchor text[3]), they were able to correctly identify the search goals in 90% of the search queries evaluated.

Jansen, Spink, and Pedersen (2005) categorized 2,600 AltaVista search queries and found that almost 50% of the queries were related to people, places or things. Approximately another 25% of the queries were related to commerce, travel, employment, and technology and the remaining 25% were related to topics such as education, sciences, entertainment, and government. Jansen et al. also reported a high incidence of navigational queries, suggesting that users are increasingly using search engines as a navigation mechanism. For instance, the three most common queries from the 2002 AltaVista data set were "Google", "Yahoo", and "eBay".

Using direct participant observations, previous research has also examined the strategies users employ to conduct web searches. Fidel & Efthimiadis (1999) studied the information

---

[3] The visible text of a hyperlink.

seeking behaviour of engineers through interviews and observations. They reported that although there were common search strategies among the participants (e.g., search queries, rapidly scanning results), they also observed that individuals had developed their own personal search strategies, which has implications for the design of information systems. Teevan, Alvarado, Ackerman and Karger (2004) found two common search strategies among participants: orienteering (approaching the task as a sequence of small steps) and teleporting (jumping directly to the desired information). The orienteering strategy was more common among participants as it allowed them to iterate towards their information goal rather than explicitly state an initial, fully articulated query.

Analysis of search engine logs has also yielded information on user search strategies. Spink, Wolfram, Jansen, and Saracevic (2001) analyzed over one million queries submitted to the Excite search engine in 1997 and found that users employed few search terms, rarely modified their queries, and rarely used advanced search features. Between 1998 and 2002, Jansen, Spink, and Pedersen (2005) observed a decrease in the number of one term queries and an increase in longer queries.

Web users' information seeking strategies may be influenced by their level of experience. In a ten month longitudinal study, Cothey (2002) examined the change in students' information seeking behaviour as they gained more experience over time. Through analysis of web browser history logs, Cothey found that as the students became more experienced they began to visit a more distinct set of web pages, accessed the Web less frequently, and exhibited a lower rate of search queries (relying more on browsing strategies). Cooper's (2001) analysis of transaction log data found that session length within an electronic library catalogue increased over a sixteen month period while the number of searches per session remained constant.

Aula and Käki (2003) and Aula, Jhaveri, and Käki (2005) studied the Web search strategies of expert users. They found that expert searchers often use multiple windows or tabs while searching to support revisitation and to maintain a search history. They also report that expert searchers tended to save links (often using bookmarks) and documents relevant to their search for later revisitation. In terms of search engine functionality, they found that even experienced web users do not used advanced search engine tools, such as Boolean

search. In an attempt to make the strategies of expert searchers more easily available to novice users, they developed Session Highlights, which supports revisitation during Information Gathering tasks (described in more detail in Section 2.1.4).

Much of the research examining users' search strategies on the Web has been conducted in the workplace or in university settings. Rieh (2003) conducted one of the first studies examining web searching behaviour in the home and found that users searched differently than in previous research conducted in the workplace. While in the home, participants searched the Web more frequently, but for shorter periods of time, and the types of searches conducted were much broader.

2.1.4 Tools to Support Information Seeking on the Web

A number of tools have been developed to assist users in their information seeking tasks on the Web. In this section, we provide an overview of systems developed to support search and information gathering activities.

Information Gathering tasks are typically supported by systems that allow users to build and manage collections of relevant web pages. Hunter Gatherer (schraefel, Zhu, Modjeska, Wigdor, & Zhao, 2002) allows users to manage relevant pieces of web page content during Information Gathering tasks. The tool allows users to easily capture and collect relevant content and store it within a collection, where it can be viewed and edited. The results of an initial field study found that while participants did not use Hunter Gatherer to create casual or serendipitous collections, they did use it for formal and planned Information Gathering tasks.

TopicShop (Amento, Terveen, Hunt, & Hix, 2000) is another interface that helps users identify relevant web sites and organize collections of web sites during Information Gathering tasks. The interface offers two web site views: icons or details. TopicShop was evaluated through a pilot study and subsequent user study, where participants were asked to evaluate a set of web sites related to a particular topic. In both studies, participants using TopicShop selected more high-quality web sites in a shorter amount of time than participants using a traditional search engine.

Session Highlights (Jhaveri & Räihä, 2005) is a workspace that enables users to create short-term collection of web pages for revisitation during Information Gathering tasks. A thumbnail representation of visited web pages is displayed in chronological order. A user evaluation of Session Highlights found the tool to be effective for both within-session and between-session revisitation.

Scratchpad (Newfield, Sethi, & Ryall, 1998) is an augmented web browser designed to support faster directed web search. The system attempts to reduce opportunities for serendipitous web browsing and provides faster navigation through two new functionalities: Dogears and breadth-first navigation. Dogears were developed as a cross between bookmarks and the history list and provides an alternative to typical backtracking (which usually happens via hyperlinks and the back button). The reduction in backtracking may potentially reduce off-topic web browsing. The Breadth-first navigation mechanism displays and pre-loads potential future links so that users can load the pages more quickly, thereby potentially improving the speed of navigation.

Based on their study of law librarians and attorneys, Komolodi, Soergel, and Marchionini (2006) designed an interface to support information seeking using search histories. Since people often use external memory aids during complex information seeking tasks, the interface supports the annotation, organization, interpretation, and management of search results. The interface also allows users to record relevance judgments (i.e., relevance of search results to a task) and relate results to a particular task.

Finally, WordBars (Hoeber & Yang, 2006) is a tool that supports query refinement and the exploration of search results. The tool displays the term frequencies of the documents returned by the query, which users can then click on to re-sort the search results. One drawback to query refinement systems is that the system cannot provide benefit if the initial query is bad.

2.2 Revisitation and Monitoring on the Web

In Chapter 7 and Chapter 8 we examine web-based monitoring in the context of task. In this section, we provide an overview of the literature related to revisitation on the Web, web

browser navigation mechanisms (typically used for revisitation), web-based monitoring, and tools to support information monitoring.

2.2.1 Revisitation

While not all revisitation occurs as the result of monitoring, revisitation plays a large role in monitoring activities. There have been four primary examinations of the revisitation behaviour of users on the Web. The results from these studies indicate that revisitation is a very common activity on the Web.

Tauscher and Greenberg (1997) studied the revisitation patterns of 23 participants whose web use was logged over a six weeks period in 1995. A revisitation rate of 58% was reported and the most common reasons provided by the participants for revisiting web pages were: to access new or updated information, to further explore web pages, for special purposes (e.g., search engine), to author a web page, and to use page links to revisit another web page. Tauscher and Greenberg also conducted a second exploration of users' revisitation patterns by analyzing the data collected during Catledge and Pitkow's (1995) study, which revealed a 61% revisitation rate.

Cockburn and McKenzie (2001) characterized revisitation patterns through a retrospective study of 17 users' history backup data over a four month period in 1999. An overall revisitation rate of 81% was reported and individual participant revisitation rates ranged from 61% to 92%. In general, a small number of dominant web pages accounted for most of a participant's revisitation behaviour.

Herder (2005) logged the web usage of 25 participants for varying periods between August 2004 and March 2005 (min: 51 days and max: 104 days). A revisitation rate of 51% was reported, which is much lower than the previous reported studies. Herder attributed this discrepancy to the way in which different researchers have calculated the revisitation rate. For example, Cockburn and McKenzie (2001) removed suffixes from dynamic URLs (e.g., name=value), meaning that repeated visits to a search engine by a user with different queries were recorded as a single URL. After following this truncation method, Herder's revisitation rate increased to 73.7%. The majority of revisitation was characterized as within-session revisitation, often consisting of backtracking. Finally, Herder noted that participants'

revisitation rate stabilized after approximately 1000 page views, which typically occurred after 10 days of logging. This is an important finding for researchers to consider when deciding on the duration of a study exploring revisitation on the Web.

2.2.2 Web Browser Navigation Mechanisms

Since web page revisitation accounts for a large percentage of users' navigation, much of the previous work investigating web browser navigation tools has centered on those tools that support both immediate and post-session revisitation (Aula, Jhaveri, & Käki, 2005). The three most commonly studied web browser navigation tools are those designed to support revisitation: back button navigation, bookmarks, and the history mechanism. New navigation mechanisms are emerging, such as embedded browser toolbars to support search activities (e.g., Google toolbar, Yahoo! toolbar), whose use has not been thoroughly studied.

Despite the many modifications and new features introduced by web browsers, the primary methods of web navigation consist of hyperlinks and the back button. Catledge and Pitkow (1995) reported that the two most commonly used methods of web navigation were hyperlinks (52%) and the back button (41%). The reported usage of other navigation mechanisms such as bookmarks, open URL, the home button, the forward button, and the history list was relatively minimal and each navigation mechanism individually accounted for 0.1% - 2% of all navigation.

Tauscher and Greenberg (1997), while studying the revisitation strategies of users, reported that hyperlinks accounted for 41% of all navigation and the back button accounted for 30%. The overall reported usage of other navigation mechanisms such as bookmarks, typed-in URLs, the home button, the forward button, the reload button, and the history list, was again relatively minimal and each navigation mechanisms individually accounted for 0.8% - 5% of all navigation.

Milic-Frayling, Jones, Rodden, Smyth, Blackwell, and Sommerer (2003) reported that 43% of all web navigation was a result of hyperlinks while the back button accounted for 23% of all navigation. The reported usage of other navigation mechanisms such as bookmarks, typed-in URLs, the home button, the refresh button, and the forward button, was relatively minimal and each navigation mechanism individually accounted for 0% - 3% all navigation.

Weinreich, Obendorf, Herder, and Mayer (2006) observed a significant decrease in back button usage from the earlier studies. While the use of hyperlinks remained fairly constant, accounting for 43% of all navigation actions, the back button only accounted for 14.3%. They also report only 0.2% of navigation was due to the history mechanism.

The previously high usage frequency of the back button motivated several researchers to examine enhanced back button functionality. Moyle and Cockburn (2003) evaluated the use of gesture-based navigation mechanisms for forward and back navigation in web browsers and found that it significantly improved participants' navigation speed. Cockburn, McKenzie, and JasonSmith (2002) examined a recency-based model of back navigation, allowing users access to all previously visited pages through the back button. This is in contrast to the standard back button stack-based navigation model. The results of their evaluation found that the temporal model was very efficient for distant navigation, but performed poorly in navigation leading to a parent page. Milic-Frayling, Sommerer, and Rodden (2003) have also developed enhanced back button functionality. They developed SmartBack, a back button navigation mechanisms that allows users to "jump" back to previously visited web pages.

The reported use of bookmarks in the literature has been relatively low, especially in comparison with back button and hyperlink usage. Much of the bookmarks research has studied the information management issues associated with the collection and use of bookmarks. Abrams, Baecker, and Chignell (1998) explored the organizational habits of bookmark users through a survey questionnaire and an analysis of 50 bookmark archives. Participants in this study identified several problems with bookmarks, including: a lack of support for the organization and management of bookmarks, problems re-finding bookmarks within a current archive, and a lack of naming descriptions. Boardman and Sasse (2004) examined the personal information management (PIM) strategies of users across a number of tools through semi-structured interviews and a longitudinal study. They developed a classification of users' bookmark management strategies, which include extensive filing, partial filing, and no filing. Jones, Dumais, and Bruce (2002) studied the "keeping" behaviours of information professionals through an observational study in the workplace. They found that bookmarks names tended to lack meaning and did not provide a descriptive reminder function.

One example of enhanced bookmark functionality is Landmarks (MacKay, Kellar, & Watters, 2005), which are an extension to traditional bookmarks that allow users to return to a specific place on a web page. A comparison against traditional bookmarks found that participants were able to re-find information significantly faster using Landmarks.

The history function appears to be one of the least commonly used navigation mechanisms (Aula, Jhaveri, & Käki, 2005; Tauscher & Greenberg, 1997; Weinreich, Obendorf, Herder, & Mayer, 2006). One of the leading issues with the history mechanism is how to represent the large number of previously viewed pages in a way that is meaningful to users. Ayers and Stasko (1995) developed a graphical history view, which allowed users to recognize previously visited hypertext documents through titles, URLs, and thumbnail images. In a similar vein, PadPrints (Hightower, Ring, Helfman, Bederson, & Hollan, 1998) is a "browser companion" that builds a graphical representation of visited web pages that users can navigate using a zooming user interface. An evaluation of found that PadPrints reduced access time during revisitation tasks. Kasssten, Greenberg, and Edwards (2002) also examined the use of thumbnail images for recognizing previously found web pages. They found that users were able to recognize web pages from the thumbnail images and that colour and layout were the primary identifiers. In addition to providing thumbnail images of previously viewed web pages, WebScout (Milic-Frayling, Sommerer, & Rodden, 2003) provides users with a representation of previous navigation events as well an archive of previous search queries and assigned labels.

2.2.3 Monitoring

Monitoring has been defined in numerous ways by several researchers and the nature of monitoring has not been extensively studied. Within the Information Science community, Ellis (1989) and Meho and Tibbo (2003) both consider monitoring as an information seeking behaviour, which is defined as maintaining an awareness of research developments through particular sources (not necessarily web-based sources). Both of these studies have focused primarily on the theoretical study of monitoring, exclusively within the context of research and library based information seeking. Choo et al. (2000) characterized conditioned viewing as an information seeking mode that is comprised of browsing, differentiating, and monitoring. They state that "The most important characteristic of conditioned viewing was

that participants regularly or frequently returned to their selected or differentiated sites/pages to check for new information (monitoring)."

Morrison et al. (2001) defined monitoring as routine visits to websites to view updated information that is not motivated by a particular goal. This research seems to suggest that monitoring is synonymous with Browsing, which is a serendipitous information seeking task with no particular goal in mind. Sellen et al. (2002) did not classify monitoring as one of their six web activities, but reported that some searches lead to monitoring and also described some Browsing activities as "checking to see what's new on a hobby-related site". Rozanski et al. (2001) do not explicitly mention monitoring in the category Do it Again; however, they do describe repeated visits to familiar web sites. Within this thesis, we define web-based monitoring as *repeated visits to a web page to view new or updated information.*

2.2.4 Tools to Support Web-based Monitoring

A number of web-based monitoring systems have been developed; however, published evaluations of these systems from an HCI perspective have been limited. WebCQ (Lee, Liu, & Cho, 2005) and WebVigiL (Chakravarthy, Sanka, & Jacob, 2004) are both server based notification systems that allow users to monitor web pages for general or specific changes to a web site. ChangeDector (Boyapati et al., 2002) is a site-level monitoring tool that detects significant changes within an entire web site and not just a single web page. The system allows users to specify the information to monitor (e.g., names, dates on an organization's web site) and can be useful in monitoring corporate web sites for "silent news" (e.g., changes to an organization's management).

In contrast to the previously described notification systems, awareness displays allow users to maintain an awareness of important information. Van Dantzich, Robbins, Horvitz, and Czerwinski (2002) developed Scope, a visualization tool that provides an awareness of information notifications. One of the major strengths of Scope is that it unifies notifications from a variety of sources, including email, instant messages, and information alerts. Notifications are displayed on a circular, radar-like interface, where the position of a notification denotes its urgency.

InfoCanvas (Plaue, Miller, & Stasko, 2004) allows users to peripherally monitor information through information art. Users can create an abstract representation of the monitored information, which are displayed on a flat-paneled LCD screen that can be hung on the wall or placed on a desk. BlueGoo (Plaue & Stasko, 2006) allows users to monitor RSS feeds on peripheral displays. The system delivers an animated collage of images that represent articles in the RSS feed. A laboratory evaluation of the system found that the display was not overly distracting to users engaged in a primary information seeking task. WebTracker (Greenberg & Boyle, 2006) is a custom notification system that allows users to select elements of any web page to monitor, such as text, images, and video. All selected elements are arranged and displayed within a single window.

From a commercial standpoint, Netmind (also known as URL-Minder and Mind-it) was an early web-based monitoring service that notified users when registered web pages were modified. ChangeDetect (2007) is a similar service that emails users when changes to specified web pages are detected. More recently, monitoring tools have been appearing as browser extensions and desktop widgets. Apple Inc. (2007) has introduced Web Clips into the Mac OS X Dashboard, which allows users to monitor snippets of web pages by clipping regions of interest. This tool is similar Greenberg and Boyle's (2006) Web Tracker. Tools such as the Morning Coffee Firefox extension (Liesegang, 2006) offer enhanced bookmark-like functionality to assist users with their habitual monitoring routines. Real Simple Syndication (RSS) is a web feed format that allows users to monitor syndicated content. Feeds can be accessed through a variety of tools, including RSS readers, aggregators, and personalized homepages, such as those offered by Google (2007), Yahoo! (2007), and NetVibes (2007).

2.3 Summary

The research presented in this chapter provides an overview of the related work in the areas of information seeking and web-based monitoring. This literature review has highlighted the gaps in the literature that motivated the research presented in this thesis. For example, Section 2.1.2 demonstrates how researchers differ in their classifications of the tasks performed by web users. Section 2.2.2 presents an overview of the uses of web browser

navigation mechanisms; however, there has been little study of the use of these tools in the context of task. Section 2.2.3 illustrates the lack of research examining web-based monitoring. Finally while a number of tools to support information seeking and web-based monitoring were presented in Sections 2.1.4 and 2.2.4, there have been few published evaluations of these tools.

The next chapter describes many of the research strategies and data collection techniques used in studying user behaviour on the Web. Many of the methodological tradeoffs and considerations discussed in the next chapter have influenced the design of the studies presented in Chapter 4 through to Chapter 8.

Chapter 3

Studying User Behaviour on the Web

Studying user behaviour on the Web is often challenging from a methodological standpoint. In this chapter, we present an overview of research strategies and data collection techniques used in studies of user behaviour on the Web and discuss many of the tradeoffs involved in the choice of methodology. We also describe the definition of sessions use by researchers studying user behaviour on the Web.

3.1 Strategies for Studying User Behaviour on the Web

McGrath (1995) describes four categories of research strategies for conducting research in the social and behavioural sciences: (1) field strategies (field experiments and field studies); (2) experimental strategies (laboratory experiments and experimental simulations); (3) respondent strategies (sample surveys and judgement studies); and (4) theoretical strategies (formal theory and computer simulation). The choice of strategy impacts the *generalizability* of the results, the *precision* of the measurements and conditions being studied, and the *realism* of the scenario in which the data is collected. McGrath states that no single research strategy can maximize all three features; choosing to maximize one strategy comes at the expense of the others and the decision of which strategy to use should be carefully considered. In the field of HCI, and specifically in the study of user behaviour on the Web, the most common research strategies include field strategies, experimental strategies, and respondent strategies.

In this section we provide an overview of the use of these research strategies for studying user behaviour on the Web.

3.1.1 Field Strategies

Field strategies are becoming an increasingly common research strategy for studying user behaviour on the Web. The primary strength of field strategies is the increase in realism as participants are observed within their natural environment, with their own tools (e.g., bookmarks, history, choice of browser), and completing tasks that are motivated by the participant and not the researcher. However, field strategies are typically conducted with a relatively small, homogenous set of participants, which can lessen the generalizability of results. Additionally, due to the natural environment in which field strategies are conducted, data collection can be difficult and researchers must often accept a loss of precision and control.

McGrath characterizes field studies as purely observational studies, with as little intervention on the part of the researchers as possible. Field studies have been successfully conducted by a number of researchers studying user behaviour on the Web (Catledge & Pitkow, 1995; Tauscher & Greenberg, 1997; Weinreich, Obendorf, Herder, & Mayer, 2006), which is typically conducted through unobtrusive web logging. However, the purely unobtrusive nature of these studies means that researchers cannot collect valuable contextual information. In field experiments, researchers may introduce a change or a new feature into the environment, thereby sacrificing some realism for improved precision. Examples of field experiments include the studies conducted by Hawkey and Inkpen (2006) and Byrne, John, Wehrle, and Crow (1999). In these cases, participants were asked to provide feedback during the study, which they would not have done otherwise. In the case of Hawkey and Inkpen's study, participants were asked to annotate their web usage at the end of the day, while Byrne's et al's participants were asked to use a talk aloud protocol as they navigated the Web in their natural environment. For simplicity, the term "field study" is used for both field studies and field experiments during the remainder of this thesis.

3.1.2 Experimental Strategies

When using experimental strategies, researchers can mandate the tasks in which a user engages and the software is standardized across all participants. Data is often collected using video/screen capture, direct observations, and logging methods, such as transaction logs. Although video recordings are easy to capture, the subsequent data coding that needs to take place can be very time consuming.

While experimental strategies are often much easier to conduct than field strategies, one major drawback to experimental strategies is that they are often lacking in realism. Typically, participants are asked to complete tasks under time constraints and on lab computers, without their usual web resources (e.g., bookmarks, web history, and toolbars). One alternative, as used by Schiano, Stone, and Bectarte (2001), is to invite participants to perform a task they already need to do. Web-based monitoring can be especially problematic to evaluate in a laboratory setting since monitoring is usually a secondary task that takes place in the background. A common solution used by researchers (McCrickard, Chewar, Somervell, & Ndiwalana, 2003; Plaue & Stasko, 2006) is to simulate a multitasking environment by having participants engage in a primary task, such as document editing or web search, while the monitoring activity become a secondary task.

3.1.3 Respondent Strategies

Respondent strategies are often used to study user behaviour on the Web because they can be administered to large and diverse population and can produce data with a high degree of generalizability. Unlike other research strategies, this method of data collection is relatively uncomplicated. The downside of these strategies is that participants are studied outside of the context of their information seeking, which can decrease the level of realism. As noted by Sellen et al. (2002), the way in which questions are asked can bias the results towards certain types of events. Teevan et al. (2004) noted that simple semantics such as the difference between asking participants what they were "looking for" versus "searching for" may influence what participants report. This is true, however, across all research strategies.

3.2 Data Collection Techniques for Studying User Behaviour on the Web

In this section we provide an overview of the data collection techniques often used when studying user behaviour on the web. This includes data collected client-side, server-side, or via proxy, as well as contextual user information.

3.2.1 Logging

Client-side logging takes place on the user's own computer. Researchers have used a number of tools to collect client-side logs, including commercial "spy-ware" tools (Kelly & Belkin, 2004; Kim & Allen, 2002), custom built logging tools (Obendorf, Weinreich, & Hass, 2004; Reeder, Pirolli, & Card, 2001; Turnbull, 1998) and custom web browsers (Claypool, Le, Waseda, & Brown, 2001). Client-side logging has been used to capture a variety of behaviours and to study a wide-range of research areas, including information seeking tasks (Choo, Detlor, & Turnbull, 2000), navigation behaviour (Catledge & Pitkow, 1995), and incidental information privacy (Hawkey & Inkpen, 2006). Client-side logging offers the richest exploration of user interactions with the web browser. However, many client-side logging tools are designed to work with a specific browser and may be time consuming and costly to update as new versions of the browser are introduced. There may also be performance issues with a custom web browser due to lack of robustness.

Server-side logs do not capture the same level of detail as client side-logs; however, there is a lower cost of implementation. One benefit of server-side logging is that researchers can study a large sample population. This method has been successfully used to study search engine transactions logs (Jansen, Spink, & Pedersen, 2005; Spink, Wolfram, Jansen, & Saracevic, 2001) and library portals (Zhang, Zambrowicz, Zhou, & Roderer, 2004). The data recorded typically includes the IP address of users and the time and address of web page requests.

With server-side logging, researchers have very little information about the participants being studied; users are typically anonymous. Jansen, Spink, Bateman, and Saracevic (2000) analyzed Excite search engine logs and acknowledged that while the data reflects real search behaviour, they "...report on artifactual behaviour, but without a context". Zhang,

Zambrowicz, Zhou, and Roderer (2004), who logged users of the library portal MyWelch, stated that additional studies in the form of interviews or surveys were needed in order to better understand users' behaviour from a "mental or cognitive" perspective as well to get a sense of other applications and multitasking. Weinreich et al. (2006) discuss their use of clickstream logs, which they state "...have limited expressiveness, as aims and tasks of the users often stay below the surface. This makes their contextual interpretation inherently difficult and additional qualitative information is needed to support a detailed task-related evaluation of the data."

Logging conducted through a proxy sever is a compromise between client-side and server-side logging. Proxy servers act as an intermediary between a user's web browser and a web server and can log interactions between the client and the server. By allowing participants to login to the proxy server instead of downloading and installing software, proxy solutions such as WebQuilt (Hong, Heer, Waterson, & Landay, 2001) allow participants to work within their normal browsing environment. However, proxy server logging does not capture the full spectrum of user interactions with the browser and may not capture access to pages that have been cached at the browser level (Barford, Bestavros, Bradley, & Crovella, 1999). Proxy sever logging may also be problematic when trying to collect fine-grained measurements. Kelly and Belkin (2004) found discrepancies between a client-side logging tool and a proxy-based logging tool while collecting web page dwell times; the timing data generated by the proxy-based logger was found to be inaccurate when compared to the client-side logging tool.

3.2.2 Contextual Information

A further consideration when conducting field research is how to capture contextual information, such as the user's setting, intentions, goals, and tasks. Contextual information can be gathered at a high level through interviews and surveys (Aula, Jhaveri, & Käki, 2005; Jones, Dumais, & Bruce, 2002; Rieh, 2003; Teevan, Alvarado, Ackerman, & Karger, 2004). Several researchers have used client-side logging methods supplemented with interviews after the fact to gather additional contextual information. For instance, Choo et al. (2000) conducted interviews after analyzing their web usage logs. The logs were used to guide the discussions with participants about the tasks they were performing. Sellen et al. (2002)

interviewed participants about their previous two days' web activities, while participants were seated in front of their browser history.

Other researchers have asked participants to annotate their web usage with contextual information. For example, Kelly and Belkin (2004) logged participants' web usage in a 14 week study of display time as a measure of user interest across information seeking tasks. One a week, participants were asked to annotate their web usage with task information and to indicate the usefulness of viewed web pages. Over the course of a week long field study investigating the privacy of incidental information on the Web, Hawkey and Inkpen (2005a; 2006) asked participants to annotate their web usage with privacy comfort levels using an electronic diary.

3.3 Sessions

In the study of user behaviour on the Web, sessions are often used to create delimited portions of user activity for further analysis. A session is generally defined as a period of continuous web use with no break in usage exceeding a specified threshold. However, the specific definition of a session tends to vary across researchers and research disciplines. In this section, we provide an overview of previous researchers' approach to session length.

Using client-side transaction logs, Catledge and Pitkow (1995) defined a session as period of continuous web use with no break is usage greater than 25.5 minutes. They found that participants averaged 9.4 sessions over the three week long study. In a later study using client-side logs of 20,000 web users, Montgomery and Faloutsos (2001) defined a session as a period of continuous web use, which begins when a user has not accessed the Web in the previous two hours. Using data collected during the month of December in 1999, they reported a median of 4 web sessions per month, per user. Based on their observed growth of web use between 1997 and 1999, Montgomery and Faloutsos also expected the number of sessions per month to double over the following four years.

Grace-Martin and Gay (2001) collected client-side logs when examining laptop use among university students. They used a delimiter of 10 minutes in determining session length and reported an average of 3 sessions per day. Hawkey and Inkpen (2005b) used both a 10 and

30 delimiter for their client-side logging data, collected primarily from university students as well. With the 10 minute delimiter, they reported an average of 9.4 sessions per day and 5.4 sessions a day with a 30 minute delimiter.

In their studies of search engine transaction logs, Jansen and Spink (2003) measured session duration as the time from when a user submitted their first query to the search engine until the user navigates away from the search engine and does not return. They reported that 52% of all sessions lasted less than 15 minutes and 26% of all sessions lasted less than five minutes. Anick (2003) analyzed search engine logs and defined a session as a period of web activity that begins with a query and continues until 60 minutes of inactivity. If after 60 minutes a user interacts with a search result, then the session is further extended.

3.4 Summary

In this chapter, we have outlined the research strategies typically used in the study of user behaviour on the Web, which include field strategies, experimental strategies, and respondent strategies. The methodological tradeoffs discussed in this chapter were carefully considered in the design of the studies presented in Chapter 4, Chapter 7, and Chapter 8. As well, we also incorporate the notion of "task sessions" in the field study, based partially on the definition of web sessions presented in this chapter.

Chapter 4

A Field Study to Explore Information Seeking Tasks on the Web

The literature review in Chapter 2 indicated there is a lack of research examining the high level tasks in which users engage on the Web and how users interact with their web browsers within the context of these tasks. As a first step, we designed a week long field study to observe the types of information seeking tasks participants engaged in, as well as differences in how users interact with their web browser during these tasks. This chapter presents the research questions and goals motivating this study, followed by the methodological approaches and data collection techniques used, including the development of a custom web browser. Results are presented in Chapter 5 and Chapter 6.

4.1 Research Questions and Goals

This research, while exploratory in nature, was conducted in order to answer the following research questions:

RQ1: What are the high level information seeking tasks in which users engage on the Web? While previous research (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002) has contributed to an initial understanding of the high level tasks in which users engage on the Web, one goal of this research was to further solidify this

understanding and explore inconsistencies in the previous research. This research question is addressed in Chapter 5

RQ2: What web browser functionalities are currently being used during web-based information seeking tasks? In particular, we were interested in the usage (or lack thereof) of web browser navigation mechanisms (e.g., auto-complete, bookmarks, history), browser functions (e.g., windows and pages loaded, use of copy/cut/paste), and search tools during Fact Finding, Information Gathering, Browsing, and Transaction tasks. We address this research question in Chapter 5 through general usage frequencies for each task and continue this exploration in Chapter 6 through a detailed analysis of the use of web browser navigation mechanisms across tasks.

RQ3: Are there differences in the patterns of use of web browser functionality across web-based information seeking tasks? We were interested in whether there are significant differences in how participants interact with their web browser between Fact Finding, Information Gathering, Browsing, and Transaction tasks. This research question is addressed in Chapter 5 where we examine the differences in the use of web browser functionality across all information seeking tasks. Together with *R2*, the goal of this research was to provide recommendations for how we can better support information seeking behaviour on the Web.

4.2 Methodology

A week long field study was conducted where participants' web usage and web browser interactions were recorded using a custom-built web browser. In order to address our research questions two types of data were collected. First, participants' web use and web browser interactions were logged over the course of the field study using a custom-built web browser. Second, participants used an electronic diary to describe and categorize their web usage according to a defined categorization. In advance of the field study, a pilot study and focus group were conducted to evaluate two electronic diary techniques and refine the task categorization provided to participants. This section outlines the pilot study and focus group, followed by a description of the field study methodology. The research described in this chapter was approved by the Social Sciences and Humanities Human Research Ethics Board at Dalhousie University.

4.2.1 Pilot Study

In preparation for the field study, a four day long pilot study was conducted with six participants all recruited from within our research lab at Dalhousie University. Participants were asked to use a custom web browser for all their web usage during the pilot, which logged all interactions with the browser (including URLs visited). Participants were also asked to categorize their web usage according to the following categories: Fact Finding, Information Gathering, Monitoring, and Browsing. A fifth task of Other was provided for all other tasks that did not fit within the given schema. Participants took part in a fifteen minute training session before beginning the pilot in which they were introduced to the task categorization and each category was carefully explained. Upon completion of the study participants completed a post-session questionnaire which allowed us to explore their experiences with the logging software and the task categorization.

One goal of the pilot study was to determine which of two electronic diary methods allowed participants to more easily and more accurately record task information related to their web usage. The first electronic diary method required users to provide task information in real-time using a toolbar available within the custom web browser. The second method required users to record their task information at the end of each day using a task diary. Participants in the pilot study used the toolbar for half the pilot (two days) and the task diary for the other half. The order in which the participants used the two different methods was counterbalanced. The results of the pilot study found that the participants were equally split on overall preference and ease of use for the two input methods. Most participants (5/6), however, reported that they felt they were more accurate in their task assignments when using the toolbar. Since the participants were equally split on the two techniques in terms of ease of use and overall preference, we decided to provide the study participants with both methods, allowing them to use either as needed.

The second goal of the pilot study was to evaluate how well participants were able to categorize their web usage according to the five task categories: Fact Finding, Information Gathering, Monitoring, Browsing, and Other. These five categories had been chosen based on previous work on information seeking behaviour on the Web (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002). Before starting the

field study, we needed to verify that the categories reflected most of the tasks in which users engage on the Web while at the same time remaining easy to understand and relatively distinct. Overall, participants struggled with the task of Monitoring because it often led to new tasks and was hard to distinguish from "re-Fact Finding" or "re-Browsing". One example given was reading online comics. A participant was unsure whether repeatedly reading the same comic strip was Browsing or Monitoring. Additionally, half of the participants reported that it was difficult to distinguish between Fact Finding and Information Gathering. Participants also used the category "Other" for several types of tasks, most notably email. Based on these results, a focus group (described in the next section) was held to refine the task categorizations. The pilot study was also an opportunity to detect bugs within the data collection tools (e.g., custom web browser, electronic diary) as well as to refine the training materials provided to participants.

4.2.2 Focus Group for Task Refinement

Ten participants from the Faculty of Computer Science at Dalhousie University (students and faculty), none of whom had taken part in the pilot study, took part in an informal focus group. We selected 40 task descriptions from the larger set of task descriptions collected during the pilot study to use during the focus group. Each task description was printed onto an index card and spread out on a large table (as shown in Figure 1). Examples of task descriptions included: *Searching for papers on direct input*, *Looking for the final Superbowl score*, and *Updating my blog*. The participants were asked to work together as a group to organize the

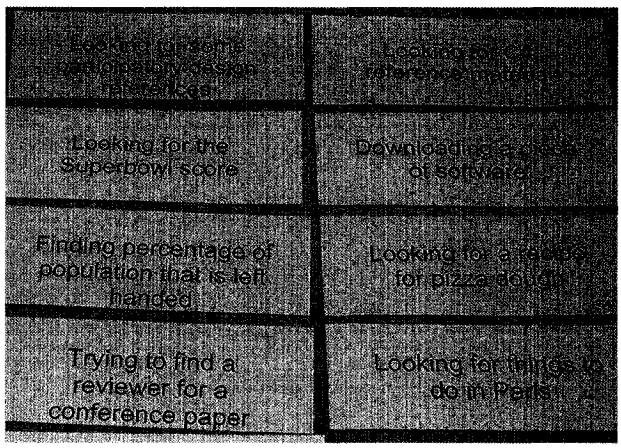

Figure 1. Cue cards, each containing a task description, were used to refine the tasks.

tasks and form a consensus on categories based on the goal of the task. Although some participants had backgrounds in web behaviour and information science research, the focus group participants were not informed of the categories used in the pilot study or in previous literature.

During the hour long session, participants re-arranged the task groupings several times. The content and number of categories fluctuated continually during the course of the session. After much discussion among the participants, the categories began to stabilize and six final categories emerged (shown in Table 2). We labelled the categorizations produced by the focus group participants as: Looking for Specific Information, Passing Time/Entertainment, Transactions/Communication, Information Gathering, Routine/Hobby, and Monitoring.

Table 2. Initial task categories after focus group.

Task	Examples
Looking for Specific Information	Location of a conference workshop Finding the percentage of the population that is left handed
Passing Time/ Entertainment	Random surfing Just browsing EBay
Transactions/ Communication	Checking my email Online banking
Information Gathering	Trying to find a reviewer to review a conference paper Looking for references on a topic
Routine/Hobby	Reading my favourite comic Reading blogs
Monitoring	Checking to see if a project page is up to date so I can send the URL to a colleague Looking up the prices of my stocks

The task categories that evolved out of this focus group were in fact very similar to the tasks reported in the literature. Based on the findings of our pilot, we hypothesized that Monitoring was a subtask of other information seeking tasks rather than a distinct information seeking task. Therefore, we eliminated the Monitoring category as it was clear that further study was needed to define the nature of users' monitoring activities. We merged the categories Passing Time & Entertainment and Routine & Hobby into a single category (Browsing) as it was difficult to clearly articulate the distinction between these two

categories; they are both serendipitous in nature and lack specific goals. The resulting task categories, shown in Figure 2, include: Fact Finding, Browsing, Information Gathering, and Transactions. Typically, Transactions such as email or banking have not been classified as information seeking tasks. However, given the growing proportion of browser activities that these tasks constitute, we felt it was important to study information seeking behaviour in the context of all web usage.

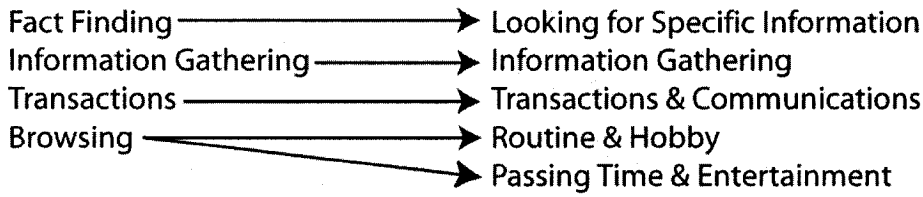

Figure 2. The final task categories.

For the ensuing field study, the following task descriptions were provided to all participants:

Fact Finding: Fact Finding is defined as a task in which you are looking for specific facts or pieces of information. These are usually short lived tasks that are completed over a single session because either you find the answer or you do not. Examples include looking for tomorrow's weather, a pizza dough recipe, or a printer driver for your printer.

Information Gathering: Information Gathering involves the collection of information, often from multiple sources. This type of task can take place over a single day or it may stretch out over several days. Unlike Fact Finding, you do not always know when you have completed the task and there is no one specific answer. Examples of Information Gathering include building a bibliography for a research paper, researching different car models when buying a new car, or planning an upcoming vacation.

Just Browsing: Browsing is defined as a serendipitous task where you may be visiting web pages with no specific goal in mind. You may allow yourself to take part for a pre-determined period of time (e.g., I have 20 minutes before my meeting). This type of task is your classic "web browsing", with no specific goal in mind other than entertainment or to

"see what's new". Sometimes this is done as part of a daily routine. Examples include reading the news, your favourite comic, or a friend's blog.

Transactions: Transactions are defined as tasks in which you are performing an online action. Often, a username/password is associated with the transaction. Examples include web-based email, banking, or posting to a message board.

Other: A final category of Other was provided to participants in the event they encountered tasks during the study in which they were not sure how to categorize or which did not fit within any of the predefined categories. Participants were also instructed to categorize their homepage as Other if they did not use it as part of task, since it loads each time the web browser loads, and these pages were not included in our analysis.

4.2.3 Field Study

4.2.3.1 Participants

Twenty-one students from Dalhousie University took part in a one week field study in March 2005. Although 23 participants were recruited, only data for 21 participants was analyzed. One of the original participants did not finish the study and another participant's data was not usable because the task descriptions were incomplete and inconsistent. All participants were paid $25 for taking part in the study.

Recruitment notices were circulated via email and stated that all university students who were laptop and Microsoft Internet Explorer (IE) users were eligible to participate. Laptop users were targeted because we could capture most of their web usage on a single machine and because it facilitated installation of the custom software. Since the web browser used during the study was a clone of IE, participants were required to be current users of IE.

The academic background of the participants was divided among Computer Science (11/21), Health Informatics (2/21), Business (4/21), Economics (2/21), Kinesiology (1/21), and Arts (1/21). Participants were also from both the graduate and undergraduate communities: Computer Science (7 grad/4 undergrad), Health Informatics (2 grad), Business (4 grad), Economics (2 grad), Kinesiology (1 undergrad), and Arts (1 undergrad). The median age

group category of the participants was 20-29 and the gender was almost evenly split with 11 males and 10 female participants. The median category of web usage reported by the participants was between 30-39 hours of web usage a week. Although Computer Science students are typically considered to be highly technical, all participants who took part in our study were experienced web users, regardless of discipline. All participants were the primary users of their laptops and five participants also reported they used a desktop (either at home or work) for some of their web usage.

4.2.3.2 Procedure

On the first day of the study, each participant met with the researcher administering the study for a one hour session. Participants signed an informed consent form which outlined the procedures in which they would be involved while taking part in the study. The custom web browser and logging tools were then installed on the participant's laptop. The custom web browser was configured with the same settings as the participant used in IE, such as auto-complete, the bookmarks toolbar and the Google toolbar. Both the demographic and web browser navigation inventory questionnaires (described in the next section) were administered at this time. The researcher then carefully described the different information seeking categories and explained how to use both electronic diary methods (i.e., the task toolbar and task diary) to record task information. Participants then took part in a short training exercise in which they were required to complete several short information seeking tasks using both electronic diary methods to categorize their web usage. Finally, participants were given printouts of the task definitions (which were also available online) and instructions for the study tools.

After a one week period, participants returned to meet with the same researcher. At this time, the software was uninstalled from the participant's laptop and all logging data was copied to a backup disk. At this time, participants completed a final post-study questionnaire.

Before we began data analysis, a single researcher manually reviewed all participants' data. We encountered some situations where the task information did not appear to match the URLs recorded. In cases where the behaviour was habitual and obvious, the researcher

changed the task information. In all other cases, the participants were contacted in order to clarify the task information.

4.3 Data Collection

Over the course of the field study, three types of participant data were collected: logged web browser interactions, qualitative task data, and questionnaire data. In this section we describe the data that was collected and the tools used for data collection.

Table 3. A listing of the web browser interactions logged during the field study.

Navigation Events†	Browser Events			
	File	**Edit**	**View**	**Misc. Tools**
Auto-Complete	New	Select All	Toggle Bookmarks	Highlight Search Terms
Back Button	Window	Find	Toggle History	Internet Options
Back Menu	Open	Copy‡	Stop	
Bookmarks	Save As	Paste‡	View Source	
Forward Button	Page	Cut‡	Privacy Report	
Forward Menu	Setup			
Google toolbar	Print			
History	Print			
Home Button	Preview			
Hyperlinks	Properties			
New Window	Close			
Other				
Reload Button				
Select URL				
Typed-in URL				

† Includes navigation conducted through button clicks, shortcut keys, and menu interactions
‡We differentiated between cut, copy, and paste that occurred within the web browser web page and within the web browser combo-boxes (the address field and Google toolbar)

4.3.1 Logged Web Browser Interactions using a Custom Web Browser

We were primarily interested in collecting participants' direct interactions with the web browser interface, such as the use of web browser navigation mechanisms (e.g., auto-complete, bookmarks, history), browser functions (e.g., windows and pages loaded, use of copy/cut/paste), and search tools. Table 3 displays a listing of the web browser interactions logged by the custom web browser. This set of interactions was partially based on Byrne et al.'s (1999) Taskonomy of WWW tasks, which was developed to better understand the range

of tasks that web browsers must support. Other classifications of web browser interactions include Oard and Kim's (2001) classification of observable behaviours, which was further extended by Kelly and Teevan (2003) and Jansen and McNeese (2005), is based on two dimensions: behaviour category (i.e., purpose of the category) and minimum scope of the object being manipulated. This classification, however, was developed in the context of implicit measures and has a larger focus on documents and content, as opposed to web browser interface interactions.

In preparation for the field study, we examined various commercial and academic logging programs and found they typically collected a small subset of interactions with the browser, such as the time and URLs of visited pages, but did not log the use of navigation mechanisms or any other interactions with the web browser. We also examined instrumented web browsers used in previous research (Claypool, Le, Waseda, & Brown, 2001; Reeder, Pirolli, & Card, 2001) but found they were either out of date or lacked standard browser functionality. Another alternative we explored was the use of a Browser Helper Object (BHO) in conjunction with IE. A BHO is a dynamic link library (DLL) file that loads every time IE loads. While the BHO allowed us to easily track the page title and URLs, as well as a subset of browser functions, it was virtually impossible to log the use of many web browser navigation mechanisms. We also explored the use of screen capture software. This approach presented two problems. First, the software created a delay on older systems and we did not want to limit our recruitment to participants with high powered machines. Second, all captured video would then have to be coded by hand, which would be extremely time-consuming.

Based on our examination of available tools for studying user behaviour on the web, we decided to build a custom web browser (shown in Figure 3), which would allow us full control over the data logging. IE was chosen as our target browser because it was the most commonly used browser at the time of the study, with a reported usage share of 87% (OneStat.com, 2004). The custom web browser was built in C# using the Microsoft .NET web browser control. The web browser control provides the web page viewing window and has several properties, methods, and events that can be used to implement features found in IE (Microsoft Developer Network, 2007). For instance, the web browser control has access to standard web browser methods, such as Navigate(), GoBack(), GoForward(), Stop(), and

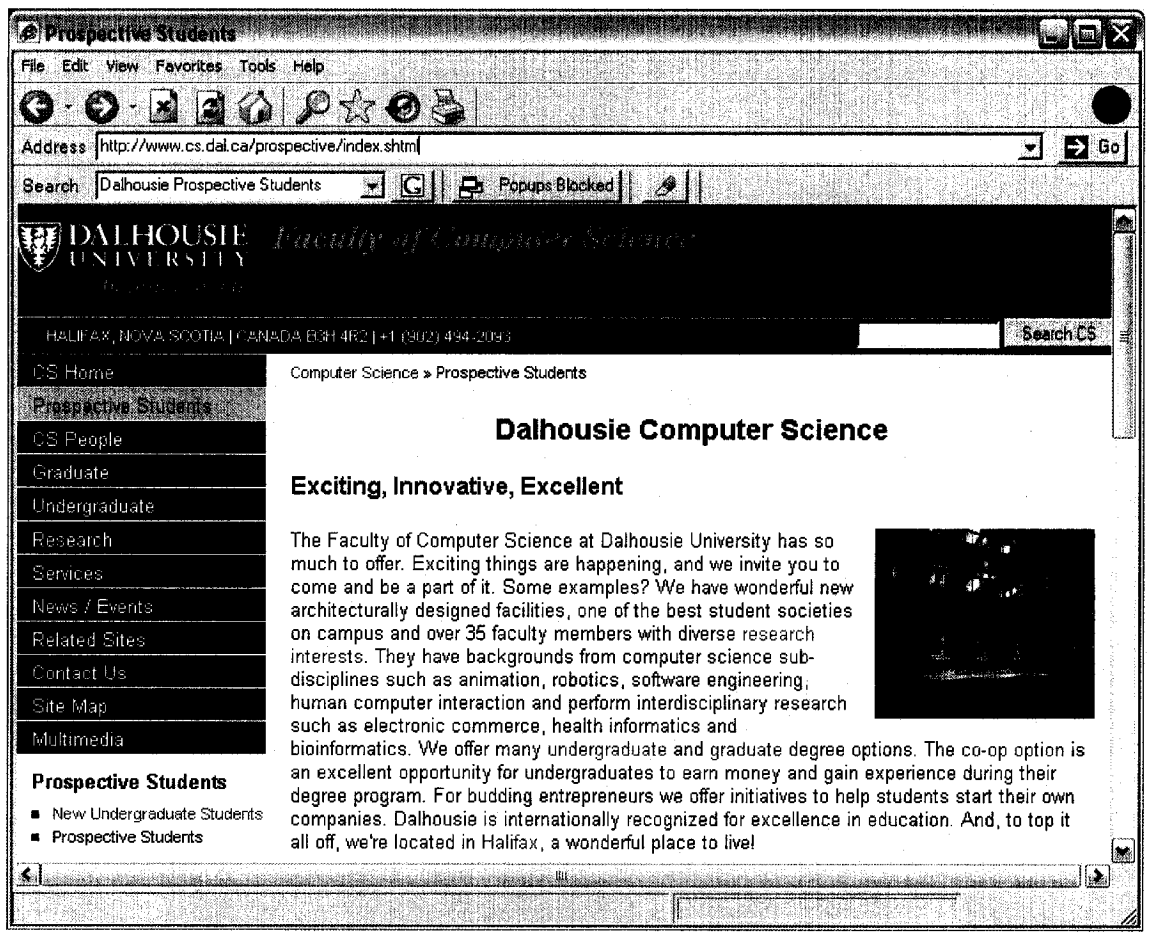

Figure 3. A screen shot of the custom-built web browser, which was built to mimic Microsoft Internet Explorer.

Refresh(). We then implemented all other web browser functionality, including bookmarks (Favorites in IE), history, menu items, and toolbars. We also implemented a Google toolbar as it is a common web browser plug-in.

Our custom browser was built to mimic IE in all areas of look and feel (e.g., shortcuts, icons, etc.). One advantage of the custom web browser approach based on the web browser control was that it was trivial to access a user's IE Favorites and previous IE history data. This meant that participants had access to the history data previously generated in IE and we were not required to import participants' Favorites into the custom web browser.

A log file was generated to provide a detailed summary of all user interactions within the web browser during each web session. Table 4 displays a sample log file. Two main types of web browser interactions were recorded: Navigation Events and Browser Events.

Table 4. Event log file with Browser and Navigation Events.

Window ID	Date Time	Page Title	URL	Event	Description
1705942	12/13/2006 19:15:10.51	MSN.com	http://www.msn.com/	**Browser Event**	OPEN SESSION
1705942	12/13/2006 19:15:16.67	MSN.com	http://www.msn.com/	**Navigation Event**	FIRST PAGE
1705942	12/13/2006 19:15:29.04	Dalhousie Prospective Students – Google Search	http://www.google.com/search?q=Dalhousie+Prospective+Students	**Navigation Event**	GOOGLE TOOLBAR
1705942	12/13/2006 19:15:30.87	Prospective Students	http://www.cs.dal.ca/prospective/index.shtml	**Navigation Event**	CLICKED LINK
1705942	12/13/2006 19:16:12.00	Dalhousie Prospective Students – Google Search	http://www.google.com/search?q=Dalhousie+Prospective+Students	**Navigation Event**	BACK BUTTON
1705942	12/13/2006 19:16:49.77	Dalhousie Prospective Students – Google Search	http://www.google.com/search?q=Dalhousie+Prospective+Students	**Browser Event**	BROWSE FAVORITES MENU
1705942	12/13/2006 19:20:06.99	CNN.com International	http://edition.cnn.com/	**Navigation Event**	FAVORITES
1705942	12/13/2006 19:39:14.72	Faculty of Computer Science – Melanie Kellar	http://flame.cs.dal.ca/~melanie	**Navigation Event**	AUTO COMPLETE
1705942	12/13/2006 19:39:14.72	Faculty of Computer Science – Melanie Kellar	http://flame.cs.dal.ca/~melanie	**Browser Event**	COPY TEXT
1705942	12/13/2006 19:39:14.72	Faculty of Computer Science – Melanie Kellar	http://flame.cs.dal.ca/~melanie	**Browser Event**	CLOSE SESSION

Navigation Events were recorded as each web page was loaded and were associated with the web browser navigation mechanism that triggered the navigation. The custom web browser differentiated between the use of the auto-complete function, selecting a URL from the drop-down address menu, and typing a URL directly into the address bar; the use of these navigation mechanisms were logged as separate navigation events. In the case of bookmarks, we recorded whether they were accessed through the side window, the drop-down menu or the links toolbar. We encountered some navigation events that could not be easily identified, such as form submissions and hyperlinks loaded through JavaScript or Flash. In these cases, we could detect that a high-level document complete event fired (i.e., a single page loaded) but could not identify the direct source of the navigation event. In these cases, the

navigation was categorized as "other". We observed that these events often occurred within web-based email and other online applications, which often use forms and JavaScript. New Window events typically consisted of new windows initiated either by the user or automatically from a script. However, the custom web browser provided a pop-up blocker so it is unlikely that pop-up advertisements accounted for a significant portion of the New Window events. A listing of all Navigation Events logged is shown in the leftmost column of Table 3.

Browser Events consisted of all other menu, button, and shortcut interactions with the web browser. This included actions such as opening and closing a window; printing or saving a document; and edit functions such as cut/copy/paste. A listing of all Browser Events is shown in Table 3.

4.3.2 Qualitative Task Data

The qualitative task data consisted of a user's task categorization (Fact Finding, Information Gathering, Just Browsing, Transactions, and Other) and a short textual description of the task (e.g., "Reading the news", "Looking for an email address"). Participants were asked to categorize all web activity recorded by the custom web browser and not just usage thought to be information seeking related. Based on the results of the pilot study, participants were given the option to provide their task information using one of two electronic diary methods: in real-time using the task toolbar shown in Figure 4 or the task diary shown in Figure 5, which was filled out at the end of each day. Participants could also use a combination of both techniques.

Participants who used the toolbar method were instructed to fill in the toolbar before beginning a new task. An auto-complete function was implemented for the textual description based on feedback received during the pilot study. Participants quickly built a small library of tasks to choose from when assigning task information for repeated tasks. Tool tips displaying task definitions were displayed when a participant hovered over one of the task buttons with their mouse.

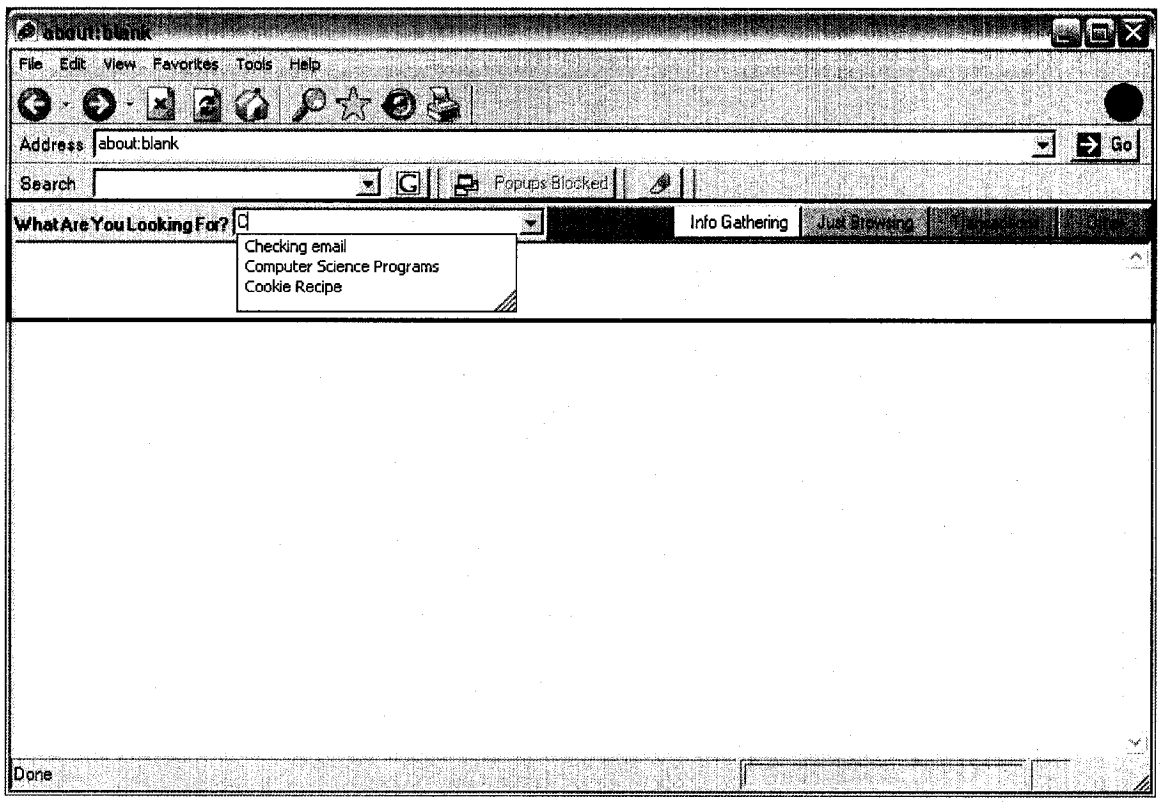

Figure 4. The custom web browser was outfitted with a task toolbar (highlighted in red) for participants to record their task information in real time.

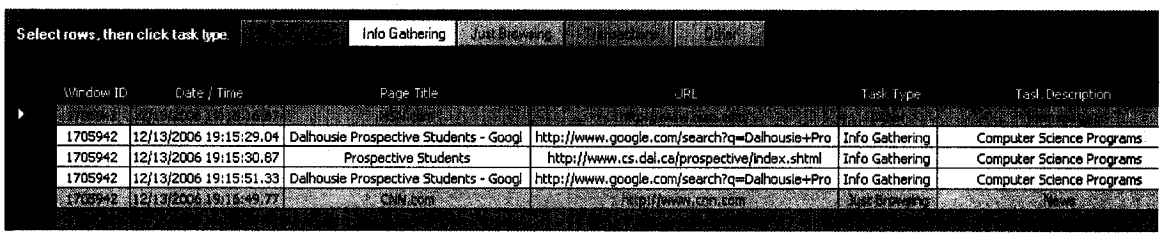

Figure 5. Participations could also use the task diary to record their task information at the end of each day.

Participants who chose to use the task diary to assign task information were instructed to do so at the end of day. The task diary, similar to the approach used by Hawkey and Inkpen (2005a), allowed participants to assign task information to multiple URLs at once. To record task type, participants selected one or multiple rows in the task diary and clicked on the appropriate task button. Textual task descriptions were entered either directly into the task diary for single entries or through a dialog box if multiple rows were selected. Similar to the task toolbar, an auto-complete function was implemented for the task diary. The items in the

auto-complete function were shared between the toolbar and the task diary. Tool tips displaying task definitions were displayed when a participant hovered over one of the task buttons with their mouse. The task diary also allowed participants to delete any web site addresses they were uncomfortable sharing with the researchers involved in the study. It was hoped that this would help encourage participants to work on the Web as they normally would.

Regardless of the electronic diary method used to collect the task information, each URL visited was associated with a task categorization and description. This information was recorded in a log file, shown in Table 5, in the following format: Window ID, Date & Time, Page Title, URL, Task Type, and Task Description.

Each participant was asked to email their data to the study researcher at the end of each day using a custom email application. This application emailed both log files to the researchers. This allowed the researchers to ensure that participants were correctly recording their data without problems. Researchers could also contact participants if more than two days passed without any data submitted to determine if there were any problems. Before analysis, timestamps were used to merge the two log files shown in Table 4 and Table 5 together.

Table 5. Task information log file.

Window ID	Date & Time	Page Title	URL	Task Type	Task Description
1705942	12/13/2006 19:15:16.67	MSN.com	http://www.msn.com/	Other	Homepage
1705942	12/13/2006 19:15:29.04	Dalhousie Prospective Students – Google Search	http://www.google.com/search?q=Dalhousie+Prospective+Students	Information Gathering	Computer Science Programs
1705942	12/13/2006 19:15:30.87	Prospective Students	http://www.cs.dal.ca/prospective/index.shtml	Information Gathering	Computer Science Programs
1705942	12/13/2006 19:16:12.00	Dalhousie Prospective Students – Google Search	http://www.google.com/search?q=Dalhousie+Prospective+Students	Information Gathering	Computer Science Programs
1705942	12/13/2006 19:20:06.99	CNN.com International	http://edition.cnn.com/	Browsing	News
1705942	12/13/2006 19:39:14.72	Faculty of Computer Science – Melanie Kellar	http://flame.cs.dal.ca/~melanie	Fact Finding	Email Address

4.3.3 Questionnaires

Participants completed three separate questionnaires over the course of the study. During the pre-study session, a demographic questionnaire (Appendix A) was used to collect participants' demographic information and current web usage. An inventory questionnaire of the web browser navigation mechanisms used (Appendix B) was also completed by participants during the pre-study session. This questionnaire was used to collect self-reported usage of web browser navigation mechanisms from participants, as well as feedback on why they used or did not use particular navigation mechanisms. Upon completion of the study, participants completed a post-study questionnaire (Appendix C) which examined any difficulties they encountered during the study, including difficulties with annotating web usage with task information, use of the two electronic diary methods and use of the custom web browser.

4.4 Summary

In this chapter we have outlined the pilot study and focus group that were conducted in advance of the field study. The results of these two preparatory studies shaped the design of the field study methodology and data collection techniques presented in this chapter. In Section 5.5, we reflect on our methodology and discuss the study limitations, many of which are the result of the tradeoffs involved in conducting naturalistic research.

The results of the data collected during the field study are subsequently presented in Chapter 5 and Chapter 6. In Chapter 5, we explore differences in users' web browser interactions across tasks and also present the Web Information Classification. In Chapter 6, we present a more detailed examination of the use of web browser navigation mechanisms in the context of task.

Chapter 5

Characterizing Web-based Information Seeking Tasks

In the previous chapter we presented the methodology and data collection tools used during the field study. In this chapter, we present the first set of analyses from the field study. We characterize participants' web tasks through an analysis of the web browser interactions captured during the task sessions. We also use the task data reported by participants, such as task type and task description, to build a high level understanding of web information tasks. The contributions of the research presented in this chapter are twofold. We first provide a characterization of the differences in how users interact with their web browsers across the range of information seeking tasks. This understanding has been used to provide implications for future support of web-based information seeking, as well to provide direction for future research in this area. We also present a high-level classification of users' activities on the Web according to their web information goals. This classification is based on observations during the field study as well as earlier models and frameworks (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002).

5.1 Results

5.1.1 General Task Characteristics

In this section we report general observations describing the characteristics of participants' task sessions. A task session is defined as a period of continuous web usage annotated with the same task information and with no break in usage greater than 25.5 minutes. This differs from the previous definitions of session presented in Chapter 3, in that task information is also used to delineate a session. The 25.5 minute timeout was chosen based on the work of Catledge and Pitkow (1995). In the case of Transactions, a new task session was identified either using the 25.5 minute lapse in activity or an explicit session logout indicated by the existence of the "logout" string in a Transaction URL (e.g., www.mail.yahoo.com/logout).

Table 6. The breakdown of tasks by participant.

Total Tasks by Participant	Fact Finding	Info. Gathering	Browsing	Transactions	Other
140	11.4%	19.3%	21.4%	47.9%	.0%
122	4.9%	9.8%	19.7%	64.8%	.8%
92	32.6%	1.1%	46.7%	19.6%	.0%
80	5.0%	17.5%	8.8%	68.8%	.0%
75	4.0%	18.7%	26.7%	50.7%	.0%
70	20.0%	15.7%	58.6%	5.7%	.0%
68	30.9%	2.9%	10.3%	55.9%	.0%
61	3.3%	9.8%	.0%	83.6%	3.3%
55	5.5%	7.3%	20.0%	67.3%	.0%
52	32.7%	9.6%	7.7%	25.0%	25.0%
52	40.4%	15.4%	5.8%	36.5%	1.9%
45	44.4%	15.6%	.0%	40.0%	.0%
42	14.3%	14.3%	16.7%	54.8%	.0%
41	31.7%	4.9%	4.9%	58.5%	.0%
39	17.9%	43.6%	35.9%	2.6%	.0%
37	.0%	13.5%	13.5%	67.6%	5.4%
32	18.8%	28.1%	21.9%	31.3%	.0%
28	60.7%	21.4%	3.6%	14.3%	.0%
25	8.0%	12.0%	36.0%	44.0%	.0%
20	25%	5.0%	.0%	70.0%	.0%
16	31.3%	.0%	12.5%	50.0%	6.3%

Overall, participants recorded 1192 task session involving 13,498 pages over the week long study. The mean number of task sessions completed per participant was 56.8 (median = 52, SD = 31.97) with a range of 16 to 140 tasks. A breakdown of the number of task sessions completed by each participant is shown in Table 6. We found there was no significant difference between the Computer Science (CS) and non-CS groups in terms of the number of task sessions completed. The CS group recorded a mean of 58.4 task sessions while the non-CS group recorded a mean of 54.1 task sessions.

For each task type, we present the breakdown of all task sessions across all participants (see Figure 6). We also examined the occurrence of repeated tasks on a per-participant basis. A task was defined as repeated if, within a participant's list of tasks, there were multiple occurrences of a task session with the same task categorization and similar task description. For instance, two task sessions categorized as Fact Finding and labelled as "checking the weather" and "weather" respectively, were recorded as repeated tasks. Finally, we examined the ways in which participants described their tasks. For example, terms used within the task descriptions, such as "finding", "searching", "looking", and "checking", are useful in characterizing the nature of a task. Table 7 shows a subset of the tasks submitted by participants during the study.

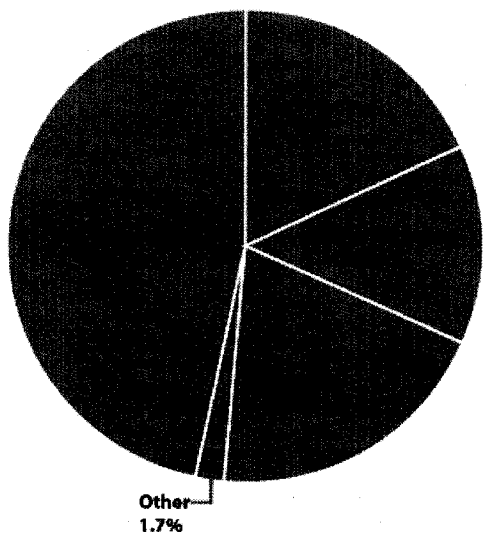

Figure 6. Breakdown of all tasks captured.

Table 7. A sample of the task descriptions collected during the field study.

Fact Finding	**Information Gathering**
Looking , Searching, or Checking for:	**Looking, Researching, Information on:**
A book in the library	A new laptop
A course mark	Admissions information
A file (for download)	Beer distribution
A phone number	Breastfeeding
A recipe	Building a computer system
A research paper	Conferences
An Email address	Health Economics
Assignment information	Help with a virus
Bus schedule	iPod prices
Definitions	Job Searching
Directions to a restaurant	Making a resume
Exam dates	Math tools on the Web
Help with a game	New wireless card
How to reference a memo	Palm OS development
How to tie a tie	Papers on policy-based network
Java documentation	Renting a car
Movie times	Risk analysis
Song lyrics	Summer school courses
The average mass of a bullet	
Weather	
Browsing	**Transactions**
Looking for, Reading:	**Checking, Updating:**
Blogs	Applying for a credit card
Browsing web site	Banking
Comics	Blogs
Entertainment	Profile information
Friend's Homepage	Document delivery request
Gaming forum	Doing an online test
Link received in email	Email
Listening to music	Logging diet and exercise
News	Online accounts
Movie trailers	Online MSN
MP3s	Online shopping
Updates on movie web site	Sending a greeting
Wasting time	Taking part in a survey

5.1.1.1 Fact Finding

Fact Finding (FF) task sessions accounted for 18.3% (218/1192) of all web usage. Looking for weather information appeared to be the most common Fact Finding task, accounting for 11.5% (25/218) of task sessions in this category. Other common Fact Finding tasks included looking for course or assignment related material, song lyrics, and specific software. Fact Finding tasks appeared to be somewhat split between ludic and school/work-related activities

Of the 218 Fact Finding task sessions, we found that 55.5% (121/218) were repeated at least once and this category had the lowest proportion of repeated tasks. There appears to be three main reasons why Fact Finding tasks were repeated: monitoring, re-finding, and task variants. When monitoring, participants were looking for specific dynamic information, such as the current weather forecast. When re-finding, participants were looking to return to a previously found piece of static information. Task variants occurred when participants were looking for related pieces of specific information, such as looking for programming resources. One example of this was a participant who labelled two tasks "looking for Java documentation" where in one case he was looking for information on hash tables while in another case he was looking for Java documentation on substrings.

When participants described their Fact Finding tasks through the task diary, they often used terms such as "checking" (e.g., checking the weather), "finding" (e.g., finding a phone number), "looking" (e.g., looking for Win XP Pro upgrade), and "searching" (e.g., searching for a journal) for information.

5.1.1.2 Information Gathering

Information Gathering (IG) task sessions accounted for 13.4% (160/1192) of all web usage. There was no single representative task but common tasks included job hunting, course or project related research, researching a new purchase (such as a computer or iPod), and course/admissions information. Many of the Information Gathering tasks were related to technology concepts, such as researching new computer equipment.

For Information Gathering task sessions, 58.8% (94/160) of tasks were repeated at least once. Information Gathering tasks appeared to be repeated because participants continued with their tasks at a later time. Since Information Gathering tasks tend to be longer in duration, they were often broken up over a day or even over several days. Among some participants, we saw Information Gathering tasks that stretched over as many as six days, such as a participant who was researching graduate school admission information.

Participants who completed Information Gathering tasks often used terms such as "finding" (e.g., finding information for a presentation), "information" (e.g., information for a project), "looking" (e.g., looking for a new laptop), "research" (e.g., risk analysis research), and "searching" (e.g., searching for iPod prices) to describe their task. While some of the terms used are common to Fact Finding task descriptions (finding, searching), terms such as information and research highlight the collection of information that takes place during information gathering tasks.

5.1.1.3 Browsing

Browsing (BR) task sessions accounted for 19.9% (237/1192) of all web usage. Browsing tasks appeared to be primarily ludic in nature, and consisted of news reading in 40.5% (96/237) of tasks in this category. Other common tasks included reading blogs, visiting gaming related sites, and reading music/TV/movie related web pages.

Browsing tasks were highly repetitive as 84.4% (200/237) of task sessions were repeated at least once. Browsing tasks were primarily habitual or monitoring tasks, such as checking the news or a friend's blog. We observed many participants who repeated the same Browsing tasks daily over the course of the study.

Participants often used the words "looking" (e.g., looking for a blog update) and "reading" (e.g., reading the news) to describe their Browsing tasks. Often when participants said they were "looking for" something during Browsing, it was in the context of browsing for a hobby or travel related interest.

5.1.1.4 Transactions

Transactions (TR) were the most frequently recorded task sessions, accounting for 46.7% (557/1192) of all web usage. Transactions were primarily made up of web-based email, accounting for 80.4% (448/557) of all Transactions and 38% of all web usage. Other types of Transactions recorded by our participants included online bill payments and blog/message board entries.

Transactions were the most often repeated task, with 95.2% (530/557) of task sessions repeated at least once. As previously mentioned, Transactions consisted mainly of email, which was often accessed by participants several times during the day.

We observed that tasks categorized as Transactions appeared to have two distinct goals. The first is the communication of information through email, blog updates, or postings to message boards. The second is the completion of online actions, such as online banking and shopping. This distinction mirrors the task categorization of Communication & Transactions that evolved out of the focus group in Chapter 4.

5.1.1.5 Other

Finally, only a small number of task sessions were categorized as Other (OT) and they accounted for 1.7% (20/1192) of all web usage. These were tasks such as viewing web pages during web development and may have been specific to our user population (i.e., mainly CS students).

5.1.2 Differences in Implicit Measures across Tasks

Implicit measures consist of the collection of user behaviour traces and web browser interactions that can be recorded without any intervention on the part of the user (Kelly & Teevan, 2003). Typically, this includes measures such as dwell time; mouse, keyboard, and scrolling activity; and interactions with a web document, such as saving or printing. Implicit measures have been studied as a non-obtrusive method for inferring user interest extensively (Claypool, Le, Waseda, & Brown, 2001; Kelly & Belkin, 2004; Morita & Shinoda, 1994; Oard & Kim, 2001). In our research, we have not used implicit measures to infer interest but rather to define task characteristics.

While implicit measures may be used on a per page basis, we examined the implicit measures recorded over the course of a task session. Previous researchers (Mat-Hassan & Levene, 2005; Seo & Zhang, 2000) have used implicit measures to explore information seeking behaviour on the Web. In this research however, we are studying users' information seeking behavior over a wide range of information needs rather than within a single portal or dataset.

We were interested in studying elements of user behaviour while users were engaging in information seeking tasks that could be collected implicitly, i.e., without any intervention from the user. We studied the following elements: dwell time; windows opened; pages loaded; use of web browser navigation tools; time of day; use of Google; use of site specific searches; and use of web browser functions. We present descriptive statistics and the results of statistical analysis where appropriate. Raw data was analyzed using nonparametric one-way ANOVAs (Kruskal-Wallis) because the data did not exhibit a normal distribution. Nominal data was analyzed using chi-square tests. An alpha value of 0.05 was used for all omnibus tests. Pairwise post hoc tests were conducted using the Mann-Whitney and chi-square tests and alpha values were determined using the Bonferroni correction (in order to decrease the possibility of Type 1 errors). Tasks labelled as Other appeared to be specific to our population and accounted for only a small percentage of all tasks (1.7%). Therefore, we only report descriptive statistics for this task type and it was not included in any statistical analyses.

5.1.2.1 Dwell Time

In a field setting, it can be problematic to accurately record dwell time, i.e., the amount of time participants spend reading and interacting with a particular web page. Although we can record the time of each page access, it is often not possible to determine where a participant's attention is directed. In this study, we were interested in the amount of time participants spent completing their information seeking tasks. Task duration was measured from the time the first page in a task session was loaded until the time in which the last page was loaded. This means that duration was only measured for task sessions in which more than one page was loaded, excluding 192 (16%) sessions. This method resulted in a smaller, but more reliable, set of task duration data. However, it is still important to note this data is not as reliable as laboratory collected task duration data.

The mean time per task session is shown in Figure 7. The mean duration recorded for Fact Finding task sessions was 481.6 seconds (SD = 1169.9). The mean duration of Information Gathering task session was 1087.6 seconds (SD = 2048.0). The mean duration recorded for Browsing task sessions was 648.1 seconds (SD = 856.5). The mean duration for Transactions was 468.7 seconds (SD = 1084.4). Finally, the mean duration for Other task sessions was 437.9 seconds (SD = 692.5).

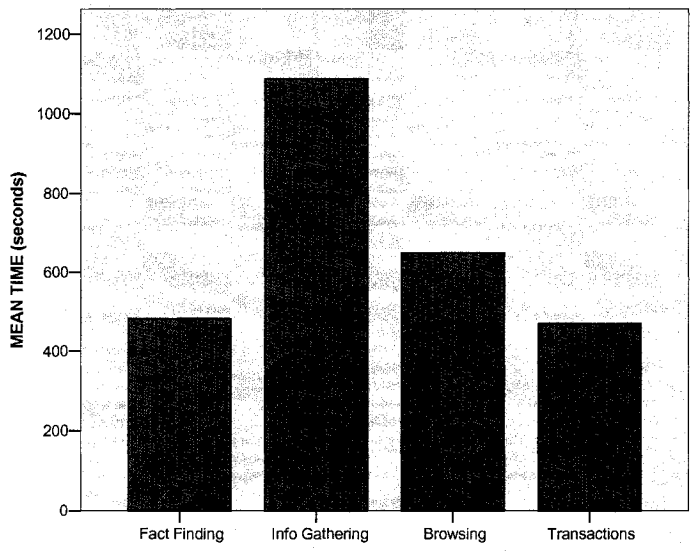

Figure 7. The mean time (in seconds) spent completing each task.

Significant differences were found for dwell time between task sessions (Kruskal-Wallis H = 40.720, $df = 3, p \approx 0.000$). Pairwise comparisons were conducted using the Mann-Whitney test and an adjusted alpha level of 0.008 was used. Information Gathering task sessions were significantly longer than both Fact Finding ($p \approx 0.000$) and Transactions ($p \approx 0.000$) but not Browsing. Browsing task sessions were also significantly longer than both Fact Finding ($p \approx$ 0.000) and Transactions ($p \approx 0.000$).

The task duration data exhibited a high degree of variability for each task type. Task duration can be influenced by the task complexity, familiarity with the task (e.g., habitual tasks), and domain knowledge (Hölscher & Strube, 2000). The duration of a Transaction task session, for instance, may depend on the amount of email an individual receives over the course of a day and the number of times the email account is accessed.

5.1.2.2 Windows Opened

The number of windows opened during each task session was calculated by counting the number of unique window IDs recorded during a single task session. The custom web browser provided a pop-up blocker so pop-up advertisements did not have a large impact on the number of windows opened. It should be noted that web-based email clients differ in the number of windows launched for email operations. For example, one email client we logged opened a new window for each composed or read message while another used the existing browser window for all operations. Therefore, the number of windows opened during Transaction task sessions may be highly variable.

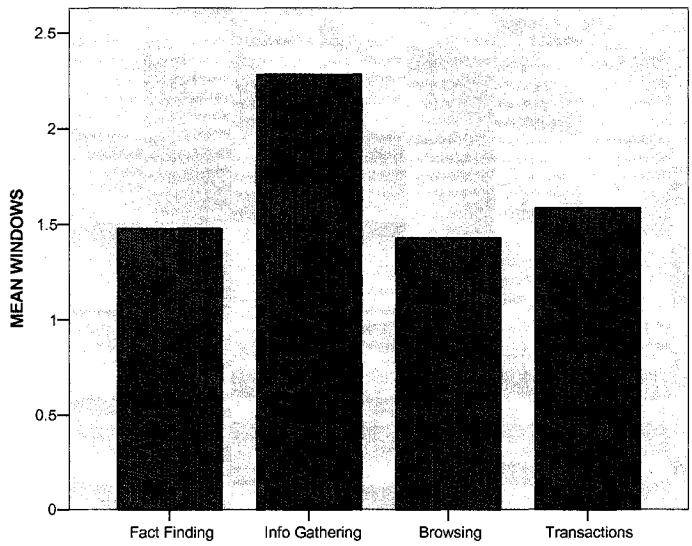

Figure 8. The mean number of windows opened during each task.

In general, a low number of windows were opened across the different task sessions; a total of 1934 windows were loaded during the field study. Figure 8 displays the mean number of windows opened across all task sessions. The mean number of windows opened during Fact Finding task sessions was 1.48 (median = 1, SD = 1.34). For Information Gathering task sessions, the mean number of windows opened during a task was 2.28 (median = 1, SD = 3.21). For Browsing task sessions, the mean number of windows opened was 1.43 (median = 1, SD = 1.05). The mean number of windows opened during Transactions was 1.58 (median

=1, SD = 1.34). Finally, the mean number of windows opened during Other task sessions was 1.35 (median = 1, SD = 0.99).

Significant differences were found for the number of windows opened between tasks (Kruskal-Wallis $H = 15.650$, $df = 3$, $p = 0.001$). Pairwise comparisons were conducted using the Mann-Whitney test and an adjusted alpha level of 0.008 was used. More windows were opened during Information Gathering task sessions than both Fact Finding (p=0.003) and Browsing ($p = 0.005$). Significant differences were also found between Fact Finding and Transactions ($p = 0.006$). Due to the small number of windows opened overall, these results do not have strong practical significance. We reflect on the small number of windows opened during the task sessions in Section 5.4.5.

5.1.2.3 Pages Loaded

The number of pages loaded during a task session was calculated by counting the number of top level frames loaded. This means that for pages with frames, only one page was counted. Similar to the number of windows opened, the number of pages loaded for Transactions were influenced by the web-based email services, some of which loaded a new page for each emailed viewed or sent while others loaded a single page for the entire session.

Figure 9 displays the mean number of pages loaded across all task sessions. A total of 13,498 pages were loaded during the field study. The mean number of pages loaded during Fact Finding task sessions was 8.1 (median = 5, SD = 9.7) During Information Gathering task sessions, the mean number of pages loaded was 31.4 (median=8, SD = 61.8). For Browsing task sessions, the mean number of pages loaded was 10.3 (median = 5, SD = 15.2). During Transactions, the mean number of pages loaded was 7.3 (median = 4, SD = 10.0). Finally, during Other task sessions, the mean number of pages loaded was 11.2 (median = 4, SD = 21.2).

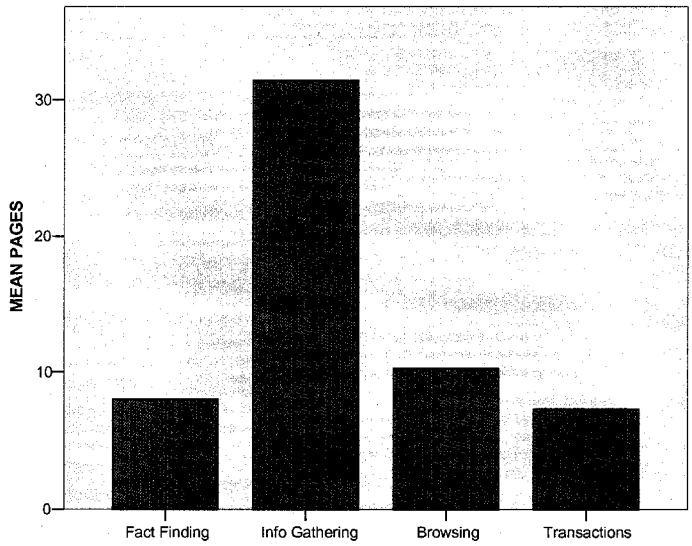

Figure 9. The mean number of pages opened during each task.

Significant differences were found for the number of pages loaded between task sessions (Kruskal-Wallis $H = 49.904$, $df = 3$, $p \approx 0.000$). Pairwise comparisons were conducted using the Mann-Whitney test and an adjusted alpha level of 0.008 was used. The number of pages viewed during Information Gathering task sessions was significantly higher than all other tasks: Fact Finding ($p \approx 0.000$), Browsing ($p \approx 0.000$) and Transactions ($p \approx 0.000$).

5.1.2.4 Web Browser Navigation Mechanisms

In this section, we conducted statistical analysis to determine whether there were significant differences in the usage of web browser navigation mechanisms to initiate new task sessions. Figure 10 shows the per task usage of web browser navigation mechanisms to initiate new task sessions.

Within Fact Finding task sessions, there were significant differences between the navigation mechanisms used [χ^2 (9, N =218) = 233.101, $p \approx 0.000$]. Typed-in URLs were the most common method (73/218 – 33.5%) for initiating Fact Finding task sessions, followed by the Google toolbar (51/218 – 23.4%) and bookmarks (32/218 – 14.7%). Pairwise comparisons (alpha = 0.005) revealed that typed-in URLs were used significantly more than all other navigation mechanisms ($p \approx 0.000$ for all), with the exception of the Google toolbar.

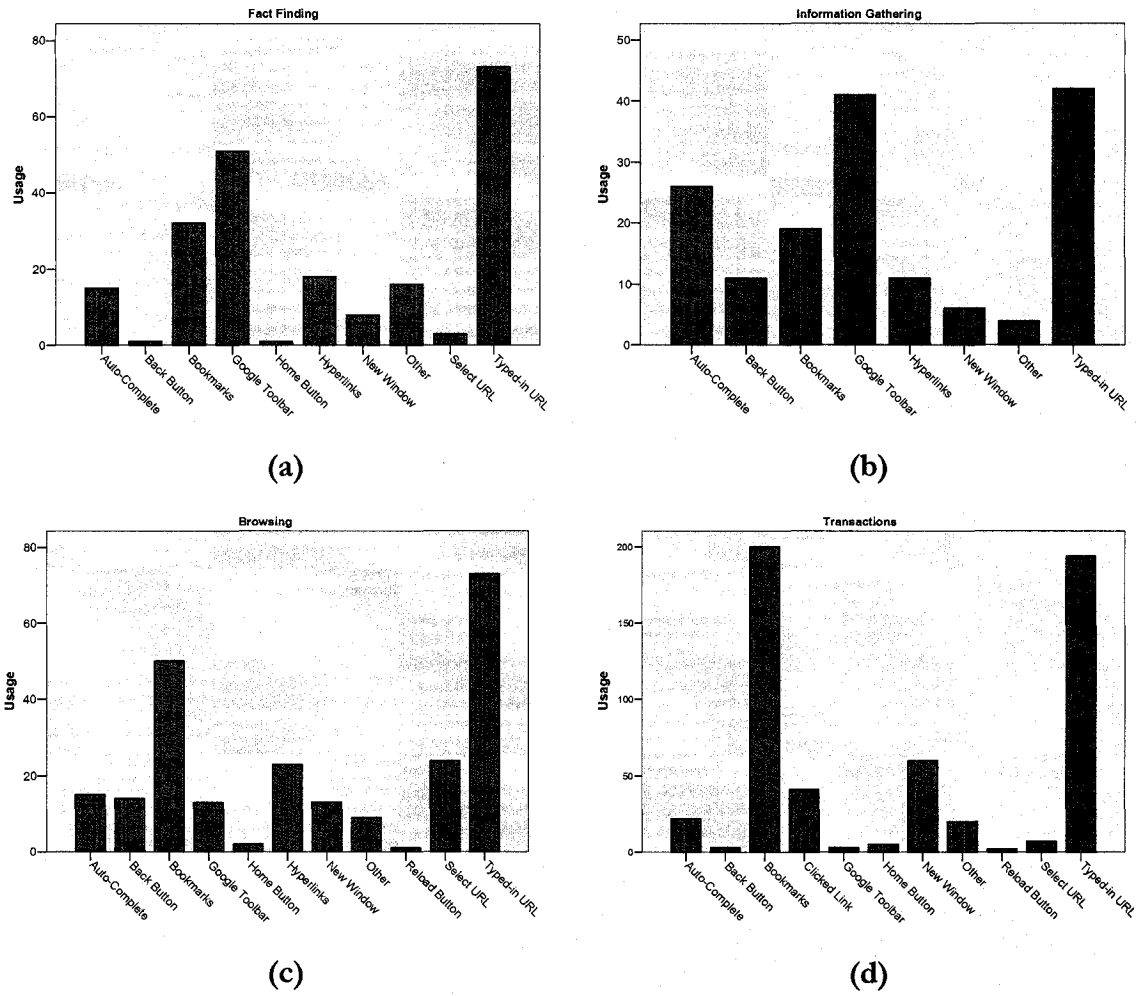

Figure 10. Shows the use of web browser navigation mechanisms to initiate new Fact Finding (a), Information Gathering (b), Browsing (c), and Transaction (d) task sessions.

Similarly, the use of the Google toolbar was significantly higher than all other navigation mechanisms ($p \approx 0.000$ for all), with the exception of bookmarks.

Within Information Gathering task sessions, there were significant differences between the navigation mechanisms used [$\chi^2(7, N=160) = 78.800, p \approx 0.000$). These tasks were commonly initiated through typed-in URLs (42/160 – 26.3%), followed by the Google toolbar (41/160 – 25.6%), and the auto-complete function (26/160 – 16.3%). The use of navigation mechanisms when initiating Information Gathering tasks appears to be evenly distributed between the most commonly used mechanisms; pairwise comparisons (alpha = 0.005) did not reveal a significant difference between the use of typed-in URLs, the Google

toolbar, and the auto-complete function. Significant differences were detected between the two most commonly used navigation mechanisms (typed-in URLs and the Google toolbar) and all other navigation mechanisms ($p \leq 0.002$), with the exception of the auto-complete function.

There were significant differences within Browsing task sessions between the navigation mechanisms used [$\chi^2(10, N=237) = 216.878, p \approx 0.000$). Browsing task sessions were most commonly initiated through typed-in URLs (73/237 – 30.8%), followed by bookmarks (50/237 – 21.1%), and selected URLs (24/237 – 10.1%). Pairwise comparisons (alpha = 0.005) did not reveal a significant difference between the use of typed-in URLs and bookmarks. However, typed-in URLs and bookmarks were used more often than all other browser navigation mechanisms ($p \approx 0.000$ for all).

Within Transactions, there were significant differences between the navigation mechanisms used [$\chi^2(10, N=557) = 1099.853, p \approx 0.000$). Transactions were primarily initiated through two mechanisms: bookmarks (200/557 – 35.9%) and typed-in URLs (194/557 – 34.8%). Pairwise comparisons (alpha = 0.005) did not show a significant difference between the use of these two mechanisms; however, bookmarks and typed-in URLs were used more often than all other mechanisms ($p \approx 0.000$ for all). Finally, task sessions labelled as Other were most commonly initiated using typed-in URLs (9/20 – 45%).

5.1.2.5 Time of Day

We categorized the time during which a task session was initiated across four time of day categories: morning (6am-11:59 am), afternoon (12:00pm-5:59pm), evening (6:00pm-11:59pm), and overnight (12:00am-5:59am). Previous research by Beitzel, Jensen, Chowdhury, Grossman, and Frieder (2004) reported that time of day had an impact on the popularity and uniqueness of topically categorized queries. In our research, we were interested in knowing, for each task type, what was the most common time of day in which that task was initiated. The bar chart presented in Figure 11 shows the proportion of task sessions by time of day. Pairwise comparisons were conducted using chi-square tests with an adjusted alpha level of 0.008.

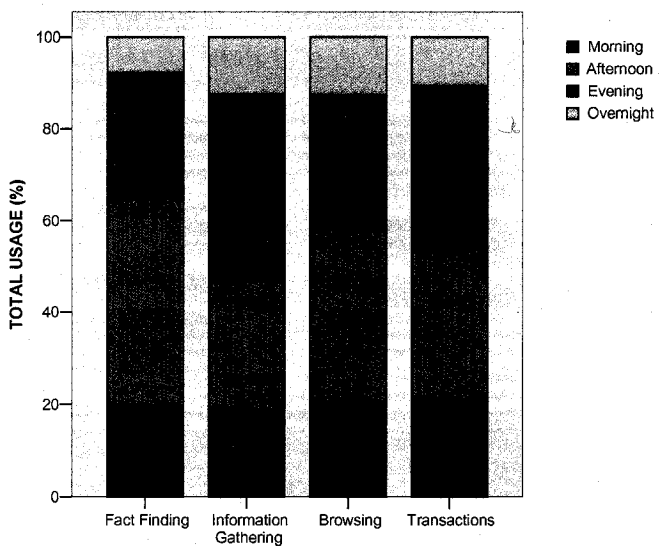

Figure 11. The proportion of tasks across time of day.

Within Fact Finding task sessions, there were significant differences between time of day [χ^2 (3, N =218) = 63.505, $p \approx 0.000$). Fact Finding task sessions most commonly occurred in the afternoons (98/218 – 45%). Pairwise comparisons revealed significant differences between occurrences in the afternoon and all other times of day: morning ($p \approx 0.000$), evening ($p = 0.003$), and overnight ($p \approx 0.000$).

Within Information Gathering task sessions, there were significant differences between time of day [χ^2 (3, N =160) = 28.750, $p \approx 0.000$). Information Gathering task sessions most commonly occurred in the evenings (65/160 – 40.6%). Pairwise comparisons revealed a significant difference between occurrences in the evenings and both morning ($p = 0.000$) and overnight ($p = 0.000$) but not the afternoon.

There were significant differences within Browsing task sessions between the navigation types [χ^2 (3, N =237) = 32.755, $p = 0.000$). Browsing tasks were most commonly recorded in the afternoon (89/237 – 37.6%). Pairwise comparisons revealed a significant difference between occurrences in the afternoon and both morning ($p = 0.001$) and overnight ($p \approx 0.000$) but not the evening.

Within Transactions, there were significant differences between the time of day [χ^2 (3, N =557) = 87.919, $p = 0.000$). Transactions were very closely split between afternoons, which

accounted for 180/557 (32.3%) of all Transactions, and evenings, which accounted for 200/557 (35.9%) of all Transactions. Pairwise comparisons did not find a significant difference between these two time periods. Significant differences were however found between these two time periods and both mornings ($p \approx 0.000$) and overnight ($p \approx 0.000$).

Finally, task sessions classified as Other were almost evenly distributed over mornings (6/20 – 30%), afternoons (6/20 – 30%), and evenings (7/20 – 35%).

5.1.2.6 Use of Google

The use of Google has become ubiquitous in today's web environment. Aula, Jhaveri, and Käki (2005) found that 95.3% of 236 survey respondents reported using Google as their primary search engine. We examined the use of Google by participants across the different task sessions. All URLs were mined for the string "Google". After eliminating Google email (GMail) and within-site searches (provided by Google and addressed in Section 5.1.2.7), we recorded the number of queries submitted to Google per task. We saw very little evidence of the use of alternate search engines (<1%), with the exception of those used for site specific searches.

Google was accessed in 78/218 (35.8%) of the Fact Finding task sessions. When Google was used within Fact Finding task sessions, the mean number of queries submitted was 2.18 (SD = 3.90). Within Information Gathering, Google was used in 66/160 (41.25%) of all task sessions. The mean number of queries submitted per Information Gathering task session was 2.72 (SD = 3.08). The use of Google dramatically declined for the remaining task sessions, occurring only in 8.43% of Browsing tasks sessions, 0.005% of Transactions, and 0.05% of Other task sessions. We found that in addition to the main Google search engine, participants also used the Google Image, Scholar, and Map searches. There were no significant difference in the use of Google between Fact Finding and Information Gathering task sessions, nor was there any difference in the number of queries between the two tasks.

We also examined the difference in the query length submitted to Google between Fact Finding and Information Gathering task sessions. The mean query length for Fact Finding task sessions was 4.72 words (SD = 2.57), compared with 3.32 words (SD = 2.26) for Information Gathering task sessions. On average, Fact Finding queries were longer than

Information Gathering queries (t (337) = 5.360, p = 0.000). Due to the nature of the task, participants often submitted very specific query strings when completing Fact Finding tasks, such excerpts of song lyrics, partial or full publication titles, and specific questions (e.g., "how do I oil the heatsink fan"). The queries submitted during Information Gathering task sessions were more vague and tended to represent general topics rather than specific questions (e.g., graduate school admission information).

5.1.2.7 Site-specific Searches

The use of site-specific searches was also examined. These were defined as searches that were conducted within a specific web site or domain. To retrieve these instances, we collected all URLs which contained the term "q=", which is a typical string used to represent queries within a URL. We then removed all Google searches processed in the previous section but included searches of individual domains powered by Google.

The most common site-specific searches were product searches within commercial web sites and searches within online databases or libraries. Overall, we saw a small number of site specific searches (27 in total), most of which occurred within Information Gathering tasks (19/27 – 70.3%). Six instances (6/27 – 22.2%) were found within Fact Finding tasks and two (2/27 – 7.4%) were found within Browsing tasks. The small amount of data collected did not warrant any statistical analysis.

5.1.2.8 Use of Browser Functions

Browser function use was logged and associated with the task session being performed in the corresponding window. We were interested in how the use of these functions (copy, paste, cut, find on page, print, save, and the creation of new bookmarks) differed across task sessions. A total of 178 browser functions were recorded across all participants and the breakdown within task sessions is shown in Table 8.

Information Gathering task sessions recorded the highest number of browser functions (97/178 – 54.5%), followed by Fact Finding (33/178 – 18.5%), Transactions (25/178 – 14.0%), and Browsing (23/178 – 12.9%). Significant differences were found between task and the following tools: creating bookmarks [$\chi^2(3, N=45) = 34.022, p = 0.000$)], using the

find function[$\chi^2(2, N=17) = 8.941, p = 0.001$)], copying[$\chi^2(3, N=39) = 17.308, p = 0.001$)], and pasting [$\chi^2(3, N=67) = 24.164, p = 0.000$)].

Pairwise comparisons (alpha = 0.008) using chi-square analysis found significant differences between Information Gathering and all other tasks for creating bookmarks, copy, and paste ($p < 0.004$ for all). Within Information Gathering tasks, the most common functions included pasting text (34/97 – 35.1%), copying text (21/97 – 21.6%), and creating new bookmarks (28/97 – 28.9%). Copied text typically consisted of html content (web page) and pasted text typically consisted of URLs and search strings pasted to the address and Google toolbar (combo boxes).

Table 8. The use of browser functions within tasks.

Function		Fact Finding	Info. Gathering	Browsing	Trans-actions
		(n=33)	(n=97)	(n=23)	(n=25)
Copy	Web page	1/218 (<1%)	17/160 (10.6%)	0	0
	Combo box	5/218 (2.3%)	4/160 (<1%)	6/237 (<1%)	6/557 (<1%)
Paste	Web page	0	2/160 (<1%)	0	1/557 (<1%)
	Combo box	13/218 (6.0%)	32/160 (20.0%)	9/237 (<1%)	10/557 (<1%)
Cut	Web page	0	1/160 (<1%)	0	0
	Combo box	0	0	0	0
Find on Page		5/218 (2.3%)	11/160 (6.9%)	0	1/557 (<1%)
Print		3/218 (<1%)	2/160 (<1%)	0	2/557 (<1%)
Save		2/218 (<1%)	0	0	0
Add Bookmark		4/218 (1.8%)	28/160 (17.5%)	8/237 (<1%)	5/557 (<1%)

5.2 Discussion

5.2.1 Summary of Task Characteristics

Based on the results presented in the previous section, we now provide a general characterization of each type of task for our sample population, recognizing that the task types are complex. Table 9 provides a summary of characteristics for each type of task. We have omitted time of day, since we expect that it is specific to the population sampled in this study.

Table 9. General task characterization.

Fact Finding	**Information Gathering**
• Shorter duration • Small number of pages viewed • Large search component • Relatively longer queries • Little use of browser functions • Typed in URLs, Google toolbar, Bookmarks	• Longer duration • Larger number of pages viewed • Large search component • Relatively shorter queries • Greatest use of browser functions • Typed-in URLs, Google toolbar, Auto-complete
Browsing	**Transactions**
• Shorter duration • Small number of pages viewed • Often repeated • Little use of browser functions • Typed-in URLs, Bookmarks, Select-URL	• Shorter duration • Number of pages and windows influenced by type email • Most often repeated • Little use of browser functions • Bookmarks, Typed-in URLs

Fact Finding task sessions were relatively short lived, lasting eight minutes on average. Just over half of all Fact Finding tasks were repeated at least once, and this was attributed to re-finding information, monitoring information, and conducting sets of related tasks. Tasks appeared to be evenly split between work/school and personal tasks. We observed a relatively small number of pages viewed during Fact Finding task sessions (average = 8). Typed-in URLs and the Google toolbar were the most common navigation mechanism used to initiate Fact Finding task sessions. The search nature of this task was reflected in the use of Google during 35% of Fact Finding task sessions and participants tended to submit

longer, more specific queries. The use of browser functions was minimal during task sessions of this type.

Participants exhibited a rich set of behaviour during Information Gathering task sessions. This task was the longest in duration, averaging 18 minutes per task session. We observed a relatively high number of pages loaded during Information Gathering task sessions (average = 31). Many Information Gathering tasks were related to participants' course or research work. Over half of all Information Gathering tasks were repeated at least once, and this was attributed to ongoing tasks spanning multiple task sessions. Typed-in URLs, the Google toolbar and auto-complete were the most common methods of initiating a new task session. We observed the largest number of Google searches and within-site searches during this task and the queries submitted to Google appeared to be shorter and more general than Fact Finding. We also observed the highest usage of browser functions within Information Gathering tasks. Participants were observed creating new bookmarks, using the copy and paste functions, and using the "find on this page" function.

The average length of a Browsing task session was ten minutes We observed a relatively small number of pages viewed during Browsing task sessions (average = 10). The most dominant characteristic of Browsing was the habitual nature of this task. On average, almost 85% of Browsing task sessions were repeated at least once and we observed a high degree of monitoring within this task. The most common methods of navigation when initiating a new task session were typed-in URLs and bookmarks, which support the repetitive/monitoring nature of Browsing tasks. Unlike Fact Finding and Information Gathering, participants seldom used Google or site-specific searches when Browsing. The use of browser functions was minimal within Browsing tasks.

Transactions differ from traditional information seeking tasks in that the user's goal is not to change their state of knowledge but instead to exchange or communicate information. The average length of a Transaction task session was close to eight minutes. Transactions were the most often repeated task, with 95% of all Transaction tasks repeated at least once, and consisted primarily of email. It is difficult to characterize the number of pages loaded and windows opened because these function were influenced by the type web-based email. However, within a single individual, we would expect the number of pages and windows

opened during Transactions would be more consistent. Transactions were commonly accessed using typed-in URLs and bookmarks. The use of Google, site-specific searches, and browser functions were minimal within Transactions. We found that while all tasks categorized as Transactions by our participants shared the same goal (to exchange information), there was a clear distinction between two types of tasks. Tasks with a communication component were more strictly defined as Communications (e.g., email, web publishing) while those tasks that were based on the exchange on information through online actions were categorized as Transactions (e.g., banking, online shopping). We elaborate further on this distinction in Section 5.3.

There was a clear division separating the four task types into two groups: search-based and revisitation-based. While Fact Finding and Information Gathering were characterized as search-based tasks with a heavy use of Google and site-specific searches, Browsing and Transactions were characterized by a heavy level of monitoring and revisitation. Between Fact Finding and Information Gathering, Information Gathering was a more complex task; participants interacted more with the web browser functions, viewed significantly more web pages, and spent longer periods of time completing this task.

When we compared our results with previous research (Choo, Detlor, & Turnbull, 2000 ; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002), we did not observe any consistent trends across the data. The most common information seeking task (excluding Transactions) within our study was Browsing, while Fact Finding was the most common task reported by Choo et al. (2000) and Information Gathering was the most common task reported by both Morrison et al. (2001) and Sellen et al. (2002). It is difficult to compare previous research due to the difference in task categories, populations, and methods of data collection. For instance, Morrison et al. (2001) may have found a higher incidence of Information Gathering because participants were asked to report an incident where they found information on the Web that led to a significant decision or action, which is a characteristic of Information Gathering in itself. The use of knowledge workers in previous research, compared with our use of university students, may also play a role in the differences in the distributions. Therefore, it is difficult to comment on whether the differences in usage across the studies are indicative of the evolution of information seeking behaviour on the Web or whether they are a result of methodological differences.

5.2.2 Monitoring

One activity that warrants further investigation is monitoring. We define monitoring as an activity that occurs when users return to previously visited pages in order to view new or updated information. We were unable to closely study monitoring because it was difficult to clearly categorize. This study simply gave us an informal view of monitoring, with no details on what information was being monitored or the goal of a monitoring activity.

Based on the observations collected during our research, we hypothesize that monitoring occurs with differing frequency across many tasks. We observed high levels of monitoring within Browsing and Transactions and lower levels within Fact Finding and Information Gathering tasks. However, as previously stated, the study was not designed to study Monitoring and therefore we did not collect a precise set of monitoring data. We also expect that the type of information being monitored is dependant upon the higher level web information task. More research is needed to better understand the role of monitoring within the context of task.

5.3 Web Information Classification

Upon completion of the study analysis, we reflected on the appropriateness of the task categorization used during the field study. Based on the task data collected during the field study, as well as previous work (Choo, Detlor, & Turnbull, 2000; Ellis, 1989; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002), we have developed the Web Information Classification (shown in Figure 12). The classification consists of three information goals: information seeking, information exchange, and information maintenance. Web information tasks consist of the set of tasks in which users engage on the Web that deal with some aspect of information, from acquisition, consumption, and distribution of information.

Information seeking tasks consist of Fact Finding, Information Gathering, and Browsing. These are tasks in which the user's goal is to "change their state of knowledge" (Marchionini, 1995). Fact Finding consists of tasks in which a user is looking for a specific piece of information. Information Gathering consists of tasks in which a user is collecting information, often from multiple sources, in order to write a report, make a decision, or

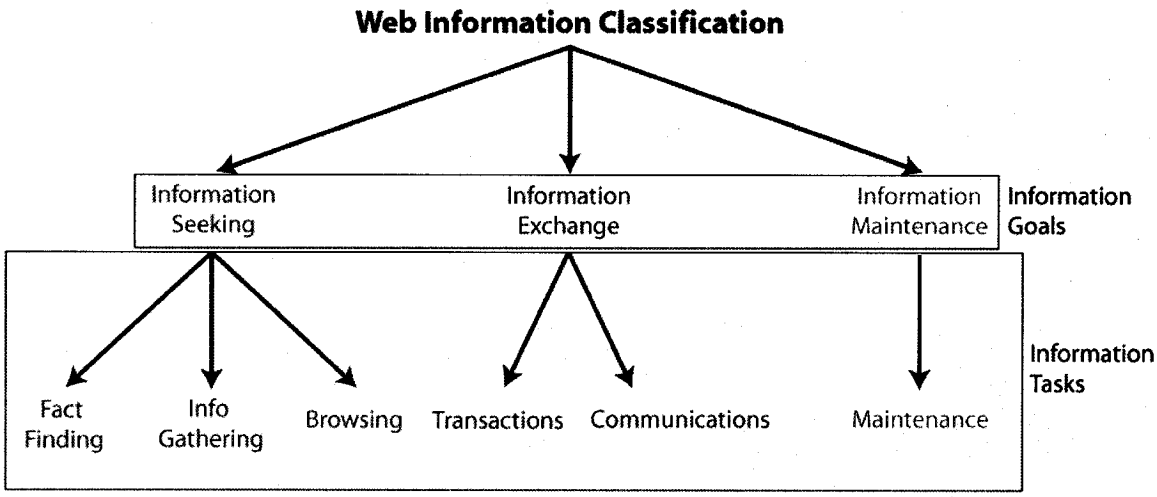

Figure 12. The Web Information Classification.

become more informed about a particular topic. Browsing consists of tasks in which there is no specific information goal in mind other than perhaps entertainment or to see what is new.

Information exchange tasks consist of Transactions and Communications. These are tasks in which the user's goal is to exchange information in a web-based setting. Transactions consist of tasks in which an online action takes place, such as banking or a web purchase. Communications consist of tasks that facilitate web-based communication, such as email, online bulletin boards, or web-based publishing such as blog postings.

Information maintenance tasks are the tasks which were classified as Other by our participants and as "Housekeeping" by Sellen et al. (2002). Maintenance tasks generally consist of visits to web pages with the goal of maintaining web resources, such as to ensure that the content appears as they should, that links are working properly, as well as updates to user profiles. Most tasks of this nature were observed during the field study when participants were updating or creating new web pages.

This classification has not been built to the exclusion of previous work, but instead has incorporated the components of many previous models and frameworks. In addition, this classification also encompasses other non-information seeking tasks. The collection of six tasks (Fact Finding, Information Gathering, Browsing, Communications, Transactions, and Maintenance) closely mirrors the work of Sellen et al. (2002), although our research was conducted with a slightly different focus. Our classification expresses user activities in terms

of web information goals and provides a validation of Sellen et al.'s (2002) previous framework. Whittaker, Terveen, and Nardi (2000) state that researchers often tend to conduct a small number of pioneering studies within a task domain, with little or no future follow-up by other researchers. It is important that researchers continue to validate and iterate on previous studies of user activity on the Web.

5.4 Implications

We have examined how participants used the features of their web browsers to engage in information seeking tasks on the Web and the results of our analysis suggest that participants interacted differently with their web browsers across the different information seeking tasks. Our results provide insight into how web browsers can better support these tasks as well as directions for future research. These findings are not strictly tied to web browsers but also provide insight into how future information systems may better support users during their information seeking tasks. In this section we discuss the implications of this research, with respect to dominant task attributes, repeated tasks, complex information seeking tasks, the history function, and window management.

5.4.1 Dominant Task Attributes

We have examined information seeking tasks in the context of two of the most dominant task attributes: use of browser functions and search versus revisitation. In Table 10, we plot the tasks on these two dimensions and see that few browser functions are used during the non search-based tasks, leaving one quadrant open. This raises the question of whether this is due to an absence of browser functions that support Browsing and Transactions, or because the functionality is simply not required during these tasks. Given the dynamic nature of web pages often viewed during Transactions (e.g., banking, travel bookings), it was surprising that we did not observe more instances of printing, saving, and copying of information. However, it could be the case that participants did use these functions because they were provided by the online applications themselves (using Java Script or Ajax) and therefore not detected by our logging tool. This is a potential topic for future research and could be investigated through the use of new logging tools that allow researchers to log Ajax-based interactions (Atterer, Wnuk, & Schmidt, 2006).

Table 10. Tasks were plotted on two dimensions: use of browser functions and search versus revisitation.

	Search	Revisitation
Low use of browser functions	Fact Finding	Transactions Browsing
High use of browser functions	Information Gathering	

5.4.2 Repeated Tasks

While we observed the highest number of repeated tasks and revisitation across Browsing and Transactions, revisitation occurred across all tasks. The nature of the revisitation differed according to the underlying task type. During Fact Finding task sessions, we observed that participants engaged in repeated tasks in order to monitor new information, re-find previously found information, and to continue with variants of a previous task. During Information Gathering task sessions, tasks were typically repeated when participants were continuing an unfinished task. Repeated tasks that occurred during Browsing and Transaction task sessions appeared to be primarily due to monitoring of particular web sites. Each of these different types of tasks requires different support. While web browsers and information systems may not be able to reliably anticipate what information users may want to re-find during subsequent Fact Finding task sessions, improved history functions (as discussed in Section 5.4.4) may better support this behaviour. Support for repeated Information Gathering sessions conducted to continue a task could be provided through saved sessions and representations of previous web browser interactions. These finding provide for support for navigation systems, such as WebScout (Milic-Frayling, Sommerer, & Rodden, 2003), which offer an archive of a user's previous navigation events and search histories. Repeated task sessions occurring as the result of Browsing and Transactions could be better supported by recognizing the repetitive and habitual nature of these tasks.

5.4.3 Complex Information Seeking Tasks

Information Gathering was the most complex task type in terms of web browser interactions that we considered. We observed the highest number of web browser interactions during

Information Gathering tasks, including functions such as copy/cut/paste, find, and the creation of bookmarks. In addition to the use of these browser functions, Information Gathering tasks had the largest number of pages viewed, were the longest in duration, and were often search-based. While these browser interactions in themselves are not particularly complex, the combination of these interactions contributes to the overall complexity of the task. A visual representation of traces of previous web browser interactions may help users to work more efficiently during a task session as well as during future task session.

5.4.4 The History Function

Over the course of the study, we did not observe any use of the history function even though participants had access to their usual IE history during the study (i.e., history collected before and during the study). This confirms previous research that reported little use of the history function (Aula, Jhaveri, & Käki, 2005; Tauscher & Greenberg, 1997). While researchers have investigated how to better represent the large number of previously viewed pages in a way that is meaningful to users (Ayers & Stasko, 1995; Kaasten, Greenberg, & Edwards, 2002), commercial web browsers have done little to ameliorate this problem. In the post-session questionnaires, many of our participants reported they found it difficult to find previously visited URLs through the history function and only used it as a last resort.

During Information Gathering task sessions, we observed a large number of pages viewed during a single task session (34.5 on average). Users can quickly accumulate a large number of history entries during a single task session; web browsers and information systems should provide mechanisms that allow users to easily revisit any of the pages viewed during the session. It is apparent from the current implementations of the history function that a simple listing of page titles and URLs is not sufficient to allow users to re-find previously viewed pages. More research is needed to determine how to provide more effective representations of previous visited pages. As we will discuss in Chapter 6, a history function that better supports individual differences among users may be more effective and appeal to a wider variety of web users. Alternatively, it may also be worthwhile to explore history entries tagged with an automatically generated task description based on the content of the visited

pages. Work by Komlodi, Soergel, and Marchionini (2006), as presented in Chapter 2, has investigated how to support information seeking through search histories.

5.4.5 Window Management

There has been little research examining the number of web browser windows viewed during a web session. In this research, we have examined differences in the number of windows that were opened across task sessions. We observed only a small number of windows opened across all information seeking tasks, ranging from an average of 2.28 windows during Information Gathering task sessions to 1.43 during Browsing task sessions. This result was surprising in that we expected to observe a much larger number of windows opened during more complex tasks, particularly during Information Gathering task sessions. It could be the case that users typically employ a single browser window per task, opening a window for each concurrent task session. Qualitative user reports from previous research have alluded to task dependant windows management strategies (Hawkey & Inkpen, 2005b; Weinreich, Obendorf, Herder, & Mayer, 2006). Alternatively, the number of windows opened could be influenced by the use of laptop users, who traditionally have smaller amount of screen real estate and may be reluctant to open a large number of browser windows. A wider survey of users is needed in order to better understand how browsers and information systems can better support windows management during information seeking tasks. In particular, the role of tabbed browsing on windows management during different information seeking task sessions must be explored. Preliminary findings from two previous studies (Aula, Jhaveri, & Käki, 2005; Weinreich, Obendorf, Herder, & Mayer, 2006) suggest that tabs are being used to view and organize multiple web pages within a search or task session.

5.5 Study Limitations

It is important that we also acknowledge the limitations of this study. We used a convenience sample consisting of university students, meaning that we cannot expect our results to be generalizable to all populations. Instead, the results of this study provide insight into how skilled web users conduct their information seeking tasks on the Web.

When designing this study, we accepted several tradeoffs, one of which was a short duration of observation (one week). From a pragmatic viewpoint, it would not have been feasible to expect users to provide detailed descriptions of the web usage for extended periods of time. Although this means we likely captured more habitual tasks and a smaller number of "new" or "one-off" tasks, in choosing this strategy we gained the ability to gather very detailed task information. The primary benefit of this methodology design was that we were able to obtain a relatively realistic view of the participants' everyday web use annotated with task information. We observed participants working with their own navigation mechanisms (bookmarks, history, toolbars, etc.) and undertaking tasks that were not motivated by a researcher.

Requiring users to annotate their web usage daily and use a custom web browser had the potential to reduce the naturalness for which we were striving. The post-session questionnaires asked participants if having to record task information changed the way they usually work on the Web and the median participant response was, "a little". When asked if the custom web browser used in the study changed the way they usually work on the Web, the median participant response was again, "a little". Figure 13 displays the distribution for the responses. However, this data is subjective and does not provide insight into how the study may have impacted participants' behaviour on the Web. A more objective measure at our disposal is the number of pages viewed during the study in comparison with previous week long field studies (with similar sample populations) examining user behaviour on the Web (Hawkey & Inkpen, 2005a; Hawkey & Inkpen, 2006). Through this comparison, we learn that our participants viewed approximately 30% less web pages. This may indicate that we only received snapshots of participants' usage on the Web and that they may have used an alternate browser in instances where they perhaps became tired of annotating their data or were viewing sensitive information.

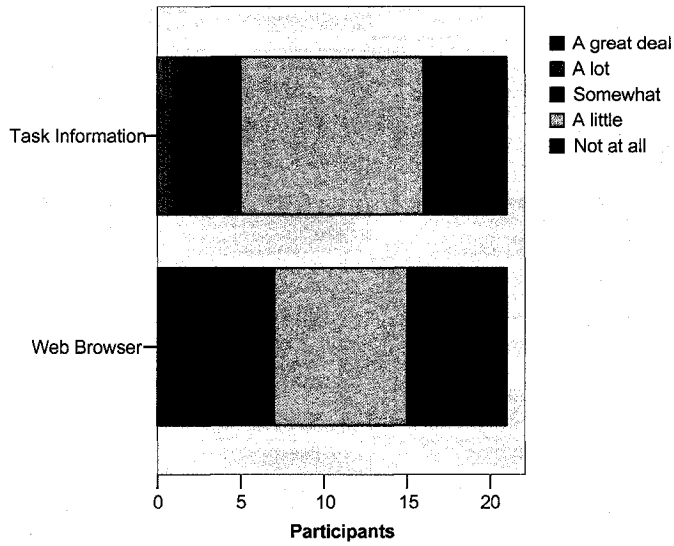

Figure 13. Participant responses to whether having to record their task information and use a custom web browser impacted how they normally work on the Web.

One aspect that we could not explore was task switching. Some participants reported that using the task toolbar to annotate their web information influenced their usual task switching habits. For instance, one participant reported that instead of switching between multiple tasks, she would sometimes fully complete one task before beginning a new task because this would then minimize the amount of task information updates required.

5.6 Summary

In this chapter, we examined how participants interacted with their web browsers across different information seeking tasks. Within each type of task (Fact Finding, Information Gathering, Browsing, and Transactions) we found several distinguishing characteristics. In particular, Information Gathering was the most complex task. On average, participants spent more time completing this task type, viewed more pages, and used the web browser functions most heavily, indicating more research is needed to support users in their Information Gathering tasks. We also gained a better understanding of the role of Transactions within our participants' web usage and observed that Transactions accounted for a large portion of their web use, primarily due to web-based email. Overall, we observed that participants used their web browsers to engage in a mix of task types on a regular basis.

Based on these findings, we have provided implications for the future support of information seeking on the Web, as well as direction for future research in this area.

Using the task descriptions and categorizations collected during the study, as well as previous work, we have developed the Web Information Classification. This classification describes three main web information goals (information seeking, information exchange, and information maintenance) and includes the following web information tasks: Fact Finding, Information Gathering, Browsing, Transactions, Communications, and Maintenance.

The data and findings from this study, which we build upon in the ensuing chapters, form the foundation of the thesis research. Data collected during this field study is used to explore a more detailed understanding of the use of web browser navigation mechanisms in Chapter 6. The Web Information Classification serves as a basis for the research presented in both Chapter 7 and Chapter 8, which focus on monitoring activities within web information tasks.

Chapter 6

The Impact of Task on Web Navigation

In the previous chapter we presented detailed statistical analyses of the differences in web browser interactions across tasks, based on the data collected during a week long field study. In this chapter we present the results of analyses from the same data set, in which we examined how factors such as task and individual differences influence the use of different web browser navigation mechanisms (e.g., hyperlinks, bookmarks, auto-complete).

Since its inception, the ways in which users interact with and navigate the Web have been shaped by the set of navigation mechanisms provided by standard web browsers, such as back and forward buttons, bookmarks/hotlinks, history, and more recently, auto-complete and search toolbars. At present, the research community has a general understanding of the usage frequencies of most navigation mechanisms (Catledge & Pitkow, 1995; Milic-Frayling, Sommerer, & Rodden, 2003; Tauscher & Greenberg, 1997). However, general usage frequency in itself does not provide a complete picture of the usefulness and effectiveness of a given navigation mechanism. Individual differences have been found to play a role in the navigation (Herder & Juvina, 2004) and organization strategies (Teevan, Alvarado, Ackerman, & Karger, 2004) of users on the Web.

We expect that factors such as task and individual differences may also play a role in the use of navigation mechanisms. In order to design more effective mechanisms, it is important that designers and researchers have an understanding of the factors that influence the use of current navigation mechanisms. For instance, when it is reported that a particular navigation

mechanism is used infrequently, in comparison with the most commonly used methods, there may be several possible reasons: *Is the navigation mechanism simply ineffective? Is the mechanism only appropriate for certain tasks? Or is the navigation mechanism preferred by a certain type of user?*

The key contribution of the results presented in this chapter is a characterization of how task and individual differences influence the use of web browser navigation mechanisms. This understanding is used to provide design implications for the future design and evaluation of web browser navigation mechanisms.

6.1 Results

In this section, we first present participants' overall use of web browser navigation mechanisms. We then explore how task and individual differences influence the use of different navigation mechanisms. In order to explore an overview of the data and relative trends, we present only descriptive statistics.

6.1.1 General Use

As previously stated, participants viewed a total of 13,498 web pages (mean = 642.8, range = 98-1733) during the week long field study. Table 11 displays the overall use of each navigation mechanism across all participants. The most common methods of web navigation were hyperlinks (41.7%), 'other' navigation[4] (23.6%), and the back button (18.9%). These three navigation mechanisms were used by all participants and together accounted for 84.2% of all navigation during the study.

The use of bookmarks (2.6%), the Google toolbar (1.3%), new window (5.2%), and typed-in URLs (4.5%) was relatively minimal, when compared with the above mentioned methods of navigation. In addition, we observed that several navigation mechanisms accounted for less than 1% of all use: auto-complete (0.9%), back drop-down menu (0.21%), forward button

[4] As stated in Chapter 4, navigation was categorized as 'other" when we could detect that a high-level document complete event fired (i.e., a single page loaded) but could not identify the direct source of the navigation event.

(0.7%), home button (0.1%), reload button (0.1%), and select URL (0.4%). We did not observe any use of the history or forward drop-down menus.

Table 11. Displays the overall use of each navigation mechanism, % of new task session navigation, % of within task session navigation, and the % of participants who were observed using the navigation mechanism.

Navigation Mechanism	Total Use	Task Session		% of Participants Who Used Mechanism
		New	Within	
Auto-Complete	120 (0.9%)	67.5%	32.5%	85.7%
Back Button	2545 (18.9%)	1.1%	98.9%	100%
Back Menu	21 (0.2%)	0.0%	100%	33.3%
Bookmarks	356 (2.6%)	84.6%	15.4%	71.4%
Forward Button	88 (0.7%)	0.0%	100%	66.7%
Forward Menu	0	—	—	—
Google toolbar	171 (1.3%)	63.2%	36.8%	76.2%
History	0	—	—	—
Home Button	13 (0.1%)	69.2%	30.8%	33.3%
Hyperlinks	5625 (41.7%)	1.7%	98.3%	100%
New Window	707 (5.2%)	12.7%	87.3%	90.5%
Reload Button	12 (0.1%)	25.0%	75.0%	28.6%
Select URL	53 (0.4%)	64.2%	35.8%	19.0%
Typed-in URL	601 (4.5%)	65.1%	34.9%	100%
Other	3186 (23.6%)	1.5%	98.5%	100%

6.1.2 Task Sessions

We examined the use of navigation mechanisms within the context of a task session. As previously stated, a task session is defined as a period of continuous web usage annotated with the same task information and with no break in usage greater than 25.5 minutes.

Overall, we observed a total of 1192 task sessions (mean = 56.8, range = 16-140) with the following breakdown: Fact Finding (18.3% - 218/1192), Information Gathering (13.4% - 160/1192), Browsing (19.9% - 237/1192), Transactions (46.7% - 557/1192), and Other (1.7% - 20/1192). Due to the relatively small number of task sessions classified as Other, we did not consider this task in our analysis.

At first glance, the usage frequencies of navigation mechanisms such as hyperlinks and the back button appear to dwarf the usage of other methods such as bookmarks and auto-complete. However, when we examined all navigation events in the context of a user's task session, we saw a division between two groups of navigation mechanisms. The first group, which we refer to as New Task Session (NTS) navigation mechanisms, were employed either when initiating a new task session or when changing navigation strategies within a session. The second group, referred to as Within Task Session (WTS) navigation mechanisms, were the primary means for navigation within a task session. A breakdown of the session navigation type (NTS and WTS) is shown for each navigation mechanism in Table 11.

6.1.2.1 New Task Session (NTS) Navigation

Web browser navigation methods categorized as NTS navigation mechanisms consisted of those commonly used to initiate a new task session: auto-complete, bookmarks, the Google toolbar, the home button, selecting a URL from the address bar drop-down menu, and typed-in URLs. Figure 14 shows the breakdown of navigation mechanisms used to initiate all 1192 task sessions. Participants most commonly used typed-in URLs (32.8%) and bookmarks (25.3%) to initiate new task sessions.

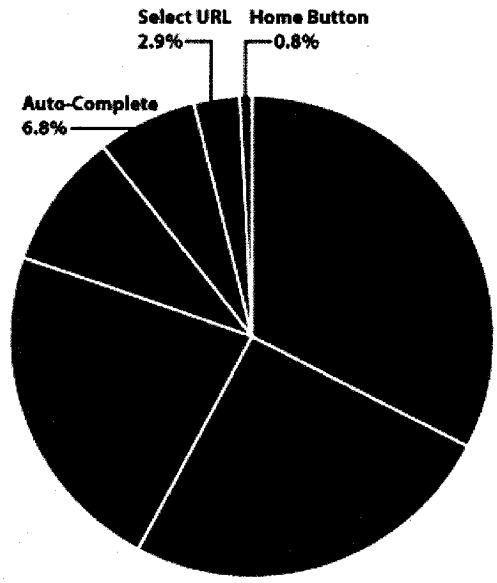

Figure 14. The breakdown of navigation mechanism to initiate new task sessions.

While NTS navigation mechanisms were primarily used to initiate new task sessions, their use was not exclusive. Task sessions were initiated by NTS navigation mechanisms in 77.7% of all tasks. As shown in Table 11, the use of these mechanisms to initiate new tasks ranged from 64.2% (select URL) to 84.6% (bookmarks).The bar chart in Figure 15 shows the breakdown of use (i.e., NTS vs. WTS navigation) for NTS navigation mechanisms.

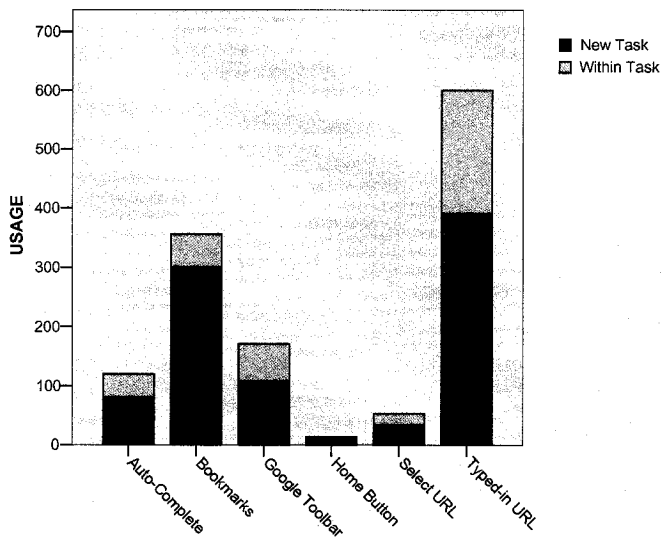

Figure 15. The breakdown of use for NTS navigation mechanisms.

The use of NTS navigation mechanisms was not consistent across participants; the only NTS mechanism used by all participants was typed-in URLs. The NTS navigation mechanisms, when not used to initiate a new task session, were typically used within a session to either branch off to a new web site (e.g., moving from one news web site to another during a Browsing task) or to return to a previously visited page (e.g., returning to Google within a Fact Finding task).

6.1.2.2 Within Task Session (WTS) Navigation

Web browser navigation methods characterized as WTS navigation mechanisms consisted of the back and forward buttons and drop-down menus; hyperlinks; new windows; the reload button; and 'other' navigation. Figure 16 shows the overall breakdown of these navigation mechanisms to navigate within a task session. Participants most commonly navigated within a task session using hyperlinks (44.9%), 'other' navigation (25.5%), and the back button (20.4%).

These mechanisms were used for almost all navigation within a task session, accounting for 96.8% of all WTS navigation. The bar chart in Figure 17 shows the breakdown of use (i.e., NTS vs. WTS navigation) for these navigation mechanisms. As shown in Table 11, the use

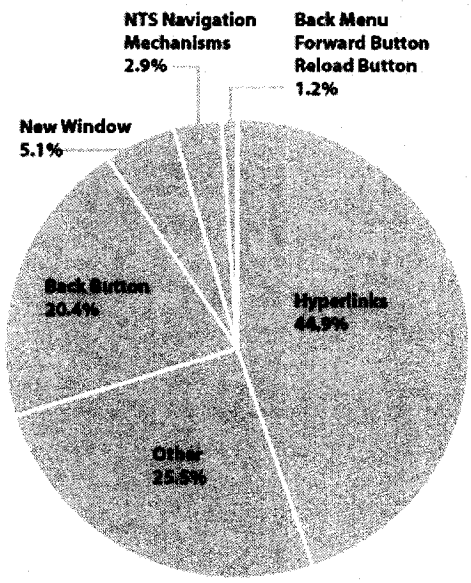

Figure 16. The breakdown of navigation mechanisms used to navigate within a task session.

of each navigation mechanism for WTS navigation ranged from 75% (reload button) to 98.9% (back button). We observed use of the three main navigation mechanisms (hyperlinks, other, back button) by all participants, which is not surprising given that hyperlinks are a fundamental aspect of hypertext.

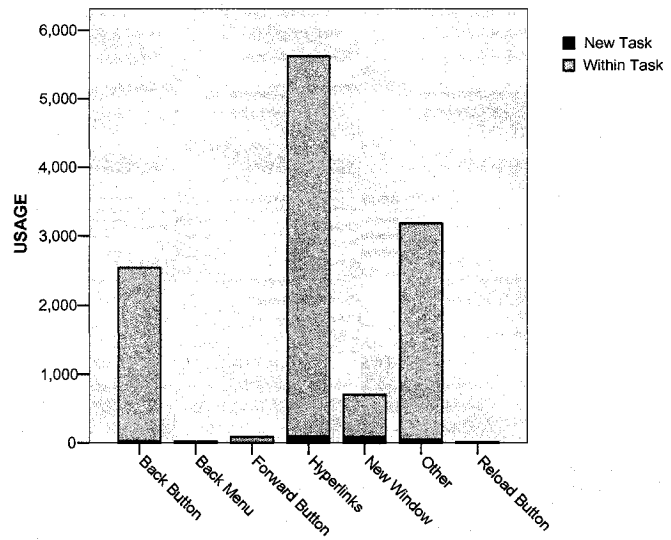

Figure 17. The breakdown of use for WTS navigation mechanisms.

6.1.3 Task Type

We investigated whether the type of task played a role in the use of navigation mechanisms by examining the mechanisms used to initiate and navigate within Fact Finding, Information Gathering, Browsing, and Transaction task sessions. Table 12 displays the proportions of navigation mechanisms used to initiate new task sessions.

With the exception of Transactions, the use of typed-in URLs was the most common method to initiate new tasks. Typed-in URLs were used to initiate 33.5% of Fact Finding tasks, 26.3% of Information Gathering tasks, 30.8% of Browsing tasks, and 34.8% of Transaction tasks. Bookmarks were the most commonly used navigation mechanism to initiate Transactions (35.9%), closely followed by typed-in URLs. For Browsing tasks, the second most common navigation mechanism was bookmarks (21.1%), followed by URLs selected from the drop-down address menu (10.1%). The Google toolbar was the second

most common way to initiate both Fact Finding (23.4%) and Information Gathering (25.6%) tasks.

Table 12. Proportion of NTS navigation mechanisms used to initiate a new task session for each task type.

	Fact Finding	Information Gathering	Browsing	Transactions
Auto-complete	6.9%	16.3%	6.3%	3.9%
Bookmarks	14.7%	11.9%	21.1%	35.9%
Google toolbar	23.4%	25.6%	5.5%	.5%
Home button	.5%	.0%	.8%	.9%
Select URL	1.4%	.0%	10.1%	1.3%
Typed-in URL	33.5%	26.3%	26.3%	34.8%
Non-NTS navigation mechanisms	19.6%	19.9%	29.9%	22.7%
Total	**100%**			

Table 13 displays the proportions of navigation mechanisms used to navigate within each type of task session. Navigation through interactions with the web page (i.e., hyperlinks, other, and new windows) was common across all tasks and was used to navigate within 73.6% of Fact Finding tasks, 63.4% of Information Gathering tasks, 73.4% of Browsing tasks, and 94% of Transaction tasks. The use of the back button was greatest within Information Gathering tasks (32.3%) and was greatly reduced within Transactions (4.5%).

Table 13. Proportion of WTS navigation mechanisms used to navigate within a task session for each task type.

	Fact Finding	Information Gathering	Browsing	Transactions
Back button	20.6%	32.3%	20.2%	4.5%
Back menu	.3%	.3%	.1%	.0%
Forward button	.5%	1.4%	.4%	.1%
Hyperlinks	43.4%	44.9%	54.3%	40.2%
New window	5.4%	3.4%	3.8%	7.9%
Other	24.8%	15.1%	15.3%	45.9%
Reload button	.0%	.1%	.0%	.1%
Non-WTS navigation mechanisms	5.0%	2.5%	5.9%	1.3%
Total	**100%**			

6.1.4 Individual Differences

We investigated whether individual differences played a role in the use of NTS navigation mechanisms. Using a k-means cluster analysis, participants were clustered based on their use of the most commonly used NTS navigation mechanisms while initiating a new task session. Figure 18 shows the four resulting clusters. The first three clusters (c1, c2, c3) consisted of 76.2% (16/21) of the participants and represented those who appeared to have a dominant method of navigation. The fourth cluster consisted of the participants (23.9% - 5/21) who did not display a dominant method of navigation.

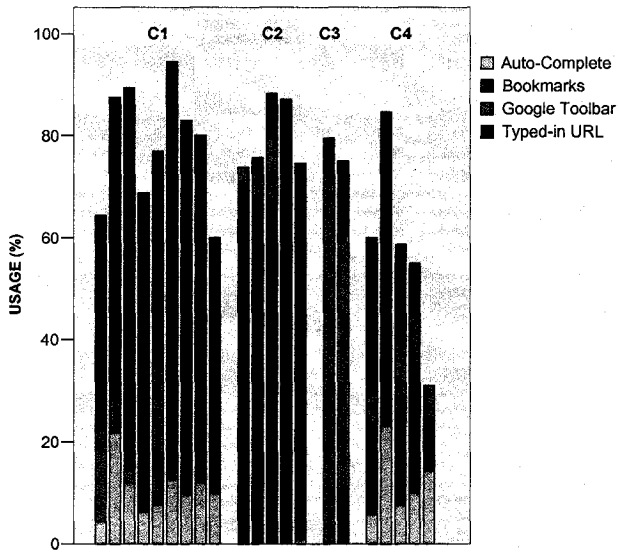

Figure 18. Participants were clustered according to their use of the four most frequently used navigation mechanisms for initiating a task session.

The first cluster (C1) made up of nine participants, consisted of the group of participants who primarily choose to directly type-in URLs to initiate a new task session. The use of typed-in URLs for initiating a new task ranged from 42.2% to 80.0% across members of the cluster. All participants in this group were also users of the auto-complete function for navigation.

The second cluster (C2) was made up of five participants and consisted of those who primarily used bookmarks to initiate a new task session. The use of bookmarks for initiating a new task ranged for 55.7% to 70.5% across members of this group. This cluster differed from the other three clusters in that *all* of the members reported they did not use a

secondary computer, meaning that frequent users of bookmarks did not move between machines. Participants were also very habitual in how they used their bookmarks. From this group, 4/5 participants used a single method to access their bookmarks (either through the side window, drop-down menu or links toolbar). This trend was found across all participants; 80% (12/15) of bookmark users chose a single method to access their bookmarks.

The third and smallest cluster (C3) was made up of two participants and consisted of those whose dominant method of navigation was the Google toolbar. The use of the Google toolbar for initiating a new task was 50.0% and 71.8%, respectively, for the two members of this group.

The fourth cluster (C4) was made up of five participants and consisted of those who did not display a dominant method of navigation. Within this group, participants exhibited varied use of the navigation mechanisms.

6.2 Discussion

Based on the data collected during the field study, we have identified three factors that play a role in the use of navigation mechanisms: task session, task type, and individual differences. Overall, we did not observe striking differences between the usage data we collected and that of previous studies. However, it does appear that the use of the back button has decreased fairly substantially, from 41% in Catledge and Pitkow's (1995) study to 18.9% found in our study. This supports Weinreich et al.'s (2006) recent report of 14.3% back button usage.

We have identified two classes of web browser navigation mechanisms: those used to start a new task session (NTS) and those used within a task session (WTS). This is an important finding for the designers and researchers of web browser navigation mechanisms. We must consider whether use of a navigation method is low simply because the type of mechanism exhibits lower relative use (e.g., NTS vs. WTS) or because the mechanism has an inherent flaw that discourages its use. For instance, when we compare the overall use of bookmarks (2.6%) with that of the back menu (0.15%) it appears that both navigation mechanisms are used relatively infrequently. However, once we examine the navigation mechanisms in their

context of use, we are able to gain a more comprehensive understanding of their use. Within the context of a new task, bookmarks were the second most common method for initiating a new task and accounted for 25.3% of all new task navigation. Our participants commented that they used bookmarks because they were fast and easy to use. The back drop-down menu, a within task navigation mechanism, accounted for only 0.17% of all within task navigation. Participants reported they did not use this navigation method because it is easier to click the back button repeatedly and that it can be difficult to recognize the desired page from the list of pages in the drop-down menu. Three participants also reported they were not even aware of the functionality provided by this navigation mechanism. This example also illustrates how the ineffectiveness of one navigation mechanisms (back menu drop-down) may impact the use of another mechanism (back button).

We observed the navigation mechanisms used by participants while engaging in Fact Finding, Information Gathering, Browsing, and Transaction tasks. The data collected suggests that the task at hand did influence the navigation methods used to initiate new tasks. Although some navigation mechanisms, such as typed-in URLs were frequently used among all tasks, participants tended to also choose methods of navigation that supported the characteristics of the task at hand. For instance, Fact Finding and Information Gathering tasks are often search-based, and this was reflected in the use of the Google toolbar, whose use was minimal among Browsing and Transactions. In Chapter 5, we reported on the repetition rate of web tasks (i.e., tasks that were conducted more than once by a participant during the study). Transactions, due to the large amount of email use, were the most often repeated tasks (95.2%), followed by Browsing (84.4%), Information Gathering (58.8%), and Fact Finding (55.5%). Tasks which were more often repeated exhibited a higher use of navigation mechanisms that support revisitation. Both Browsing and Transactions exhibited a higher use of bookmarks than the less often repeated tasks and we also observed the highest use of select-URLs during Browsing tasks.

We observed four main groups of users based on their choice of a dominant navigation mechanism for initiating new tasks. The first three groups consisted of those who consistently used typed-in URLs, bookmarks, and the Google toolbar. The fourth group consisted of those participants with no one dominant method of navigation. Malone (1983) observed individual differences in how people organize their information and classified

people as either filers or pilers. Filers organize their information in a logical structure while pilers do not adhere to any structured organization. Teevan, Alvarado, Ackerman, and Karger (2004) observed a similar distinction between users when observing their web and email search behaviour. Although further study is required, the use of navigation mechanisms may be coupled with the way in which URLs and web information are stored and retrieved. For instance, bookmarks may be well suited for filers, while a searchable history may be more appropriate for pilers.

6.3 Implications

Based on our findings, we present implications for both the evaluation and future design of web browser navigation mechanisms.

6.3.1 Evaluation and Interpretation in Context

We have demonstrated that the use of current web browser navigation mechanisms must be interpreted in context. When interpreting the effectiveness of navigation mechanisms, we cannot make direct comparisons between the use of NTS and WTS navigation mechanisms. Similarly, when evaluating new navigation mechanisms, it is important to understand how the use of a given mechanism may be influenced by the type of task in which a user is engaging. Therefore, researchers should give careful consideration to the types of tasks users are assigned to perform when evaluating new navigation mechanisms in a laboratory environment. It is also important to recognize that when evaluating NTS navigation mechanisms, users may not have as many opportunities to use these mechanisms when compared to WTS navigation mechanisms. Therefore, tasks and scenarios used in evaluations should provide participants with ample opportunities to use the navigation mechanisms.

6.3.2 Supporting Task Characteristics

We theorize that web browser navigation mechanisms can provide more effective and efficient navigation by better supporting the underlying characteristics of the tasks in which users engage on the Web. We observed two principal underlying task characteristics during

our study: search and revisitation. Fact Finding and Information Gathering were primarily search-based tasks while Browsing and Transaction tasks often consisted of repeated visits to web sites.

Standard web browser navigation mechanisms currently provide limited support for search-based tasks through search toolbars, such as the Google and Yahoo! Search toolbars. Search-based tasks (i.e., Fact Finding and Information Gathering) account for a considerable portion of most users' web use. Therefore, next generation navigation mechanisms should provide support for more efficient navigation and interaction with search engines that can also be customized to users' specific interests.

While many NTS navigation mechanisms currently support revisitation, their functionality is limited. We observed that participants revisited web pages in order to both re-find static information as well as to monitor dynamic information. Users' navigation needs may differ depending on the type of revisitation and information being sought and next generation navigation mechanisms should aim to better support this behaviour.

6.3.3 Supporting Individual Differences

We observed that individual differences seemed to impact the use of web browser navigation mechanisms. Future navigation mechanisms should be designed with an understanding of the individual differences that exist and provide variations of each mechanism to reflect these differences. For instance, menu-based bookmarks may only appeal to those users who prefer to file their information. For those web users who prefer to type most URLs, they may prefer a version of bookmarks that can be accessed through typed commands (or shortcuts), as opposed to menu interactions. Similarly, for those participants who exhibited heavy use of the Google toolbar, they may prefer to use searchable bookmarks. This is also supported by our observations that dominant bookmark users accessed their bookmarks through a single method (either through the side window, drop-down menu or links toolbar).

The history mechanism has the potential to offer effective methods for web site revisitation. However, as reported in Chapter 5, the history function is rarely used. Many of our participants reported they found it difficult to re-find a previously visited URL through the

history mechanism and only used it as a last resort. A history mechanism that better supports multiple user preferences, such as filing, searching, and typed commands, may better appeal to a variety of web users.

6.4 Summary

Logging task data to examine the use of navigation mechanisms is a largely unexplored area. In this chapter, we have conducted an analysis of the use of web browser navigation mechanisms in the context of task sessions. This analysis identified three factors that influence the selection of navigation mechanisms by the users: task session, task type, and individual differences. These findings have implications for the evaluation and design of new and improved web browser navigation mechanisms.

Chapter 7

An Exploration of Web-based Monitoring in the Context of Task

In Chapters 5 and 6, we examined the types of tasks users engaged in and how they used their web browser to complete these tasks. One open problem identified in Chapter 5 was the role of monitoring within the Web Information Classification. Web-based monitoring occurs when users return to previously viewed web pages to view new or updated information. In advance of the field study, we had expected that monitoring was an independent information seeking task, based on previous research (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001). It became apparent during our pilot study that monitoring was not a distinct information seeking task and overlapped with a variety of other tasks. For example, it was very difficult for participants to distinguish between monitoring and "re-Fact Finding" or "re-Browsing". As a result, monitoring was omitted from the field study since further research was needed to define the nature of web users' monitoring activities.

To better understand web-based monitoring, and its role within the Web Information Classification, we conducted a set of 40 semi-structured interviews. We chose to study a wider range of users than during our previous study and recruited participants from four sample populations: technical students, non-technical students, technical workers, and non-technical workers. The primary contribution of the research presented in this chapter is an examination and characterization of users' web monitoring behaviours within the context of

the Web Information Classification. The results of this study have been used to develop recommendations, both general and task specific, for the design of future web-based monitoring tools.

7.1 Research Questions

To better understand monitoring and its role within the Web Information Classification, we proposed the following research questions:

RQ1: How does monitoring fit into the Web Information Classification? Is monitoring an independent information seeking task or an activity within a task?

RQ2: What are the general characteristics of monitoring activities?

RQ3: What are the characteristics of monitoring activities, by task type?

We hypothesized that monitoring is not an independent information seeking task, but instead an activity within each web information task. Therefore, we expect that different monitoring activities will require different support, depending on the underlying task type.

7.2 Methodology

A respondent research strategy was chosen to study web monitoring in the context of task. In particular, semi-structured interviews were chosen because they allow in-depth exploration of users' behaviours and habits while offering researchers a good balance of control in a casual environment. This research strategy also allowed us to study a larger sample population than in our previous study.

In this section we first describe the participants who took part in the study, followed by the study procedure and the methods of data collection. The study described in this chapter was approved by the Social Sciences and Humanities Human Research Ethics Board at Dalhousie University.

7.2.1 Participants

Forty participants were recruited to take part in hour long *in situ* semi-structured interviews during March and April of 2006. Ten participants were recruited from each of the following sample populations: technical students, non-technical students, technical workers, and non-technical workers. A breakdown of the participants belonging to the four sample populations is shown in Table 14. While users with strong technical backgrounds (i.e., computer science students, graduates, and information technology workers) do not reflect typical web users, we expected that they may have more sophisticated monitoring behaviours. This is a demographic that is often regarded as early adopters of new technology and may well be predictors of future users. We also recruited a mix of students and office workers with non-technical backgrounds (which may reflect behaviours and habits of more typical web users), in order to gain a better overall view of monitoring behaviours on the Web. Participants were paid $10 for taking part in the study.

Table 14. A breakdown of the participants' demographics.

Sample Population	Gender		Median Age Group	Discipline/ Occupation
Technical students	7	3	21-30	Bachelors: Computer Science (1) Masters: Computer Science (5) Ph.D.: Computer Science (4)
Non-technical students	5	5	21-30	Bachelors: Science (1), Commerce (4) Masters: MBA (3), MA (1) Ph.D.: Industrial Engineering (1)
Technical workers	7	3	21-30	Web Developer (2) System Administrator (1) Software Engineer (5) Data Analyst (1) Technical Manager (1)
Non-technical workers	2	8	41 – 50	Cartographer (1) Coordinator (4) CFO (1) Lab Manager (1) Administrative Secretary (1) Consultant (2)
Total	21	19		40

7.2.2 Procedure and Data Collection

Before beginning the interview, participants completed a demographic questionnaire (shown in Appendix D for students and in Appendix E for workers) that was designed to collect their demographic information, current level of web use, and current monitoring activities. During the interview, an interview guide was used to keep the interview on track and to help maintain consistency and reliability in the data (Bernard, 2000). Participants were asked to describe up to three work/school monitoring activities and three personal monitoring activities. For each activity, participants were asked to describe the goal, the type of information being sought, and the following attributes:

1. Use of navigation mechanisms (e.g., bookmarks)
2. User login required
3. Search query required
4. Duration of monitoring activity (responses were categorized as long term, short term, or short term iterations [i.e., short term tasks that reoccurred])
5. Frequency of monitoring activity (response were categorized as continually throughout the day, daily, every couple of days, weekly, bi-weekly, monthly, sporadically, and due to an event)
6. Follow-up activities (response were categorized as either electronic, non-electronic, or both)
7. Work/school or personal (since some students worked part-time and some workers were part-time students, we merged "School" and "Work" into a single category)

A full copy of the interview is available in Appendix F. At the end of the interview, participants were shown two existing monitoring tools. They were asked if they could see themselves using the tool and to describe their impressions of the tool.

All interviews were audio recorded and the recordings were transcribed by an individual not associated with the study or our research lab. Using the transcripts, each reported monitoring activity was coded for task type (Fact Finding, Information Gathering, Browsing, Communications, Transactions, Maintenance) based on the participant's description of the activity and the information being monitored. In order to develop a characterization of the different monitoring activities, the transcripts were also coded for the previously mentioned attributes (1-7).

7.3 Results

A total of 173 monitoring activities were reported. On average, participants reported four monitoring activities (range = 2 - 6). With the exception of five activities, the reported monitoring activities were easily categorized into the six web information tasks. This supports our hypothesis that monitoring is not a distinct information seeking task but instead an activity within web information tasks. Figure 19 displays the breakdown of all reported monitoring using the tasks from the Web Information Classification. The most common type of monitoring activity was Browsing (74/173), followed by Fact Finding (42/173), Information Gathering (18/173), Transactions (16/173), Communications (11/173), and Maintenance (7/173). The five monitoring activities that did not belong to a single category consisted of a combination of different monitoring activities (5/173).

We next present a characterization of monitoring within each of the web information tasks. Table 15 displays for each task type, the use of navigation mechanisms, the anticipated duration of the monitoring activity, and the frequency of the monitoring activity. The most common responses are highlighted in the table. Appendix G shows the full list of monitoring activities reported, categorized by task.

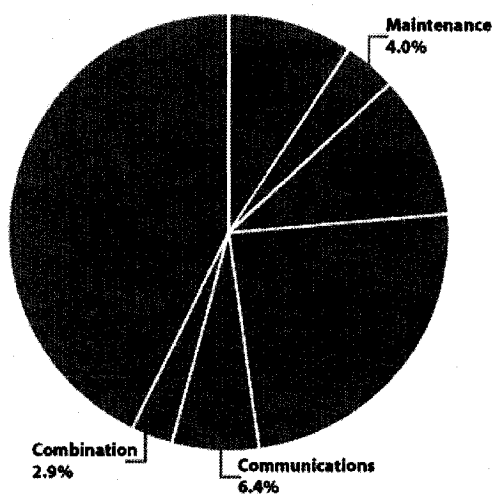

Figure 19. Categorization of all reported monitoring activities.

Table 15. Displays the breakdown of navigation mechanisms, frequency, and duration of monitoring activities, by task type.

	Browsing	Fact Finding	Info. Gathering	Tran-sactions	Communi-cations	Main-tenance
Total Tasks	74	42	18	16	11	7
Navigation Mechanism						
Auto-Complete	5 6.8%	2 4.8%	3 16.7%	2 12.5%	1 9.1%	0
Bookmarks	41 55.4%	17 40.5%	7 38.9%	11 68.8%	3 27.3%	3 42.9%
Emailed Link	4 5.4%	4 9.5%	2 11.1%	0	0	1 14.3%
Homepage	5 6.8%	6 14.3%	1 5.6%	0	1 9.1%	2 28.6%
Search Engine	1 1.4%	5 11.9%	3 16.7%	0	0	0
Select URL	0	0	0	1 6.3%	1 9.1%	1 14.3%
Typed-in- URL	17 23.0%	8 19.0%	2 11.1%	1 6.3%	2 18.2%	0
Other	1 1.4%	0	0	1 6.3%	3 27.3%	0
Duration						
Long Term	67 90.5%	30 71.4%	14 77.8%	14 87.5%	11 100%	7 100%
Short Term	4 5.4%	7 16.7%	2 11.1%	1 6.3%	0	0
Short Term Iterations	3 4.1%	5 11.9%	2 11.1%	1 6.3%	0	0
Frequency						
Continual	7 9.5%	4 9.5%	1 5.6%	0	9 81.8%	3 42.9%
Daily	32 43.2%	14 33.3%	4 22.2%	6 37.5%	2 18.2%	1 14.3%
Couple Days	15 20.3%	5 11.9%	4 22.2%	3 18.8%	0	1 14.3%
Weekly	5	6 14.3%	0 .0%	1 6.3%	0	1 14.3%
Bi-Weekly	5 6.8%	1 2.4%	1 5.6%	1 6.3%	0	0
Monthly	3 4.1%	1 2.4%	1 5.6%	1 6.3%	0	1 14.3%
Sporadic	4 5.4%	1 2.4%	1 5.6%	3 18.8%	0	0
Event	2 2.7%	10 23.8%	4 22.2%	1 6.3%	0	0
Unknown	1 1.4%	0	2 13.2%	0	0	0

7.3.1 Browsing

Monitoring activities were categorized as Browsing when participants reported they were monitoring information to see what was new or to stay up-to-date. Browsing was the most common type of monitoring activity, accounting for 42.8% (74/173) of all reported monitoring activities. Participants reported that it was a way of relaxing or taking a break. The most common Browsing monitoring was online news (29.7%), followed by blogs and wikis (12.2%), sports news (10.8%), and books/movies/music (8.1%).

Participants reported most commonly navigating to pages they monitored using bookmarks (55.4%) and typed-in-URLs (23.0%). Participants reported that they rarely needed to login to the monitored web pages; only 9.5% of Browsing monitoring required a user login. Similarly, user queries were seldom required; only 2.7% of Browsing monitoring required a user query.

Browsing monitoring differed from many other monitoring activities in that participants reported they often visited several web sites within a single monitoring activity (36.5%). Four participants also reported using some sort of aggregation tool, such as the LiveJournal friends list and personalized homepages (e.g., Google, Yahoo).

Five participants described particular navigational rituals when conducting their monitoring activity. For example, one participant described how she categorized her bookmarks (for pages she monitored) into folders, such as daily, weekly and RSS, on her browser toolbar. Each morning, she would access the pages in her "daily" folder from top to bottom and once a week she would complete a similar ritual for her "weekly" sites.

> I have a daily flag and a weekly-ish flag. I haven't made any more. Anything less than that, I don't bother making a category for it. And the RSS is its own pull-down. Weekly-ish and the RSS stuff.

The second participant reported he conducted his Browsing monitoring in the same way several times during the day, by loading all the pages bookmarked on his bookmarks toolbar and viewing them from the leftmost bookmark to right. Two participants reported that each morning they would open the same group of web pages using the Firefox "Open in Tabs" bookmark functionality. The fourth participant reported he visited the same parts of the

Wikipedia web site in the same order every day. It is worth noting that 4/5 of these participants were from the Technical Student demographic.

Most Browsing monitoring activities were long term tasks (90.5%), meaning they would continue indefinitely, or for as long as they were students or employed in the same capacity (for non-personal monitoring activities). Only a small number of short term (5.4%) and short term iteration (4.1%) tasks were reported. Participants stated they would cease short term monitoring activities because they became bored or tired of the activity. Short term iterations tended to start and stop, depending on the participant's interest in the activity. For instance:

> It was going on for a couple of months while I had extra money and once I spent it, I stopped. It will pick up when I have the need or the money.

Participants reported that a little under half (43.2%) of all Browsing monitoring was conducted on a daily basis followed by every couple of days (20.3%). While most participants were relaxed about the frequency of their Browsing monitoring, some participants reported they tended to be somewhat compulsive about the activity. For instance, two participants described their news reading behaviour and use of Facebook as addictions:

> It's something like an addiction so I have to do it or it feels like something is missing in my day

> It's an addiction [Facebook] so probably like 3 times a day - also part of nightly routine

Participants were asked to describe any follow-up actions based on the monitoring activity. Just over half (56.8%) of the activities were tagged as electronic follow-up and this included actions such as forwarding links, commenting on a blog, an online purchase, or follow-up on a topic using the Web. Only 6.8% of follow-up actions had a non-electronic follow-up action and 2.7% had both an electronic and non-electronic follow-up. A third (33.8%) of Browsing monitoring had no follow-up.

Participants reported that the majority of their Browsing monitoring was personal (75.7%). While we might expect that all Browsing monitoring activities are personal in nature, 23.0%

were work/school related. For example, one participant described how she enjoyed browsing a portal site related to her research, with no particular goal in mind, simply because it was enjoyable. Only a single monitoring activity (1.4%) was related to both Work and Personal monitoring.

RSS feeds are frequently used for syndicating news web sites and blogs. RSS has the potential to support Browsing monitoring since these types of web sites are often viewed during Browsing. Only eleven participants mentioned RSS during the interview and over half (6/11) stated that they did not find RSS useful. The most commonly cited reason for disliking RSS was the lack of pictures/details and "too much information coming at you". Three participants reported they would consider using RSS if it were to be improved. Of the 11 participants who mentioned RSS, three reported they had heard of RSS but did not know what it was or how to use it. Only 2/11 participants reported they were satisfied with their use of RSS to complete their Browsing monitoring.

7.3.2 Fact Finding

Monitoring activities were categorized as Fact Finding when participants reported they were monitoring specific or factual pieces of information, such as updated snippets of information or the appearance of a new file. Fact Finding was the second most common type of monitoring activity, accounting for 24.3% (42/173) of all reported monitoring. One participant described the difference between a Browsing and a Fact Finding monitoring activity, both related to sports; he reported that Browsing monitoring was more of a superficial sweep to see "what's new" while the Fact Finding monitoring was strictly a scan for numbers. The most common Fact Finding monitoring activities were checking for new assignments/grades/notes (19.0%), weather information (16.7%), status (typically of an application such as citizenship, graduate school, or conference student volunteer) (11.9%), and factual sports information (e.g., games scores) (11.9%).

Participants reported that they most commonly navigated to pages they monitored using bookmarks (40.5%) and typed-in-URL (19.0%). Only 33.3% of Fact Finding monitoring required a user login and participants reported they rarely (9.5%) submitted queries to access the updated information.

Almost three quarters (71.4%) of Fact Finding monitoring was categorized as long term. In contrast to Browsing monitoring, participants indicated that short term monitoring activities would end once the desired information appeared or when the information was no longer being updated. For instance one participant reported he had been monitoring Canada's medal count during the Olympics while another was waiting for the release of a particular iPod patch. Participants also reported they completed several iterations of short term tasks (11.9%). Examples of this task type include one participant who reported he would monitor the current price of an e-bay auction whenever he discovered a new item of interest for sale on eBay. Three participants described another task of this nature, which was monitoring details of annual conferences (e.g., submission dates, conference fees).

A third (33.3%) of all Fact Finding monitoring was conducted on a daily basis and 23.8% of Fact Finding monitoring did not occur with any regularity, but instead occurred when a participant was motivated by a particular event. For instance, some participants reported they only monitored the weather forecast if there was a weather event forecast while others stated they would monitor their grades after an exam.

While only four monitoring activities occurred continually throughout the day, three participants reported they felt compelled to engage in the activity. For instance, one participant who reported she checked her grades continually throughout the day stated:

> I would check 10 times a day [when waiting for a grade] (...) It is habit. I just feel like I have to because there might be something there, you know? If I'm on the [University] site it's something I have to do.

Another participant described how often he checked for a decision on his grad school application:

> When I wake up, when I go to bed, after every class, before every class, in the break between classes, if I'm studying every 90 minutes.

Participants were asked to describe any follow-up actions based on the monitoring activity. Almost half (47.6%) of the tasks had electronic follow-up and this most commonly consisted of printing or saving the new information (e.g., new assignment). Non-electronic follow-up was reported in 31.0% of Fact Finding monitoring. This follow-up typically consisted of

"real world" actions, such as dressing according to the weather forecast or deciding what time to watch a sporting event on TV. Finally, 4.8% involved both electronic and non-electronic follow-up, and 16.7% had no follow-up.

The reported Fact Finding monitoring activities were closely split between personal and work-related. Participants reported that 57.1% of their Fact Finding monitoring was work related while 42.9% was personal.

7.3.3 Information Gathering

Monitoring activities were classified as Information Gathering when participants were gathering updated information for research or decision-making purposes. Information Gathering was the third most common type of monitoring activity and accounted for 10.4% (18/173) of all reported monitoring. The most common monitoring activity was keeping up-to-date on research literature (38.9%), followed by monitoring of business commodities (e.g., currency, produce) (11.1%), and employment opportunities (11.1%).

Participants reported that they most commonly navigated to pages they monitored using bookmarks (38.9%), search engines (16.7%), and auto-complete (16.7%). Participants were rarely (5.6%) required to login to complete their monitoring activities. However, 38.9% reported they submitted search queries to access the updated information.

Participants reported that 77.8% of their Information Gathering monitoring activities were long term tasks. Two of the Information Gathering monitoring activities were categorized as short term: looking for new employment and looking for a new house. Participants reported these tasks would end once they had reached their goal. Interestingly, one employment monitoring activity was categorized as a long term task because the participant reported that he anticipated that this task would continue indefinitely. However, at this point, the task may become more of a Browsing monitoring activity. Finally, two Information Gathering monitoring activities were labelled as short term iterations. For example, one participant reported he would monitor iPod forums to conduct research before purchasing a new iPod accessory. The frequency of Information Gathering monitoring was split between occurring daily (22.2%), every couple of days (22.2%) and triggered by a particular event (22.2%).

In terms of follow-up actions, participants reported that all Information Gathering activities entailed a follow- up action. The follow-up was electronic for 61.1% of all Information Gathering monitoring and non-electronic for the remaining 38.9%. The most common types of electronic follow-up included printing and saving documents and applying for jobs. Non-electronic follow-up activities were often related to a purchase, such as particular commodities or product.

While Information Gathering monitoring is typically thought to be of non-personal nature, almost a third (27.8%) was tagged as personal. For instance, one participant reported weekly monitoring activities in order to place Pro Line sports bets. Another personal Information Gathering monitoring activity included a participant who took her movie selection seriously and regularly studied movie reviews and news in order to decide what movie she would see on the weekend.

7.3.4 Transactions

Monitoring activities were characterized as Transactions when participants reported they were monitoring information that would either initiate a new online transaction or was the result of a previous transaction. This activity differs from most other monitoring activities in that the exchange of information is often a one-to-one relationship, meaning that the information has been dynamically created for a specific user by an information provider. Transactions accounted for 9.2% (16/173) of all monitoring. The most common monitoring activity was checking online account balances (e.g., banking, credit cards) (68.8%) and online purchases (12.5%).

Participants reported that they most commonly navigated to pages they monitored using bookmarks (68.8%) and auto-complete (12.5%). In many cases (75.0%), participants were required to login to complete their monitoring activities. Participants were rarely (12.5%) required to submit search queries to access the updated information.

Participants reported that most (87.5%) of their Transactions monitoring was long term and these were typically banking related monitoring activities. Over a third (37.5%) of Transaction monitoring occurred on a daily basis, as part of a daily routine. Additionally,

participants reported that 18.8% of Transactions monitoring occurred every couple of days and another 18.8% occurred sporadically.

In terms of follow-up actions, participants reported that most Transactions monitoring entailed a follow-up activity. This follow-up activity was electronic for 50.0% of tasks, non-electronic for 25.0% of tasks, and both electronic and non-electronic for 18.8% of tasks. Only a single task (6.3%) had no follow-up. The most common types of electronic follow-up consisted of an online transaction, such as paying bills or purchasing airfare. Non-electronic follow-up activities often entailed telephoning a particular company with respect to the transaction being monitored (e.g., incorrect bank balances). Participants reported that the 81.3% of their Browsing monitoring tasks were personal and 18.8% were work-related.

7.3.5 Communications

Monitoring activities were categorized as Communications when participants reported they were monitoring a web page for new or incoming communications. Only 6.4% (11/173) of the monitoring activities were categorized as Communications. Of these activities, 10/11 (90.1%) consisted of web-based email. All participants were email users and only some chose to report email as a monitoring activity. Therefore, while it is true of all reported monitoring activities, the percentage of Communications monitoring reported is not necessarily representative of the total percentage of Communications monitoring that occurs on the Web. The only non-email monitoring activity was reported by a teaching assistant who monitored a bulletin board for new postings from students as it was the primary method of communication between her and her students.

Participants reported that they most commonly navigated to pages they monitored using bookmarks (27.3%) and through an email notification tool (27.3%). All monitoring required a user login while none of the monitoring activities required users to submit search queries.

All Communications monitoring was long term in nature. Participants reported they engaged in this activity continually throughout the day (81.8%) or daily (18.2%). In terms of follow-up activities, participants reported that all follow-up was electronic and typically consisted of managing and replying to new communications.

Participants reported that 72.7% of Communications monitoring was both personal and work-related. Two participants (18.2%) reported it was strictly personal and a single activity (9.1%) was described as strictly work related (the bulletin board-based communications).

7.3.6 Maintenance

Monitoring tasks were categorized as Maintenance when participants reported they were monitoring a web page for the purpose of maintaining it, such as checking for spelling mistakes, layout problems, or accurate content. Only 4.0% (7/173) of all reported monitoring activities were classified as Maintenance. The most common activity was monitoring university departmental pages for accurate content and to ensure there were no errors or spelling mistakes (42.9%).

While only a small number of Maintenance tasks were reported, participants stated that they most commonly navigated to pages they monitored using bookmarks (42.9%) or the page they monitored was their homepage (28.6%). Only 28.6% of monitored pages required a login and none required search queries.

All monitoring tasks reported by participants were categorized as long term monitoring tasks. Participants reported that 42.9% of their Maintenance monitoring activities were conducted continually throughout the day. The remaining four tasks occurred, daily, every couple of days, weekly, and monthly, respectively. Participants reported that all tasks entailed electronic follow up, which typically involved making changes to the website. In some cases it was the participant who did so, in other cases they contacted the person responsible for maintaining the website. All but one (85.7%) Maintenance monitoring task was work-related in nature.

7.3.7 Combined Monitoring Activities

In addition to the above monitoring tasks, participants also described five tasks (2.9%) that did not fit within a single task category. These remaining tasks were categorized as: Fact Finding & Browsing (1/5); Fact Finding & Information Gathering (1/5); Fact Finding & Communications (1/5); and Fact Finding, Browsing & Maintenance (2/5). One combined monitoring activity was described by a participant who worked for the university's

computing services, as monitoring the University's online computer store for several reasons: to see what was new and new product specials (Browsing); to monitor the price of a particular product, perhaps waiting for a sale (Fact Finding); and to make sure there were no errors on the page (Maintenance).

7.3.8 Monitoring Tools: Forecastfox and Page Update Checker

At the end of the interview participants were shown two brief screen capture videos introducing two Firefox extensions, each of which supports a particular monitoring task. The first tool introduced was Forecastfox (Mozdev, 2004) and allows users to monitor a region's weather through a series of icons positioned in a region of the web browser (e.g., status bar, toolbar). The second tool introduced was Update Page Checker (Mozdev, 2006), which notifies users, through a small icon in the status bar, when a previously tagged web page has updated. Participants were informed that these were freely available Firefox extensions and that they had not been developed by the researcher involved in the study. After viewing the screen capture, participants were asked if they could see themselves using the tool and to describe what they liked and disliked about the tool.

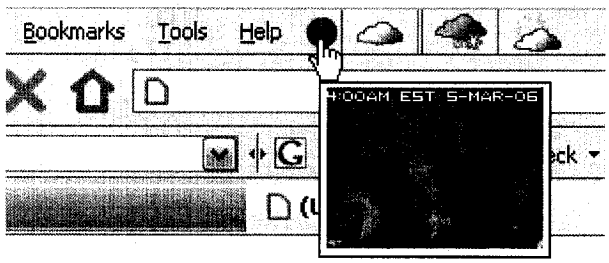

Figure 20. Forecastfox provides constant updated weather information.

7.3.8.1 Forecastfox

When introduced to Forecastfox (shown in Figure 20), participants expressed varying degrees of enthusiasm for the tool. While some participants reported they were currently using it or a similar widget, others remarked they would never use such a tool. Participants expressed three primary concerns regarding the tool: accuracy of the weather data, loss of screen real estate, and the method of notification.

Accuracy of the weather data: Five participants reported they were hesitant to use such a tool because they were not confident the forecast, which does not come from the national meteorological service, would be accurate. For instance, one participant stated:

> *Participant:* So I don't know. Maybe I would use this. I might use this, [but] it depends on where it is from, if it's a reliable source.
>
> *Interviewer:* So it needs to be reliable?
>
> *Participant:* Yeah, like I trust weather.ca. I trust their forecasts, but then I'm wondering where this is coming from?

Screen real estate: Participants voiced opinions that the tool encroached on their web browser real estate. While two participants reported it was "nice and small", six participants expressed concerns with the size and location of the tool.

Method of Notification: Participants were split on Forecastfox's method of information delivery. While three participants reported they liked hovering over the icons to receive more information (i.e., temperature), three other participants reported they did not like this functionality. Two of these participants would rather have the information displayed directly on the icon while the third would rather click to receive the information. Three other participants also expressed that they found the pop-ups annoying, as stated by one participant:

> But what I find annoying are all the pop-ups like at times when I don't need to know what the weather's like.

7.3.8.2 Page Update Checker

When introduced to the Page Update Checker tool (shown in Figure 21), participants had fewer concerns about screen real estate but voiced concerns about the method of notification and back-end functionality. Participants also described tasks for which they would find the tool useful.

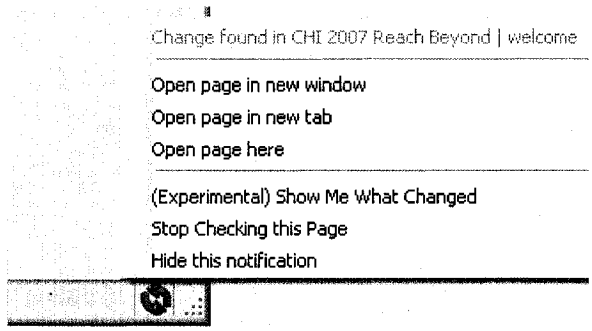

Figure 21. The Page Update Checker extension notifies users of updates to monitored web pages.

Functionality: Six participants questioned the functionality of the tool immediately, by asking if the tool was sensitive to new ads (it is) or just dramatic changes, and whether it is a bother to setup. The technical participants (students and workers) in particular, were very perceptive in their evaluation of the tool:

> I'm wondering how it would work with ads and things like that, like if you have Google ads on your page

Notifications: Seven participants expressed some discomfort with the pop-up notification. Most felt the pop-ups had the potential to be annoying or distracting. Two participants would have preferred email notification; one participant in particular commented that since he worked from two computers (home and work), email would be more convenient since he would not have to install and configure the tool twice. Six participants stated that they were not sure what the icon meant and would prefer different icons for different web sites.

Representative Tasks: While the participants did find several aspects of both the back-end and interface problematic, three participants described current monitoring tasks for which a tool such as Page Update Checker would be useful. The tasks were all Maintenance based monitoring activities.

7.4 Discussion

In this section, we first present an overall characterization of monitoring activities, followed by a discussion of monitoring as a compulsive behaviour and barriers to use of monitoring tools.

7.4.1 General Characterization of Monitoring Activities

Our motivation in conducting this research was to determine how monitoring fits into the Web Information Classification. The results of the interviews support our hypothesis that monitoring is an activity that occurs within each web information task. Browsing was the most common type of monitoring activity, followed by Fact Finding, Information Gathering, Transactions, Communications, and Maintenance. We also observed that some monitoring tasks consisted of a combination of different tasks.

We were also motivated by the need for a better understanding of how users conduct their web-based monitoring activities. While there were notable differences in the types of information being monitored, the characteristics (use of navigation mechanisms, duration, and frequency) of the different monitoring tasks were often similar, with some exceptions.

7.4.1.1 Navigation

In Chapter 6, we examined the use of web browser navigation mechanisms across task type and found that while typed-in-URLs were the most common way to initiate a new task, the choice of navigation mechanisms was dependant upon the task at hand. For instance, the Google toolbar was most commonly used during Fact Finding and Information Gathering and bookmarks were most commonly used during Browsing and Transactions. However, the respondents who took part in the semi-structured interviews reported that bookmarks were the most common navigation mechanism, across all task types. This seems to indicate that it is not only task type which impacts the choice of navigation mechanism, but also whether the task is a monitoring activity. Therefore, despite the recognized issues associated with bookmarks (Abrams, Baecker, & Chignell, 1998; Jones, Dumais, & Bruce, 2002), our research indicates that bookmarks are still a popular navigation mechanism, especially during

monitoring activities. This supports an older finding (Choo, Detlor, & Turnbull, 2000) that monitored web sites are often bookmarked.

7.4.1.2 Queries and Logins

The use of queries and search engines was very common among both Fact Finding and Information Gathering in Chapter 5 ; however, our participants reported very little use of search engines or queries during Fact Finding monitoring. It would appear that users employ a search engine to initially retrieve new information but are likely to create a bookmark if they anticipate returning to the page for monitoring purposes. User logins were not required for most monitoring tasks with the exception of information exchange monitoring tasks (i.e., Communications and Transactions).

7.4.1.3 Duration

Participants reported that the majority of all monitoring activities were long term tasks, meaning they would continue indefinitely, and this ranged from 75% of Fact Finding monitoring to 100% of Communications monitoring. While short term monitoring tasks were not infrequent, they were most common during Fact Finding, Information Gathering, and Transactions. Short term monitoring tasks had very different motivations behind their cessation. Within Browsing, participants reported the task came to an end simply because they lost interest, while short term Fact Finding monitoring came to an end when the anticipated information became available. Finally, participants reported that Information Gathering monitoring concluded upon the completion of a particular goal.

7.4.1.4 Frequency

Monitoring activities were most commonly performed daily. Fact Finding monitoring was also conducted as the result of a particular event and Communications and Maintenance monitoring often occurred continually throughout the day. Participants reported that their monitoring routine was impacted by external factors, such as the end of business hours and the time change between east and west coast. Two participants reported that they set reminders in their calendars to remind them to engage in the monitoring task on a monthly basis.

7.4.1.5 Follow-up

One aspect where monitoring activities differed was follow-up actions. For instance, Browsing monitoring had little non-electronic follow-up and a third of Browsing monitoring did not have any follow-up at all. This seems to indicate that the information monitored during Browsing is not consumed in the same way as it is for Fact Finding or Information Gathering, which had much higher reported levels of non-electronic follow-up. The information gained during Fact Finding or Information Gathering monitoring activities seemed to have a larger impact on participants' every day actions whereas the information gained during Browsing monitoring was simply used to be aware of "what was going in the world". Maintenance and Communications monitoring tasks consisted of only electronic follow-up, which typically consisted of changes or edits to the web page being monitored (Maintenance) and email (Transactions).

7.4.1.6 Work, School, and Personal Monitoring Activities

With most monitoring task types, there was a predominant type (personal, work/school, both). Browsing and Transactions were predominantly personal while Information Gathering and Maintenance were primarily work related. Participants described most Communications monitoring as both personal and work related in nature. Fact Finding monitoring was the only activity without a clear division as the task nature was split between personal and work.

7.4.2 Monitoring as a Compulsion

While we did not explicitly set out to monitor the compulsive behaviour that sometimes accompanies monitoring, we observed some participants for whom this was a compulsive activity. Five separate participants, all of whom were from the student demographic, reported at least one monitoring activity that they described as a compulsion or an addiction. Even participants for whom it was not a compulsion reported they often conducted some monitoring activity much more frequently than was needed.

We were able to categorize monitoring activities described as a compulsion into two categories: short term and long term. While this is an area that requires further research, the

properties of these two types of monitoring activities provide different implications for new monitoring tools.

Two participants reported compulsively monitoring a web site while awaiting new or updated information. In both cases, these were short-term activities, consisting of monitoring the status of an application (Fact Finding). However, participants reported that once the information arrives, the compulsive activity will cease. Therefore, this is not a monitoring activity that needs to be mitigated but instead needs better support so that users can feel comfortable that they will be immediately notified of any changes on the web page.

Three participants reported compulsively monitoring for entertainment purposes, such as news and sports web sites (Browsing). All participants reported that these were long-term activities. In contrast to short term monitoring, there is no end for these activities, although participants stated that they were trying to cut down on the activity since it was encroaching on their schoolwork or research. While serious Internet addictions have been studied in the Psychology literature (Seo & Zhang, 2000), more research is needed to investigate if and how web browser tools can help users achieve a better balance between productive and non-productive (i.e., entertainment, hobby) web use.

> I would say that is my most frequent task (...) I'll check my e-mail and take a quick glance when I get home at night or in the afternoon and depending on which day it is, what time my classes are. I'll go to doing that and it usually takes importance over my homework which is stupid, I've got to change that.

7.5 Design Implication

The findings of this study have highlighted several implications for the future design of web monitoring tools. General recommendations are presented, followed by task specific recommendations.

7.5.1 General Implications

In our discussion with participants, they reported several barriers to their use of monitoring tools: accuracy, loss of screen real estate, method of notification, tool functionality and setup

effort. Therefore, we have developed general recommendations to help mitigate these factors. These recommendations apply to all types of monitoring activities.

7.5.1.1 Accuracy

Participants articulated two dimensions to accuracy: accuracy of the notification and accuracy of the data. Monitoring tools that use "push" notifications must be consistently accurate in their notifications (i.e., the information must have been updated when the tools says). Monitoring tools that rely on third party data, such as Forecastfox, will not be used if users are not confident about the accuracy of the data. For instance, five participants reported they were hesitant to use a tool such as Forecastfox because they were not confident the forecast, which does not come from the national meteorological service, would be accurate. Therefore, is it important that the third party data used is known as an accurate source. Where possible, tools should allow users to choose their preferred data source.

7.5.1.2 Loss of Screen Real Estate

Six of the study participants expressed concerns with the size and location of the Forecastfox tool, by stating that their web browsers were already cluttered and did not want to accommodate any further icons or toolbars. Therefore, monitoring tools and notifications must remain relatively inconspicuous within the main browser window and appear only as needed.

7.5.1.3 Method of Notification

There are many factors to consider when designing a method of notification for monitoring tools, such as push vs. pull, privacy concerns, and the rate of notifications. Seven participants expressed some discomfort with the pop-up notification displayed by the Page Update Checker tool. Participants felt the pop-ups had the potential to be annoying or distracting. One participant stated:

> I would be embarrassed to have something cued up in the corner of my computer saying that you have a friend post on live journal

Notifications should only display a limited set of information so that the monitored information is not easily viewable by colleagues and friends.

Users can quickly become insensitive to constant notifications if pages are updated regularly. For instance, one participant reported:

> It's a balance between getting too many notifications that you don't even notice anymore, for example, eBay they email you every 2 seconds saying there is this offer or that offer, I've become totally insensitive to them.

The choice of push vs. pull notifications are task dependant. Push notifications are not appropriate for monitoring activities where the web pages update frequently. Instead, widgets that allow users to monitor the rate of recent updates would be more appropriate. This particularly applies to Browsing monitoring tasks, which is addressed in more detail in the following section.

7.5.1.4 Functionality

The type of functionality required in a monitoring tool is specific to the type of monitoring activity being supported. In the next section, we discuss the type of functionality needed to support each monitoring task type. However, all monitoring tools should be sophisticated enough to ignore vacuous updates, such as updated ads.

7.5.1.5 Setup Effort

The effort involved in installing and configuring monitoring tools was often cited as a reason why participants did not currently use any monitoring tools. One participant described why he had never personalized his Google news page, even though it only involves minimal setup with long term benefits:

> I know you can modify the page to show you the ones you like best but I've never. It's too much effort for the amount of time you spend.

Monitoring tools should be quick to install and adding or editing a monitoring task should be seamless. In terms of short term tasks, users need to be able to easily quit monitoring a

site. Participants stated that even if a task was expected to last indefinitely, they still were not willing to expend much time in configuring the tool.

7.5.2 Task Specific Implications

The results of our study suggest that task dictates the type of information being monitored. This means that a variety of monitoring functionality is needed to better support the range of monitoring activities that exist. We have capitalized on our finding to develop the following set of task-specific design recommendations.

7.5.2.1 Browsing

Browsing monitoring was characterized by the monitoring of web pages in order to see "what was new". This is a monitoring activity that is difficult to support. Users can quickly become overwhelmed by the amount of new information appearing on news web sites and blogs, which was one reason why some participants reported they choose not to use RSS. One new approach for RSS is personalized homepages, such as those offered by Google, Yahoo, and Netvibes (2007). These tools offer a more minimalist approach to RSS by simply displaying links to the newest articles. Users also have control over the layout and organization of the content on the homepage.

While tools can employ learning techniques to predict or anticipate a user's interests, Browsing tasks by nature are often serendipitous. A user may simply be keeping up to date with what is going on in the world with no particular topic of interest. Therefore, it may be difficult, if not impossible for tools to accurately predict what is of interest to users engaging in a Browsing monitoring activity. Also, with the exception of pages that update infrequently or unexpectedly, notifications of updated content are not appropriate for Browsing monitoring tasks.

Based on the results of our interviews, we hypothesize that the key to supporting Browsing monitoring is not through notifications or updates of new content but instead through providing an awareness of a web page's update activity. For instance, a small icon or menu item could provide users with a graphical or textual metered view of updates to a particular web page. Since we observed that bookmarks are commonly used during monitoring

activities, this information could also be integrated within the bookmarks list. For instance, recent versions of both Firefox and Safari offer RSS functionality embedded within the bookmarks.

7.5.2.2 Fact Finding

Fact Finding monitoring was characterized by the monitoring of specific pieces of information. These monitoring tasks differ dramatically from Browsing monitoring in that users know exactly what information they are looking for. The results of our interviews indicate that in order to support Fact Finding monitoring, it is not enough for a tool to notify users that the content on the page has changed, even if the change is significant. Therefore, the functionality provided by the Page Update Checker tool demonstrated during the study may not be appropriate for Fact Finding monitoring activities.

We hypothesize that in order to effectively support Fact Finding monitoring tasks, tools should (1) allow users to easily identify the factual information they would like to monitor, (2) notify users when that information has been updated, and (3) facilitate the delivery of this updated information. Based on participant responses during the interviews, there are several options for notifications. In particular, participants suggested emails containing updated information or a discrete notification within the web browser. While it may not be feasible or appropriate for the notification to display the updated information, the notification should communicate which web page has changed and facilitate easy access to the updated information. For instance, a user clicking on the notification could be either transported to the location on the web page where the updated information resides or the updated information could be delivered directly to the user through a small window or pop-up. Tools must be flexible enough to reflect a range of individual differences and task variations by allowing participants to customize features such as the method of notification and the frequency of the updates.

7.5.2.3 Information Gathering

These monitoring activities were characterized by the monitoring of new content in order to support an ongoing research-based task. Participants reported that 38.9% of Information

Gathering activities required search queries to access the updated information. Therefore, tools to support Information Gathering monitoring may benefit from using stored queries and notifying users when new content is available. This is a service that is currently available from some content providers, such as digital libraries, but it is content and site specific. Since Information Gathering tasks tend to occur less frequently, overt notifications would only serve to interrupt an unrelated task. Therefore, email may be the most appropriate form of notification for this monitoring activity.

7.5.2.4 Communications

Communications monitoring activities were characterized by the monitoring of web pages for new and updated communications. Communications monitoring is fairly narrow in scope (i.e., primarily email) and appears to be well supported by current email notification tools.

7.5.2.5 Transactions

Transactions monitoring was characterized by the monitoring of web pages either in anticipation of, or following, an online transaction. Due to the dynamic nature of the data, as well as security concerns, it is a monitoring task that is difficult to support using current tools and technology. Future research is needed to determine whether users have an interest in tools to support this monitoring task.

7.5.2.6 Maintenance

Maintenance monitoring was characterized by the monitoring of web page changes with the intent of maintaining the information on the web page. While this was one of the lesser reported monitoring activities, participants who did engage in Maintenance monitoring articulated several requirements for supporting this activity. Unlike Fact Finding and Browsing, this is a monitoring activity where it may simply be enough to know that the page has changed. Therefore, a discrete notification of a page change (to the exclusion of ads), may prove useful for Maintenance monitoring tasks. As evidence, one of our participants reported after the study that he had begun to use the Page Update Checker tool for a Maintenance monitoring task.

One final consideration for the support of Maintenance monitoring is that users must be able to access a page as typical users do. As stated by one participant:

> I want to experience it the same way as web users experience it, sort of on purpose. Even if there was some widget or something you could do this or that, I want to see it exactly the way as other people see it.

While it may be tempting to provide innocuous markers of new content on the page, it is important that the look and feel of the page is not modified in such a way that the user does not see a true representation of the web page.

7.6 Updated Web Information Classification

In Chapter 5, we presented the Web Information Classification, which was developed based on the results of our field study. Since the results of the semi-structured interviews indicated that monitoring is an activity within all web information tasks, we have refined the classification to include monitoring activities.

Table 16 presents a summary of the monitoring data collected from the field study (Chapter 5) and the semi-structured interviews. Using data from both studies, we estimated the level of monitoring that occurs within each task (e.g., low, medium, and high). While we cannot make any concrete conclusions about monitoring from the field study data, we do have several informal observations with respect to the frequency of repeated tasks and the portion of those repeated tasks that appeared to be related to monitoring. The data from the semi-structured interviews provides some indication of the level of monitoring within each task; however, the percentage of reported monitoring activities may not be representative of the absolute breakdown of monitoring on the Web. For example, only 25% of participants reported email as a monitoring activity during the semi-structured interviews. However, it is reasonable to assume that a large portion of web users do monitor web-based email since 95.2% (20/21) of participants who took part in our field study were users of web-based email. For this reason, we estimated that Communications has a 'high' level of monitoring.

Table 16. A summary of findings from Chapter 5 and Chapter 7 related to the frequency of monitoring activities.

Task	Chapter 5 Findings	Chapter 7 Findings	*Estimated* Level of Monitoring
Fact Finding	55.5% of Fact Finding tasks were repeated and monitoring appeared to be one reason why tasks were repeated.	Fact Finding accounted for 24.3% of all monitoring activities.	Medium
Information Gathering	58.9% of Information Gathering tasks were repeated and repeated tasks appeared to be the result of tasks taking place over several days.	Information Gathering accounted for 10.4% of all monitoring activities.	Low
Browsing	84.4% of Browsing tasks were repeated and repeated tasks appeared to be habitual and of a monitoring nature.	Browsing accounted for 42.8% of all reported monitoring activities	High
Transactions	95.2% of Transactions were repeated and email accounted for 80.4% of all Transactions.	Transactions accounted for 9.2% of all monitoring activities.	Low
Communications		Communications accounted for 2.9% of all monitoring.	High
Maintenance	N/A	Maintenance accounted for 4.0% of all monitoring activities.	Low

Based on Table 16, we developed an updated version of the Web Information Classification (see Figure 22), which includes monitoring as an activity within each information task. A shaded gradient is used to represent the estimated amount (low, medium, high) of monitoring that takes place within each task.

We expect that as the web continues to evolve, so will the Web Information Classification. For example, since this research was conducted, new web applications that facilitate online word processing have been released (e.g., Google Documents and Spreadsheets). If the use of these applications becomes mainstream, the Web Information Classification will need to be modified to represent these new information goals and tasks.

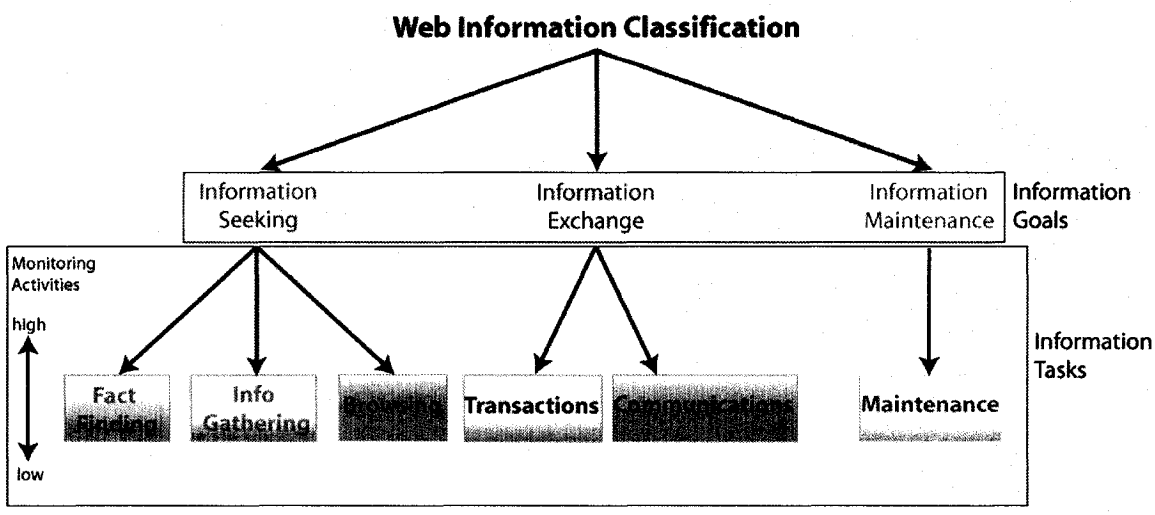

Figure 22. The Web Information Classification is updated to reflect the role of monitoring within information tasks.

7.7 Limitations

One concern when conducting semi-structured interviews is that the interviewer will introduce bias into the interview, thereby influencing the responses of the participants. Bias most commonly occurs through weighted questions, inappropriate prompting, and rephrasing of questions, and deviation from the interview guide (Noyes & Baber, 1999; Oppenheim, 1992). Bias can also occur during the transcription of the participants' responses. In order to reduce the chance of bias, we implemented a number of proactive steps. One of the most effective tools against bias is proper care and preparation. Therefore, the interviews were heavily piloted and an interview guide was used in all interviews. While often repetitious, care was taken not to deviate from the interview guide wording, although we did allow participants to deviate, if they chose. All interviews were transcribed by an individual who was not associated with the study or our research lab.

Another concern when studying user behaviour on the Web through interviews is the inaccuracy or bias associated with self-reported data. While we did not aim to capture a complete picture of participants' monitoring behaviour, we did capture a piece of their behaviour. Conducting interviews *in situ* did help alleviate this concern as participants were able to demonstrate how they actually conducted the task. It is important to note that

participants, particularly those who were interviewed in their workplace, may have been hesitant to discuss personal monitoring activities. For instance, many participants felt the need to qualify some of their monitoring activities with statements such as "on my lunch break" and "download music legally". We hypothesize that participants may have felt the need to qualify their behaviour in case the interview was overheard by coworkers. While future work is needed to study monitoring behaviour from a more objective and quantitative perspective, this study does provide an initial exploration into participants' real world monitoring tasks.

7.8 Summary

We conducted a set of semi-structured interviews in order to determine: (1) if monitoring is an independent information seeking task or an activity within several tasks; (2) the general characteristics of monitoring activities; and (3) the characteristics of each monitoring activity, by task type. Previous research suggested that monitoring may be synonymous with Browsing or an independent information seeking task; however, the results of our study indicate that it is indeed an activity within each web information task (Fact Finding, Information Gathering, Browsing, Transactions, Communication, Maintenance). While the general characteristics of monitoring (use of navigation mechanisms, duration of the activity, and frequency of the activity) did not differ greatly across tasks, there were discernable differences in the types of information being monitored. The implications of our findings suggest that very different functionality is required to support different information monitoring activities. To this end, we have developed recommendations that provide insight into how monitoring activities may be more effectively supported. Using the results from both the field study presented in Chapter 5 and the results of the semi-structured interviews, we have also updated the Web Information Classification (presented in Chapter 5) to include monitoring as an activity within all web information tasks.

Chapter 8

An Evaluation of Task-specific Monitoring Tools

In Chapter 7, we presented the results of semi-structured interviews that were conducted to examine the role of monitoring within the Web Information Classification. We found that monitoring is an activity that occurs within all types of web information tasks and that the type of information monitored differs across tasks. In this chapter, we present three task-specific monitoring tools that were developed using the findings from the semi-structured interviews. We chose to focus on the monitoring activities that had the greatest potential for support and this consisted of Fact Finding, Maintenance, and Browsing monitoring activities. We developed prototype versions of the tools and evaluated their use during a laboratory study. The primary contribution of the research presented in this chapter is an evaluation of the tool functionality that was developed based on the design recommendations presented in Chapter 7. As well, we reflect on the recommendations developed in Chapter 7 and provide further insight into future directions for the design of task-specific tools to support web monitoring activities.

8.1 Research Questions

For each prototype monitoring tool, the study reported in this chapter was designed to address the following research questions:

RQ1. Is the tool appropriate for the task type? We hypothesize that different monitoring activities require different types of support. Therefore, three monitoring tools were developed to support task-specific monitoring activities (Fact Finding, Maintenance, and Browsing). For each tool, we were interested in whether the tool effectively supports the type of information being monitored as well as the user behaviour accompanying the monitoring activity.

RQ2. Can we identify features to improve the usability and utility of the monitoring tools? We were interested in how the design of each tool's user interface and back-end functionality could be further improved.

The results of this study will serve to evaluate the tool functionality that was developed based on the design recommendations developed in Chapter 7. We will also reflect on these recommendations to provide further insight into the design of web-based monitoring tools.

8.2 Prototype Monitoring Tools

In this section we present three prototype monitoring tools that we designed, developed, and evaluated during a laboratory study. The design of these tools was informed by the findings from the study presented in Chapter 7. While the ability to personalize the delivery of the information was an important recommendation, we did not incorporate this functionality into the prototype tools because we wanted a consistent environment for all participants during this initial evaluation. Details on the implementation of the tools are given in Section 8.3.4.

8.2.1 Text Clip: A Monitoring Tool to Support Fact Finding Tasks

Fact Finding monitoring is characterized as the monitoring of specific pieces of information, including text, images, and files. Based on the design recommendations listed in Table 17 , we developed the Text Clip tool, which allows users to monitor specific clips of text within a web page. Table 17 also describes the specific tool functionality that was developed based on the recommendations.

Table 17. For Fact Finding monitoring activities, the findings and implications from Chapter 7 are listed, as well as the newly implemented tool functionality.

Findings from Chapter 7	Implication and Recommendations from Chapter 7	Implemented Functionality
Fact Finding monitoring is characterized by the monitoring of specific or factual pieces of information.	Fact Finding monitoring tools should only notify users when specific information on a web page has changed. In order to monitor factual information, this means that Fact Finding monitoring tools must allow users to: (1) Identify information to monitor (2) Notify users when the information has changed (3) Facilitate the delivery of the updated information	(1) The tool allows users to highlight the information they would like to monitor. (2) A notification icon appears when the browser detects that the previously highlighted information has changed. (3) Users can view the updated information by clicking on the notification icon.
Due to privacy concerns, some participants were uncomfortable with notifications automatically displaying updated information.	Notifications should display a discrete notification of updated content.	A notification icon alerts users to updated information but they must click on the icon to view the updated information.

To monitor a clip of text, the user highlights the desired text on the web page (Figure 23a) and selects *Monitor Clip* from the web browser's *Tools* menu (Figure 23b). When the browser detects that the highlighted information has changed, a notification appears in the bottom right hand corner of the web page (Figure 24a). The notification displays a small icon that represents the web site (for example, a web site's favicon.ico image) that has changed and a different icon is displayed for each distinct web page. Clicking on the notification icon will display a small pop-up containing the updated information (Figure 24b). Users can navigate directly to the page by clicking on the pop-up, which will also close the pop-up. Alternatively, users can click on the notification icon a second time to close the pop-up. In both cases, the notification icon will disappear.

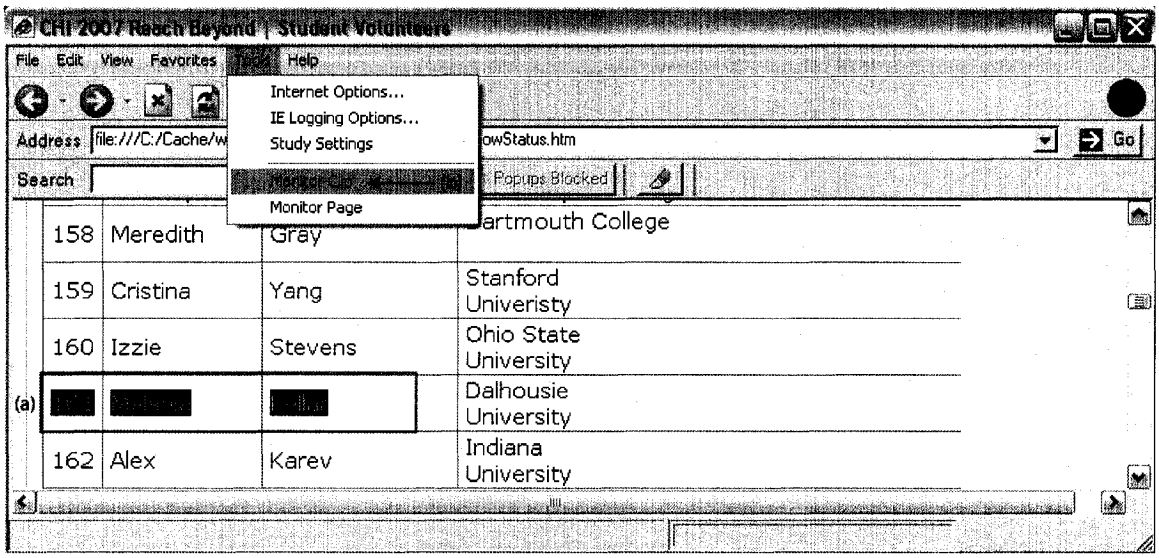

Figure 23. The user selects the text to be monitored (a) and then selects *Monitor Clip* from the *Tools* menu (b) to initiate the Text Clip tool.

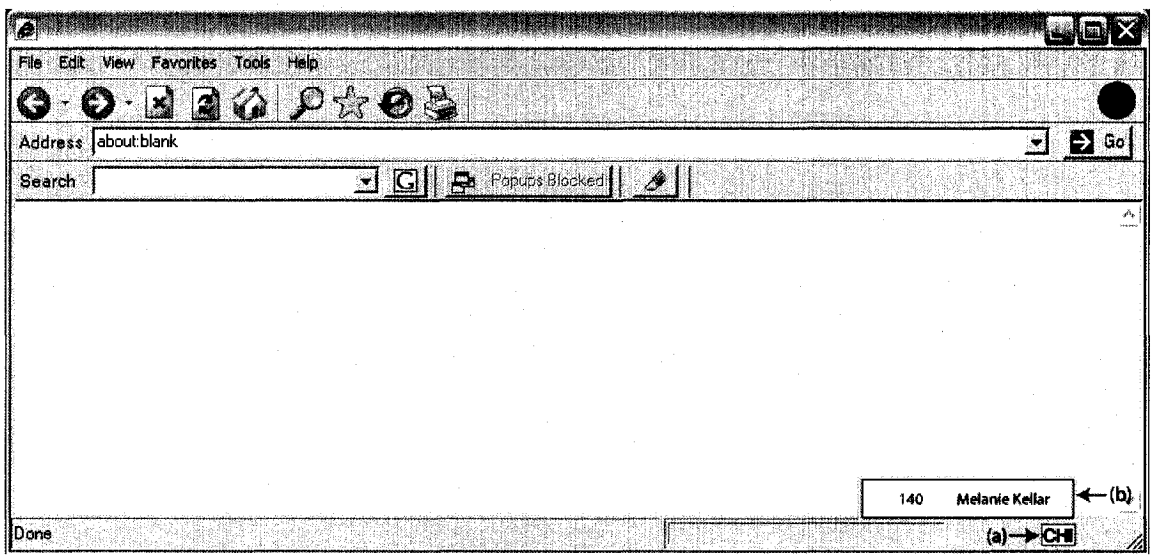

Figure 24. The Text Clip tool alerts users when a pre-identified piece of information on the page has changed. Users click on the notification icon (a) to view a popup, which displays the updated information (b).

8.2.2 Page Updated: A Monitoring Tool to Support Maintenance Tasks

Maintenance monitoring is characterized by the monitoring of web page changes with the intent of maintaining the information on the web page. For example, some participants in our semi-structured interviews reported monitoring their departmental web page for changes (by other users) in order to make sure the new information was correct and typo-free. Based

on the design recommendations listed in Table 18, we developed the Page Updated tool, which allows users to monitor any changes to a particular web page. Table 18 also describes the specific tool functionality that was developed based on the recommendations.

Table 18. For Maintenance monitoring activities, the findings and implications from Chapter 7 are listed, as well as the newly implemented tool functionality.

Findings from Chapter 7	Findings from Chapter 7	Implemented Functionality
Maintenance monitoring is characterized by the monitoring of a web page for the purpose of maintaining it (e.g., spelling mistakes, layout problems).	Users are usually interested in any changes to the web page; therefore, Maintenance monitoring tools should notify users when any information on the page has changed.	A notification icon appears when any information changes on a web page.
Since this is an activity often conducted by users who are responsible for the maintenance of a web page, participants need to see the page as a typical user would.	It is important that the look and feel of the web page is not modified in such a way that the user does not see a true representation of the web page.	The tool does not modify the web page to indicate what has been modified on the web page.
Due to privacy concerns, some participants were uncomfortable with notifications automatically displaying updated information.	Notifications should display a discrete notification of updated content.	A notification icon alerts users to updated information but they must click on the icon to view a message that the page has changed.

To monitor any changes to a web page, the user must navigate to the page and select *Monitor Page* from the web browser's *Tools* menu (Figure 25). When the browser detects that the web page has changed, a notification will discretely appear in the bottom right hand corner of the web page (Figure 26a). The notification displays a small icon (for example, a web site's favicon.ico image) that represents the web site that has changed and a different icon is displayed for each distinct web page. Clicking on the notification will display a small pop-up notification stating that the page has changed (Figure 26b). Users can navigate directly to the page by clicking on the pop-up, which will also close the pop-up. Alternatively, users can click on the notification icon a second time to close the pop-up. In both cases, the notification icon will disappear. The Page Updated tool is a more general version of the Text Clip tool. The functionality provided by the Page Updated tool is similar to that of the Page Update Checker (Mozdev, 2006) introduced in Chapter 7.

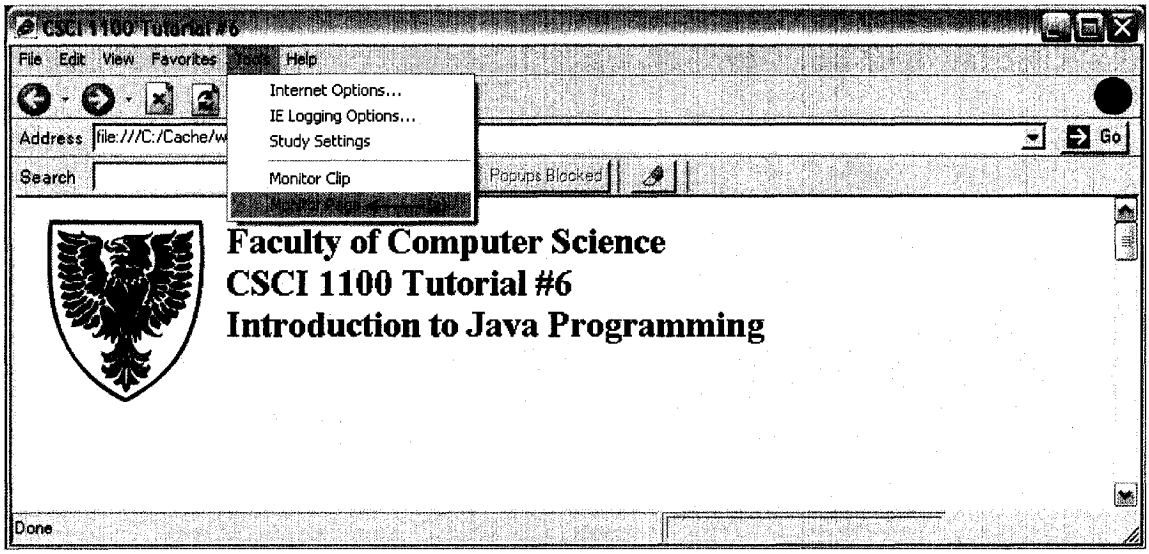

Figure 25. The user selects *Monitor Page* from the *Tools* Menu (a) to initiate the Page Updated tool.

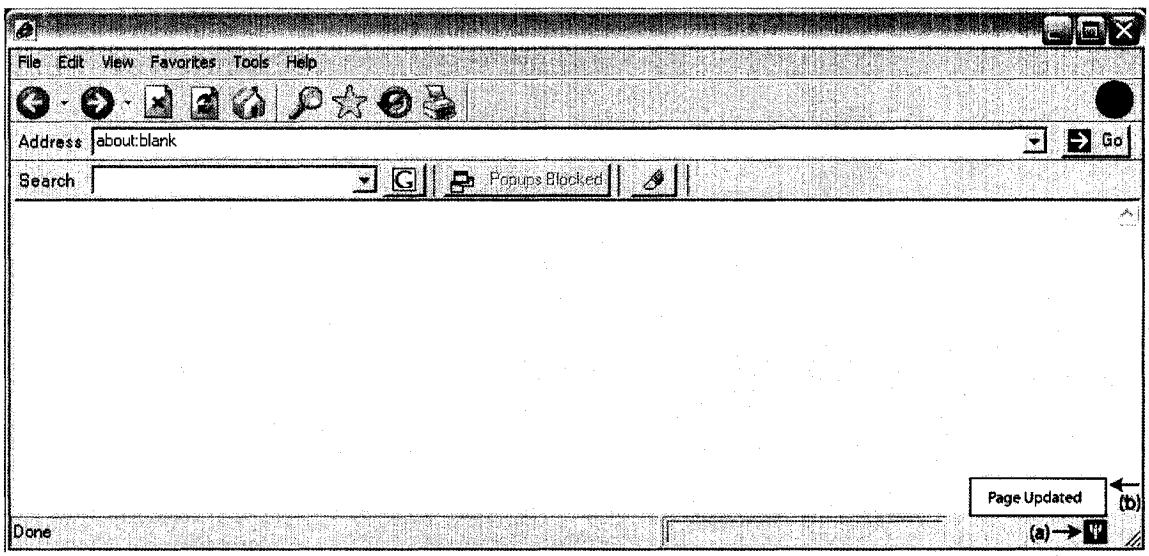

Figure 26. The Page Updated tool alerts users when any changes are made to the web page. Users click on the notification icon (a) to view a popup, which notifies them that changes have been detected on the monitored web page (b).

8.2.3 Enhanced Bookmarks: A Monitoring Tool to Support Browsing Tasks

Browsing monitoring was characterized by the monitoring of web pages in order to see "what was new" or to keep up-to-date. Based on the design recommendations listed in Table

19, we developed the Enhanced Bookmarks tool, which provides users with an awareness of the rate of change of bookmarked web pages. Table 19 also describes the specific tool functionality that was developed based on the recommendations. In contrast to the two previously presented tools, the Browsing monitoring tool is not a notification tool, but instead an awareness tool. Enhanced Bookmarks provide users with a visual representation of how much the web page has changed since their last visit (see Figure 27a). The visual representation consists of bars that estimate the percentage of the page that has changed. For instance, the enhanced bookmark shown in Figure 27b shows a change of 100% since the user's previous visit. This visualization could also be provided for bookmarks viewed through the side browser window or for bookmarks displayed on the toolbar.

Table 19. For Browsing monitoring activities, the findings and implications from Chapter 7 are listed, as well as the newly implemented tool functionality.

Findings from Chapter 7	**Findings from Chapter 7**	**Implemented Functionality**
Browsing monitoring is characterized by the monitoring of information to see what is new or to stay up-to-date. Previous research has also shown that Browsing tasks are often serendipitous (Choo, Detlor, & Turnbull, 2000).	Due to the serendipitous nature of Browsing activities, it may be very difficult for monitoring tools to accurately predict what is of interest to a user. Therefore, instead of identifying content of interest, tools should provide users with an awareness of a web page's update activity.	Users are provided with an abstract representation of the percentage of the web page content that has changed since their previous visit.
Bookmarks are commonly used to navigate during Browsing monitoring activities.	Given the long term nature of monitoring activities and the common use of bookmarks, functionality to support Browsing monitoring could be provided through bookmarks.	The abstract representation (mentioned above) is integrated with the bookmark list.
Most Browsing monitoring activities were long term.		
Some participants found RSS overwhelming because it can be too much information to process.	With the exception of pages that updated infrequently, Browsing monitoring tools should not overload users through continual notifications of updates to web pages.	The tool is an awareness tool and does not issue any notifications to the user.

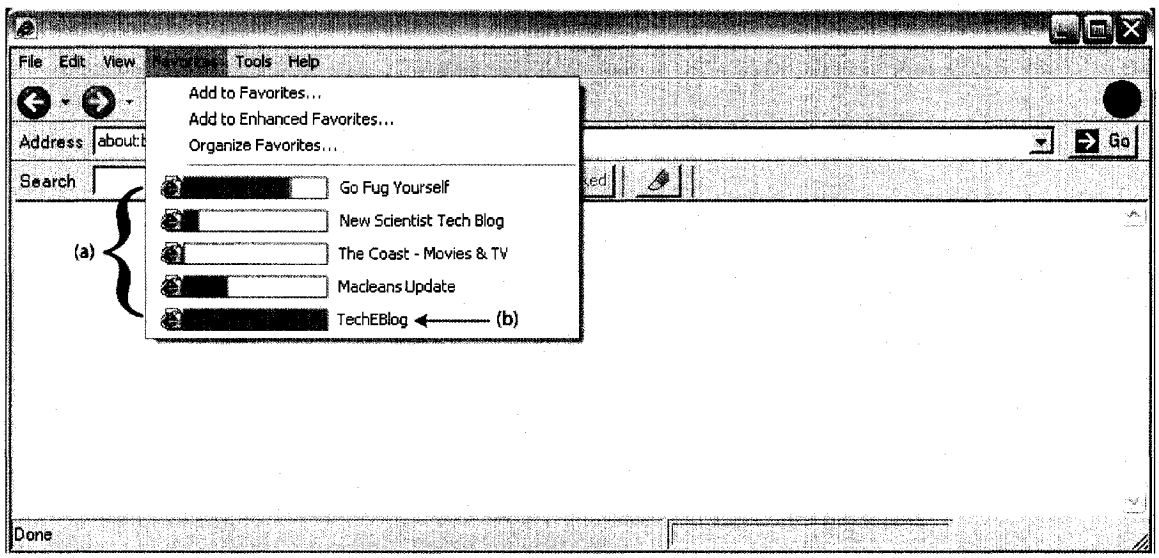

Figure 27. The Browsing monitoring tool consists of Enhanced Bookmarks (a), which provide a visual representation of how much the page has changed since the user's last visit. For example, the enhanced bookmark for the web site "TechEBlog" (b) indicates a 100% change since the user's last visit.

8.3 Methodology

Twenty participants were recruited to take part in an hour long laboratory experiment between December 2006 and January 2007. In this section we first describe the participants who took part in the study and the tasks they were asked to complete. This is followed by an outline of the study procedure, the study instruments and materials, and the data collection techniques. The research described in this chapter was approved by the Social Sciences and Humanities Human Research Ethics Board at Dalhousie University.

8.3.1 Participants

A laboratory study was conducted with 20 participants (6 female) recruited from within the Dalhousie University community. Recruitment notices were circulated by email and stated that anyone who regularly monitors information on the Web was eligible to participate. Participants were also required to have strong English language skills due to the nature of the tasks completed during the study. All participants who took part in the study were remunerated $10 for their time.

The median age group category of the participants was 21-30. Eighteen of the participants were students and the remaining two participants consisted of a staff member and a post doctoral fellow in Engineering. The academic breakdown of the 18 students was divided between Computer Science (11/18), Commerce (3/18), Arts (2/18), Law (1/18), and Mathematics (1/18). Most participants were frequent web users; the mode response to the question "how many hours a week do you spend on the Web" was "40+ hours a week". The most common web browser used by participants was Firefox (14/20), followed by Internet Explorer (4/20), and Safari (2/20). Half (10/20) of our participants reported using some type of tool to help them monitor information on the Web. Just over half (6/10) of these participants were Computer Science students. The remaining 4 were students in Journalism (1), Commerce (2), and Law (1). All but 1 of the 10 participants who said they used monitoring tools used some form of RSS. Examples of RSS tools used include RSS readers, personalized homepages (e.g., Google homepage, NetVibes (2007)), Firefox LiveBookmarks, and the Mac Dashboard.

8.3.2 Tasks

Studying monitoring behaviour in the laboratory is challenging because it is difficult to realistically impose web monitoring tasks on participants. Studying web-based monitoring in a controlled environment also presents technical challenges, such as how to control the rate of changes or updates used in the study. Since we were primarily concerned with exploring participant perceptions towards the utility and usability of the three monitoring tools, we chose to conduct a laboratory study where we manipulated cached web pages to simulate typical web monitoring activities.

In designing the study, we wanted to replicate a multi-task workplace environment where a user may be working on a primary task but also monitoring information on the Web in the background. Therefore, during the study participants were asked to complete a primary task as well as a number of secondary monitoring tasks. The primary task consisted of a written review of a movie, book, or video game. Participants were asked to complete the review using a web-based word processor (Google Documents) since the monitoring tools that we were evaluating were displayed within the web browser. The quality of the review was not important to us. The review was simply a vehicle to occupy participants and create a

situation where it made sense to use the monitoring tools. The primary task was described to participants as follows:

> You are responsible for writing a review of a recent movie you have seen, a book you have read, or a video game you have played. Your review should first provide a brief summary of the plot followed by your review of the movie/book/game. If you need help with your review, check out the "Writing a Movie Review" bookmark or you can look for help on the Web. This task is your top priority during the study. You will be using Google Documents to write your book or movie review. Please feel free to use the Web as you normally would (e.g., open multiple windows, use a search engine, etc.).

Similar to the experimental approach used by McCrickard, Chewar, Somervell, and Ndiwalana (2003), participants were told this was their primary responsibility during the study. Participants were also told they were to take part in secondary monitoring tasks. This included two Fact Finding monitoring tasks and one Maintenance monitoring task. The three monitoring tasks were described to participants as follows:

Task 1: Fact Finding Monitoring - Student Volunteer Wait-list Status (FF1-CHI)

> You have signed up to be a student volunteer for the upcoming 2007 ACM Human-Computer Interaction conference (CHI 2007). The conference is being held in San Jose, California. Being a student volunteer is a coveted position as all of your conference fees are waived and you are provided with free accommodations for the duration of the conference.
>
> There are hundreds of students who apply and only a limited number are selected. Unfortunately, you are currently on the wait list. There are only 120 people on the accepted list. The good news is that the wait list often moves quickly. It is important that you monitor your status on the list as you only have a limited amount of time to confirm your attendance once you have been accepted as a student volunteer

The predefined updates are shown in Table 20.

Table 20. The three predefined updates to the CHI Student Volunteer web page are shown.

Update Number	Updated Position
0	Position 150 on student volunteer list (wait list)
1	Position 140 on the student volunteer list (wait list)
2	Position 135 on the student volunteer list (wait list)
3	Position 119 on the student volunteer list (accepted list)

Task 2: Fact Finding Monitoring Task 2 - eBay Guitar Hero Auction (FF-eBay)

> You are currently selling a copy of the PlayStation2 game Guitar Hero II – (the highly sought after version with the red wireless guitar) so you are monitoring the current price of the auction. You paid $110 for it yourself, so you are really hoping to make your money back, and then some.

The predefined updates are shown in Table 21.

Table 21. The predefined update to the eBay auction web page is shown.

Update Number	Updated Auction Price
0	US $101.99
1	US $115.01

Task 3: Maintenance Monitoring - Web Site Course Content (MA-Dal):

> You are the lead instructor for the first year Intro to Java Programming course (CSCI 1100) which entails supervising three teaching assistants (TAs). Your TAs are responsible for developing the tutorial material and posting it to the Web. However, they are sometimes sloppy with their work so you need to monitor the page for updates and double check new content for typos and errors If new material is posted, please take note of any errors on the page made by your TAs. Last week they had a total of 6 typos. Did they do a better job this week?

The predefined updates are shown in Table 22.

Table 22. The predefined update to the Dalhousie Tutorial web page is shown.

Update Number	Updated Information
0	Week 1-5 tutorials available
1	Week 6 tutorial available

Participants were also asked to participate in a fourth monitoring task, which was a Browsing monitoring task. The Browsing monitoring task consisted of monitoring a set of five web pages (listed in Table 23), consisting of blogs and entertainment related web sites. Participants were first primed on the web sites at the beginning of the study and then asked to return to the web sites halfway through the study, during which time the web sites had undergone some updates. The updated web pages changed anywhere from 0% to 100%; Table 23 lists the percentage changed for each page.

Table 23. The five web sites participants monitored during the Browsing monitoring task.

Web Site	Description	Percentage Changed
Go Fug Yourself www.gofugyourself.com	A humorous blog critiquing celebrity fashions	75%
MacLeans.ca Weekly Update with Scott Feschuk www.macleans.ca/feschuk	A blog providing a daily commentary on current news events (e.g., political, celebrity)	30%
TechEBlog www.techeblog.com	A blog covering the latest news on technology and gadgets	100%
New Scientist Technology Blog www.newscientist.com/blog/technology	A blog covering the latest news in technology	10%
The Coast www.thecoast.ca	The Movies & TV section of a local online community newspaper	0%

8.3.3 Procedure

Upon arrival participants were asked to sign a consent form and complete a demographic questionnaire. Figure 28 provides a visual representation of the study timeline. The study was divided into 5 parts, which are described below.

Part I: Browsing Primary Session

Part I was 5 minutes in duration. Participants were told this was the "warm-up" portion of the study and asked to familiarize themselves with the five cached web sites using traditional bookmarks, shown in Table 23. It was stressed to participants that it was not necessary to memorize the content on the web sites (i.e., they were not going to be asked any questions about it) but that they may be asked to return to the web sites later in the study. The enhanced bookmark functionality was not introduced at this point.

Part II: Training Session

Participants took part in a training session where the three monitoring tasks (FF-CHI, FF-eBay, and MA-Dal) were introduced. For each monitoring task, participants were shown the web page and the information they were expected to monitor. The Fact Finding and Maintenance monitoring tools were then introduced; their use was demonstrated and participants were also shown how to configure the monitoring tools. Participants were told that they might be asked to answer questions regarding the monitored information at the end of the study.

Part III: Primary Task and Secondary Monitoring Tasks

Part III was 10 minutes in duration. Participants began their primary task (i.e., the book review) using Google Documents and were told they may receive notification of updated content. The first notification (FF-CHI-1) appeared two minutes into Part III and the second notification (FF-eBay-1) appeared 8 minutes into Part III. After 10 minutes had elapsed, participants were interrupted and told they were being given a browsing break.

Part IV: Browsing Session

Part IV was 5 minutes in duration. Participants were introduced to the enhanced bookmark functionality and asked to browse the five cached web sites using the Enhanced Bookmarks. Participants were told they only had to visit those web sites that were of interest to them.

Part V: Primary Task and Secondary Monitoring Tasks

Part V was 10 minutes in duration. Participants were told to return to their primary task and that they may receive further notification of updated information. During Part V, a notification occurred at the 4 minutes mark (MA-DAL-1), shortly followed by the next notification at 4.5 minutes (FF-CHI-2), and the final notification appeared after 7 minutes (FF-CHI-3).

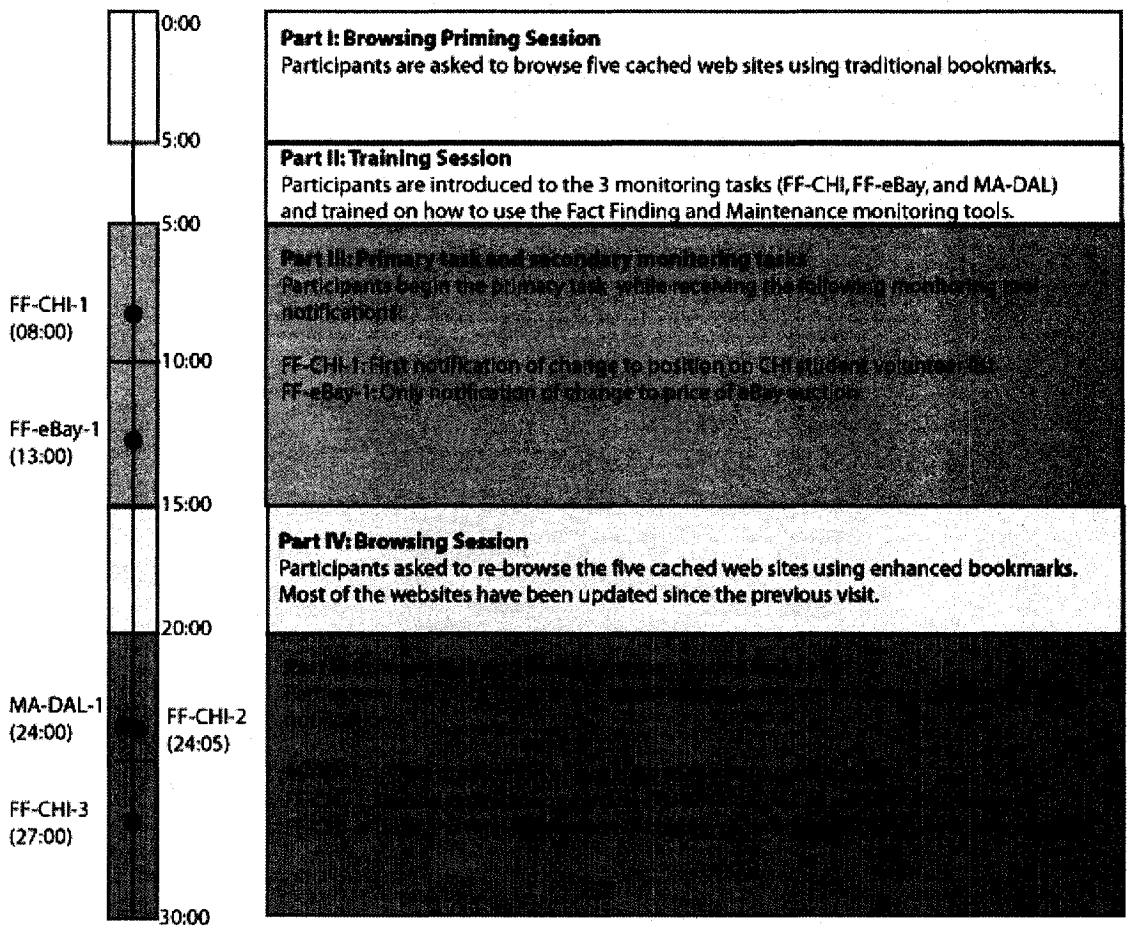

Figure 28. A visual representation of the study timeline.

8.3.4 Study Instruments and Materials

In this section we describe the instruments and materials used during the study, which include a custom-built web browser and cached web pages.

8.3.4.1 Custom Web Browser and Monitoring Tools

In Section 4.3, we presented a custom web browser that was built to mimic IE and was used during our fist study. The same web browser was used during this study to implement and evaluate the prototype monitoring tools. All the front-end functionality described in Section 8.2 was implemented within the web browser; however, a series of predefined functions provided the illusion of a working backend system. Figure 29 shows a pseudo code depiction of how the monitoring tool functionality was implemented.

```
If Part I:    //Warm-Up Browsing Session
              Display traditional bookmarks for five (cached) browsing
                    sites

If Part II:   //Training Session
              Display training notifications

If Part III://Primary & Secondary Task Session
              Set Timer_1 = FF-CHI-1 Notification = 2 minutes
              Set Timer_2 = FF-eBay-1 Notification = 8 minutes

              If Timer_1:
                    Overwrite cached CHI page with "updated" page
                    Display FF-CHI-1 Notification
              If Timer_2:
                    Overwrite cached eBay page with "updated" page
                    Display FF-eBay-1 Notification

If Part IV:   //Browsing Session
              Overwrite cached pages with "updated" pages
              Display Enhanced Bookmarks for five browsing sites

If Part V:    //Primary & Secondary Task Session
              Set Timer_1 = MA-Dal-1 Notification = 4 minutes

              Set Timer_2 = FF-CHI-2 Notification = 4 minutes, 5 seconds
              Set Timer_3 = FF-CHI-3 Notification = 7 minutes

              If Timer_1:
                    Overwrite cached Dal page with "updated" page
                    Display MA-Dal-1 Notification
              If Timer_2:
                    Overwrite cached CHI page with "updated" page
                    Display FF-CHI-2 Notification
              If Timer_3:
                    Overwrite cached CHI page with "updated" page
                    Display FF-CHI-3 Notification
```

Figure 29. Pseudo code representing how the monitoring tool functionality was implemented.

For the Fact Finding and Maintenance monitoring tools, the time of all notification events was predefined within the web browser in advance of the study. For example, the first Fact Finding notification was scheduled to appear two minutes after the commencement of Part III. All notifications occurred at the same time for all participants. At the beginning of each study part, the researcher conducting the study indicated the study part using the dialog box shown in Figure 30.This initiated a timer within the web browser. When the timer hit the previously scheduled time, the appropriate notification was displayed. At this time, the old version of the web page was overwritten with an "updated" cached version of the web page. In Section 8.3.4.2, we describe the use of cached web pages in more detail.

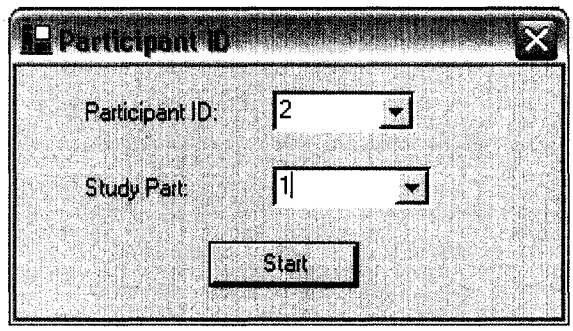

Figure 30. The dialog box used to indicate the participant ID and study part.

8.3.4.2 Cached Pages

All pages monitored by participants in the study were cached. Copies of the original pages used in the Fact Finding and Maintenance monitoring tasks were manually altered (in advance of the study) in order to simulate updated pages. For example, two versions of the eBay page were required. The first had the original auction price while the second displayed an updated auction price. For the pages used in the Browsing monitoring tasks, different versions were downloaded on two separate dates. The newer set of pages acted as the "updated" version. All cached web pages were downloaded using SuperBot (Sparkleware, 2006), which downloaded all images, ads, videos displayed on the target page. The tool also downloaded all pages that were within one hyperlink (i.e., depth = 1) from the target page.

8.3.5 Data Collection

Over the course of the study, we used three main methods of data collection: web browser logs, participant questionnaires, and a semi-structured interview.

8.3.5.1 Logged Web Browser Interactions

Over the course of the study, the custom web browser logged all web pages visited and the use of all web browser navigation mechanisms (i.e., how participants navigated to each web page). We also logged all interactions with the monitoring tools, including the time in which a notification appeared and any user interactions with the notification (e.g., viewing or closing the notification pop-up). An example log file is shown in Table 24 and displays the Window ID, Date & Time, Study Part, URL, Event (Browser Event or Navigation Event), and Description (of the event)

Table 24. Event log file with Browser and Navigation Events.

Window ID	Date Time	Study Part	URL	Event	Description
7275894	12/13/2006 19:15:10.51	3	www.cs.dal.ca	**Browser Event**	NotificationAppear-CHI1
7275894	12/13/2006 19:15:16.67	3	www.cs.dal.ca	**Browser Event**	ViewedPopUp-CHI1
7275894	12/13/2006 19:15:21.23	3	www.cs.dal.ca	**Browser Event**	PopUpClosed-CHI1
7275894	12/13/2006 19:17:05.18	3	www.cs.dal.ca/~melanie	**Navigation Event**	AUTO COMPLETE

8.3.5.2 Researcher Field Notes

Field notes were collected by the researcher conducting the study, who was seated next to the participant during the study. The researcher noted interesting observations, the order of the sites visited during the browsing sessions, and any questions posed by participants. The coding sheets used are shown in Appendix H.

8.3.5.3 Questionnaires

A demographic questionnaire (see Appendix I) was administered to participants at the beginning of the experimental session. The purpose of this questionnaire was to collect general background information about the participants and to learn about any monitoring tools they were using. Upon completion of the study, participants were also asked to complete a post-study questionnaire for each monitoring tool (Appendix J) that examined the tradeoffs between interruption, reaction, and comprehension experienced by participants while using the monitoring tools as well as user satisfaction data. These questionnaires were loosely based on the questionnaires used by McCrickard et al. (2003) and van Dantzich et al. (2002).

8.3.5.4 Interviews

A brief post-study interview (Appendix K) was conducted upon completion of the study. The interviews collected information on participants' comprehension of the monitored information, as well as preferences and general feedback on the monitoring tools. The researcher conducting the study took notes during the interview and the interviews were also audio recorded if the participant consented.

8.4 Results

In this section, we present the results of the laboratory experiment. We use quantitative and qualitative data collected from the post-study questionnaire, post-study interview, and web browser logs to examine the following factors: participant use of the monitoring tools, comprehension of monitored information, appropriateness of the monitoring tools for the task type, and participant preferences and feedback. The results are presented mainly using descriptive statistics but we also present the results of statistical analyses, where appropriate.

8.4.1 Use of Monitoring Tools

8.4.1.1 Notification Tools

For the notification-based Text Clip and Page Updated tools, we were interested in participants' behaviour after the appearance of the notification icons (shown in Figure 31a). In particular, we were interested in whether participants found the information delivered by the updated information pop-up (shown in Figure 31b) to be sufficient or if participants also viewed the web page containing the updated information. In advance of the study, we expected that many participants would chose to simply view the updated information pop-up when using the Text Clip tool (i.e., Fact Finding monitoring). Alternatively, we expected that most participants would view the web page (via the updated information pop-up) when using the Page Updated tool (i.e., Maintenance monitoring task). Using the web browser log data collected during the study, we examined how participants interacted with the tools upon the appearance of the notification icon. In particular, we were interested in whether participants: (1) chose not to interact with the notification, (2) viewed the updated information pop-up and then closed it; or (3) viewed the updated information pop-up and then viewed the corresponding web page by clicking on the pop-up.

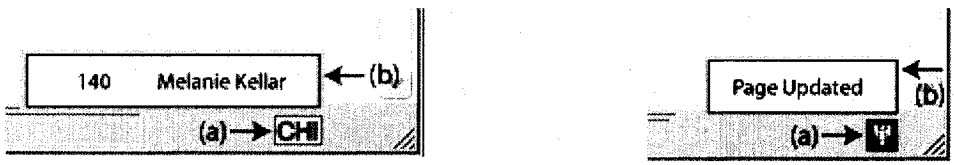

Figure 31. The Text Clip and Page Updated tools consist of the notification icon (a) and the updated information pop-up (b).

We also examined the time between the appearance of the notification icon and when a participant viewed the updated information pop-up. A summary of the logged user interactions with the two notification-based monitoring tools is shown in Table 25.

Table 25. Average time for participants to view the monitoring tool notifications and their follow-up actions.

Order	Notification Event	Time to view Pop- up Notification (sec)		Participant Follow-up Actions			
		M	*SD*	Did not View Pop-up Notification	Only Viewed Pop-up Notification	Viewed Pop-up Notification and Web page	Total
1	FF-CHI-1	71.9	99.0	2 (10.0%)	7 (35.0%)	11 (55.0%)	20
2	FF-eBay-1	72.3	159.9	1 (5.0%)	7 (35.0%)	12 (60.0%)	20
3	MA-Dal-1	40.2	60.0	1 (5.0%)	2 (10.0%)	17 (85.0%)	20
4	FF-CHI-2	64.9	62.1	3 (5.0%)	7 (35.0%)	10 (50.0%)	20
5	FF-CHI-3	20.7	26.8	3 (5.0%)	8 (40.0%)	9 (45.0%)	20

The first notification of updated information was delivered by the Text Clip tool in Part III and alerted participants to the first change in their information on the CHI student volunteer web page. After the appearance of the notification icon, 90% (18/20) of participants viewed the updated information pop-up. The remaining two participants (10%) either did not notice the notification icon or chose to ignore the icon. For the 18 participants who did view the updated information pop-up, the time between the appearance of the notification icon and viewing the updated information pop-up ranged from 3 to 364 seconds (M = 71.9; SD = 99.0). Thirty-five percent (7/20) of participants chose to *only* view the updated information pop-up, while 55% (11/20) of participants navigated to the updated CHI student volunteer web page by clicking on the pop-up.

The second notification of updated information was delivered by the Text Clip tool in Part III and alerted participants to the first (and only) change in the price of their eBay auction. After the appearance of the notification icon, 95% (19/20) of participants viewed the updated information pop-up. One participant (5%) either did not notice the notification icon or chose to ignore the icon. For the 19 participants who did view the updated information pop-up, the time between the appearance of the notification icon and viewing the updated information pop-up ranged from 3 to 573 seconds (M = 72.3; SD = 159.9). Thirty-five percent (7/20) of participants chose to *only* view the updated information pop-up while 60%

(12/20) of participants navigated to the updated eBay web page by clicking on the pop-up. The high variance (SD = 159.9) in the viewing time is attributed to two participants (10%) who did not view updated information pop-up until Part V (469 and 573 seconds later, respectively).

The third notification of updated information was delivered by the Page Updated tool in Part V and alerted participants to the first (and only) change on the Dalhousie tutorial web page. After the appearance of the notification icon, 95% (19/20) of participants viewed the updated information pop-up. One participant (5%) either did not notice the notification icon or chose to ignore the icon. For the 19 participants who did view the updated information pop-up, the time between the appearance of the notification icon and viewing the pop-up ranged from 3 to 206 seconds (M = 40.2; SD =60.0). Ten percent (2/20) of participants chose to *only* view the updated information pop-up. Eighty-five percent (17/20) of participants navigated to the updated tutorials web page by clicking on the updated information pop-up.

The fourth notification of updated information was delivered by the Text Clip tool in Part V and alerted participants to the second change in their information on the CHI student volunteer web page. After the appearance of the notification icon, 85% (17/20) of participants viewed the updated information pop-up. Three participants (15%) either did not notice the notification icon or chose to ignore the icon. For the 17 participants who did view the updated information pop-up, the time between the appearance of the notification and viewing the pop-up ranged from 1 to 159 seconds (M = 64.9; SD = 62.1). Thirty-five percent (7/20) of participants chose to *only* view the updated information pop-up and 50% (10/20) of participants navigated to the updated CHI web site by clicking on the updated information pop-up. This notification icon appeared five seconds after the appearance of the Dalhousie tutorial web page notification icon (presented in the previous paragraph). This was purposefully orchestrated to observe how participants would deal with two notification icons appearing in close temporal proximity to each other. Of the 17 participants who viewed both pop-ups of updated information, 59% (10/17) viewed the tutorials pop-up first and 41% (7/17) viewed the CHI pop-up first. Participants did not voice any concerns with having to deal with two notification icons appearing at the same time.

The fifth notification of updated information was delivered by the Text Clip tool in Part V and alerted participants to the third and final change in their information on the CHI student volunteer web page. After the appearance of the notification icon, 85% (17/20) of participants viewed the updated information pop-up. Three participants (15%) either did not notice the notification icon or chose to ignore the icon. For the 17 participants who did view the updated information pop-up, the time between the appearance of the notification icon and viewing the pop-up ranged from 2 to 109 seconds (M = 20.7; SD = 26.8). Forty percent (8/20) of participants chose to *only* view the updated information pop-up, while 55% (9/20) of participants navigated to the updated CHI web page by clicking on the pop-up.

A chi square analysis of participants' use of the Text Clip tool did not reveal any significant differences between the number of participants who chose to view *only* the updated information pop-up and the number of participants who chose to view both the pop-up and the updated web page. We were also interested in whether participants had a preferred method of interacting with the Text Clip tool. It would appear that 50% (10/20) participants did have a preferred method; 20% (4/20) of participants viewed the updated information pop-up exclusively (if they viewed a notification) and 30% (6/20) of participants consistently viewed both the pop-up and the updated web page (see Table 27). While we do not have data on why participants chose one method over the others, personal preferences likely played a role in participants' interactions with the two notification tools. A chi square analysis revealed a significant difference in participant's use of the Page Updated tool. Significantly [χ^2 (1, N =19) = 11.842, p = 0.001)][5] more participants viewed both the updated information pop-up and the updated web page than just the pop-up. However, this finding is not surprising since the Page Updated pop-up simply indicated that a change had occurred on the page.

[5] We omitted the participant who did not interact with the notification icon.

Table 26. Individual participant interactions with the Text Clip tool.

PID	FF-CHI-1	FF-eBay-1	FF-CHI-2	FF-CHI-3
2	Viewed Pop-up & Page	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up
3	Viewed Pop-up	Viewed Pop-up & Page	Viewed Pop-up & Page	Viewed Pop-up & Page
4	Viewed Pop-up & Page	Viewed Pop-up	Did not view	Viewed Pop-up & Page
5	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up
6	Viewed Pop-up	Viewed Pop-up & Page	Did not view	Viewed Pop-up & Page
8	Did not view	Did not view	Did not view	Did not view
9	Viewed Pop-up	Viewed Pop-up & Page	Viewed Pop-up & Page	Did not view
10	Viewed Pop-up	Viewed Pop-up & Page	Viewed Pop-up	Viewed Pop-up
11	Viewed Pop-up & Page	Viewed Pop-up	Viewed Pop-up & Page	Viewed Pop-up & Page
12	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up
15	Viewed Pop-up & Page	Viewed Pop-up & Page	Viewed Pop-up & Page	Viewed Pop-up
16	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up
18	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up	Viewed Pop-up
20	Viewed Pop-up & Page	Viewed Pop-up & Page	Viewed Pop-up	Viewed Pop-up

The goal of this analysis was not to compare how quickly participants reacted to the notification icons. Instead, we hypothesized that the reaction time may provide insight into participants' engagement with the primary task (i.e., if participants were interacting with the notifications within 2 seconds of their appearance, this might indicate they were not engaged in the primary task). The average time for participants to react to the appearance of a notification icon ranged from 20.7 seconds (FF-CHI-3) to 72.3 seconds (FF-eBay-1). While at first glance this appears to be a wide range, a one-way repeated-measures ANOVA did not reveal any significant differences between the reaction times. The average times for participants to react to the notification tools seem reasonable and suggest that most participants were engaged in their primary task and interacted with the notification tools in a casual manner.

8.4.1.2 Awareness Tools

We were interested in whether or not participants' use of traditional bookmarks differed from their use of the Enhanced Bookmarks. We expected that the Enhanced Bookmarks

would allow participants to engage more efficiently in their Browsing monitoring activities, meaning participants would be less likely to visit web pages with few updates or no updates. We recorded the order in which participants accessed the traditional bookmarks during Part I and also the order in which they accessed the same sites using the Enhanced Bookmarks in Part IV. We also asked participants in the post-study interview how they chose the order of the sites they visited. While the order of the bookmarks varied between participants, the order did not change within participants for Part I and Part IV.

In Part I, participants used traditional bookmarks. According to the logged data collected, all 20 participants viewed all five web sites in sequential order of the bookmarks (i.e., viewed the bookmarks from the top of the menu to the bottom). However, only 70% (14/20) of participants reported they viewed the five web sites in sequential order. The remaining participants reported that they viewed the sites that interested them the most (20% - 4/20) or visited those they had visited before first (10% - 2/20). It is not clear if participants simply did not remember how they chose the order of the web sites they visited or if by coincidence the bookmarks did appear in order of their interests.

In Part IV, participants used the Enhanced Bookmarks. According to the logged data collected, 15% (3/20) of participants viewed all 5 web sites in the same order, 10% (2/20) of participants viewed all 5 web sites but in a different order, and 75% (15/20) of participants viewed less than five web sites in differing orders (mean = 3.2; median = 4). Not surprisingly, the Enhanced Bookmark showing a 100% change in content (TechEBlog) was selected first by 55% (11/20) of all participants. Similarly, the Enhanced Bookmark showing the second greatest amount of change (Go Fug Yourself – 75%) was selected second by 55% (11/20) of participants. Of the 11 participants who selected the TechEBlog Enhanced Bookmark first, 73% (8/11) selected the Go Fug Yourself Enhanced Bookmark second. When asked how they chose the order of the sites they visited, 15% (3/20) of participants reported (accurately) that they visited the bookmarked web sites from top to bottom. Half (10/20) of the participants reported they visited those that had changed most and 35% (7/20) reported they chose the sites to visit based on both their interest in the topic and the amount of change displayed by the bookmark tool.

We hypothesize that the Enhanced Bookmarks have the potential to allow users to browse the Web more efficiently since many participants ignored web sites with little or no changes. For instance, 65% (13/20) participants ignored the web site (the Coast) with no updated content. While these preliminary findings are promising, more research is needed to better understand the impact of Enhanced Bookmarks on the efficiency of web users.

8.4.2 Comprehension of Monitored Information

We were interested in how well the monitoring tools supported participants' understanding of the monitored information. Participants were asked a series of questions (shown in Table 27) immediately following the conclusion of the study, designed to measure their comprehension of the information monitored.

Table 27. Participant responses to questions about the information monitored during the study.

	Correct Reponses	Incorrect Reponses	"I don't know"
Fact Finding – Position on the student volunteer list			
What was your final position on the Student Volunteer List?	14 (70.0%)	6 (30.0%)	0 (0.0%)
Did you make it on to the accepted list?	15 (75.0%)	4 (20.0%)	1 (5.0%)
Fact Finding – Amount of eBay auction			
What was the final price of your eBay auction?	17 (85.0%)	2 (10.0%)	1 (5.0%)
Did you make your money back?	16 (80.0%)	1 (5.0%)	3 (15.0%)
Maintenance – Typos on newly updated tutorial page			
Was the Tutorial page updated?	20 (100.0%)	0 (0.0%)	0 (0.0%)
Do your TAs need to make any edits?	17 (85.5%)	0 (0.0%)	3 (15.0%)

The Fact Finding monitoring task where participants were asked to monitor their status on the CHI student volunteer list was the most complex monitoring activity since the monitoring tool did not communicate any information about whether the participant had been accepted onto the list. Participants had to either view the page to see if they had moved onto the accepted list or remember the cut-off themselves. We were initially concerned that

participants might be aware of their position on the list but might not be aware of whether or not they had been accepted onto the list. However, this was not the case. Seventy percent (14/20) of participants knew their final position and 75% (51/20) of participants knew they had been accepted onto the student volunteer list.

Participants were slightly more accurate in their responses to the second Fact Finding monitoring task where they were asked to monitor the price of an eBay auction. We found that 85% (17/20) of participants correctly reported the current price of the auction and 80% (16/20) of participants knew that the current price of the auction meant that they had recovered the money initially spent on the item. All participants knew that the tutorials page used in the Maintenance monitoring task had been updated and 85% (17/20) of participants knew that edits to the page were required.

Across all monitoring tasks, the overall correct response rate from participants was 82.5%. This indicates that most participants who used the notification tools did not have any trouble comprehending and retaining the information.

8.4.3 Appropriateness for Task

One of our primary motivations for this study was to evaluate the appropriateness of each monitoring tool for the given tasks. That is, how well does each tool support the characteristics of the underlying task (i.e., Fact Finding, Maintenance, and Browsing). Data used in this analysis was collected through the post-study questionnaire and interviews. Questionnaire data reported in this section is based on a Likert scale, ranging from 1 to 7, where 1 is "strongly disagree" and 7 is "strongly agree". In reporting the results we have collapsed the Likert scale responses into three response categories: agree, neutral, and disagree. Table 28 shows a mapping of the categories to the questionnaire (i.e., Likert scale) responses.

Table 28. Likert scale responses from the post-study questionnaires were collapsed into three general categories: disagree, neutral, and agree.

Likert Scale Response	Category
1 - 2	Disagree
4 – 6	Neutral
7 - 8	Agree

During the post-study interviews, we asked participants to state for each monitoring task they completed during the study, which monitoring tool would they have chosen to use. Participants were also told they could suggest a new tool or choose no tool at all. For both Fact Finding monitoring tasks, there was strong evidence that the Text Clip monitoring tool was appropriate for the task. For example, when monitoring their position on the CHI student volunteer list, 85% (17/20) of participants reported that if given the choice, they would choose to use the existing Text Clip tool. Similarly, 90% (18/20) of participants reported that if given the choice during the eBay monitoring task, they would choose to use the existing Text Clip tool.

Participant responses on the post-study questionnaire also supported the appropriateness of the Text Clip tool for Fact Finding monitoring (see Table 29). Results from the questionnaire indicated that 95% (19/20) of participants agreed that the tool functionality was appropriate for the kind of information they were monitoring. Eighty percent (16/20) of participants indicated they would be willing to use the Text Clip tool in their own usage. Participants were also asked during the post-study interview how they could envision using the monitoring tools in their own everyday web usage. In the case of the Text Clip tool, we were interested in whether participants would suggest Fact Finding monitoring tasks. Participants provided a number of examples that were classified as Fact Finding monitoring activities, although some activities could also be classified as Transactions monitoring (i.e., monitoring factual information in anticipation of, or as the result of, an online transaction). Examples provided by participants include using the Text Clip tool to monitor airfare (5% - 1/20), the price and current stock of an item at an online store (15% - 3/20), their bank balance (5% - 1/20), stock prices (5% - 1/20), current class enrolment numbers (10% - 2/20), grades (5% - 1/20), and course assignments (5% - 1/20). One participant also asked about the possibility of using the Text Clip tool for monitoring changes to an image.

Table 29. Collapsed Likert scale data from the post-study questionnaire investigating participants' perceptions of the appropriateness of the Text Clip tool for Fact Finding monitoring.

	Disagree	Neutral	Agree
The tool functionality was appropriate for the kind of information I was monitoring	0 (0.0%)	1 (5.0%)	19 (95.0%)
I would be willing to use a similar tool within my own web usage	2 (10.0%)	2 (10.0%)	16 (80.0%)

When asked what tool they would have chosen to use while monitoring the appearance of new information on the tutorial web page (Maintenance monitoring), 70% (14/20) of participants reported they would have chosen the Page Updated tool. Two participants (10%) indicated they did not have a strong preference between the Page Updated tool and the Enhanced Bookmarks.

On the post-study questionnaire (see Table 30), 70% (14/20) of participants agreed that the Page Updated tool functionality was appropriate for the kind of information they were monitoring. Seventy percent (14/20) of participants indicated they would be willing to use the Page Updated tool in their own usage. During the post-study interviews, participants reported they could see themselves using the Page Updated tool for tasks such as online forums (10% - 2/10), reading online reviews (5% - 1/20), and the local power company's outage page (5% - 1/20). The monitoring tasks mentioned by participants appear to be more similar to Browsing monitoring than Maintenance monitoring. Two participants (10%) reported they would not use the Page Updated tool because it was not appropriate for their current web tasks (i.e., they had no need for this functionality).

Table 30. Collapsed Likert scale data from the post-study questionnaire investigating participants' perceptions of the appropriateness of the Page Updated tool for Maintenance monitoring.

	Disagree	Neutral	Agree
The tool functionality was appropriate for the kind of information I was monitoring	2 (10.0%)	4 (20.0%)	14 (70.0%)
I would be willing to use a similar tool within my own web usage	3 (15.0%)	3 (15.0%)	14 (70.0%)

During the post-study interviews, 50% (10/20) of participants reported that if given the choice, they would have chosen to use the Enhanced Bookmarks tool during the Browsing monitoring task. Three participants (15%) reported they would have preferred to use RSS and another three participants (15%) reported they would have chosen to use the Page Updated tool.

On the post-study questionnaire (see Table 31), 70% (14/20) of participants agreed that the Enhanced Bookmarks functionality was appropriate for the kind of information they were monitoring. Sixty-five percent (13/20) of participants indicated they would use the Enhanced Bookmarks in their own usage. Participants reported they could envision using the Enhanced Bookmark tool for sites that do not have RSS (3) and sporting news (1), which are sites often visited during Browsing monitoring activities. One participant reported he could envision using either the Page Updated tool or the Enhanced Bookmarks to monitor an online newspaper that only updates periodically. All of these monitoring tasks appear to be Browsing monitoring activities.

Table 31. Collapsed Likert scale data from the post-study questionnaire investigating participants' perceptions of the appropriateness of the Enhanced Bookmarks tool for Browsing monitoring.

	Disagree	Neutral	Agree
The tool functionality was appropriate for the kind of information I was monitoring	2 (10.0%)	4 (20.0%)	14 (70.0%)
I would be willing to use a similar tool within my own web usage	0 (0.0%)	7 (35.0%)	13 (65.0%)

8.4.4 Participant Preference and Feedback

Data on participant preferences and feedback on the monitoring tools was collected through the post-study questionnaires and interview. During the post-study interview, participants were asked to rank the three monitoring tools in order of preference. An inverted weighted ranking scheme was used to examine the results, where a tool was assigned a weight of one if ranked first by a participant, a weight of two if ranked second, and a weight of three if ranked third. Using this weighted ranking scheme, the Text Clip monitoring tool was ranked first, followed by the Enhanced Bookmarks, and the Page Updated tool (see Table 32). One

participant declined to rank the tools because he stated that he could not compare the tools since "they are for different things". As in Section 8.4.3, we have collapsed the post-study questionnaire Likert scale responses into three response categories: agree, neutral, and disagree. Participants also had a number of suggestions for improving the interface and functionality of the three monitoring tools

Table 32. Participants' weighted ranks of three monitoring tools.

Weights	Text Clip	Weighted Text Clip	Page Changed	Weighted Page Changed	Enhanced Bookmarks	Weighted Enhanced Bookmarks
1	12	12	1	1	6	6
2	4	8	12	24	5	10
3	3	9	6	18	8	24
Sum		**29**		**43**		**40**

8.4.4.1 Text Clip Tool

As shown in Table 32, the Text Clip tool was ranked first. The results of the post-study questionnaire, shown in Table 33, confirm these results: 85% (17/20) of participants reported they liked the tool functionality, 85% (17/20) reported they liked the design of the user interface, and 90% (18/20) reported that the tool was easy to use. Participants also had the most suggestions for this tool. The most common suggestion for the tool's functionality was the ability to create thresholds. For instance, four participants stated they would have liked to be able to set a threshold so that the tool would only display a notification when the information they were monitoring reached a certain value. Two participants also stated they would like for the tool to provide some history of the previous values for the information being monitored. Four participants reported that they would have liked a sound notification for both of the notification tools (i.e., the Text Clip and Page Updated tool). Participants reported that they sometimes found themselves glancing in the corner of the screen in anticipation of a notification and that an auditory notification would alleviate this.

Table 33. Collapsed Likert scale data from the post-study questionnaire investigating participant preferences for the Text Clip tool.

	Disagree	Neutral	Agree
Overall, I liked the functionality provided by the monitoring tool	0 (0.0%)	3 (15%)	17 (85.0%)
Overall, I liked the design of the user interface	0 (0.0%)	3 (15.0%)	17 (85.0%)
Overall, the tool was easy to use	0 (0.0%)	2 (10.0%)	18 (90.0%)

Two participants remarked on the complexity involved in monitoring textual information where several of the textual strings are dynamic. While monitoring the price of an eBay auction is fairly trivial from an implementation standpoint, monitoring a user's status on the student volunteer page is much more complex. For example, on the CHI student volunteer page, we selected the text string <position> <first name> <last name> to monitor (e.g., <140> <Melanie> <Kellar>). However, in practice, the monitoring tool would need to know which of the text string(s) to use as an anchor. In other words, would the user want to be notified when there is a change to the text adjacent to their position? Or would the use want to be notified of changes to the text adjacent to their name? It is assume that it would be the latter and in practice, the user would need to be able to communicate this to the monitoring tool.

8.4.4.2 Page Updated Tool

As shown in Table 32, the Page Updated tool was ranked third among the three tools. The results of the post-study questionnaire, shown in Table 34, indicate that participants were somewhat less enthusiastic for this tool. In particular, only 65% (13/20) of participants reported they liked the tool functionality. Seventy-five percent (15/20) of participants reported they liked the design of the user interface and 70% (14/20) reported that the tool was easy to use. We attribute the lower scores with respect to functionality to participants' desire for some indication of what had changed on the page. As evidence, the most common comment from participants when asked about the Page Updated tool was that they would have liked some indication of what had changed on the page (25% - 5/20). Two participants (10%) also indicated they would have liked to have known when the page had changed.

Table 34. Collapsed Likert scale data from the post-study questionnaire investigating participant preferences for the Page Updated tool.

	Disagree	Neutral	Agree
Overall, I liked the functionality provided by the monitoring tool	1 (5.0%)	6 (30.0%)	13 (65.0%)
Overall, I liked the design of the user interface	1 (5.0%)	4 (20.0%)	15 (75.0%)
Overall, the tool was easy to use	0 (0.0%)	6 (30.0%)	14 (70.0%)

8.4.4.3 Enhanced Bookmarks

As shown in Table 32, the Enhanced Bookmarks tool was ranked second among the three tools (marginally higher than the Page Updated tool). The results of the post-study questionnaire are shown in Table 35. Eight percent (16/20) of participants indicated they liked the functionality provided by the tool and all participants (20/20) reported that they liked the design of the interface. However, only 65% (13/20) of participants reported the Enhanced Bookmarks were easy to use. In fact, several participants (9/20) reported that they were either unsure of what the bars displayed by the Enhanced Bookmarks represented or they were unsure of what had changed on the page since their last visit. We expect that if participants became regular users of the Enhanced Bookmarks, they would develop a better understanding of the visual representation provided by the bookmarks. Also, if these were pages regularly visited by a user, they would most likely recognize new content more easily.

Table 35. Collapsed Likert scale data from the post-study questionnaire investigating participant preferences for the Enhanced Bookmarks tool.

	Disagree	Neutral	Agree
Overall, I liked the functionality provided by the monitoring tool	1 (5.0%)	3 (15.0%)	16 (80.0%)
Overall, I liked the design of the user interface	0 (0.0%)	0 (0.0%)	20 (100.0%)
Overall, the tool was easy to use	5 (25.0%)	2 (10.0%)	13 (65.0%)

8.5 Discussion

In this section, we first discuss the appropriateness of the monitoring tools for the task type. We then discuss the usability and utility of the three monitoring tools. Finally, based on the results of the study, we reflect on the design recommendations from Chapter 7.

8.5.1 Task-specific Monitoring Tools

Based on our previous study of web users' monitoring behaviour through semi-structured interviews, we hypothesized that different monitoring tools need different types of support. No participants reported that they would have preferred to use the same tool for all four monitoring tasks. One participant even remarked that he could not rank the three monitoring tools in order of preference because they all served different purposes. Another participant stated "Each [of the tools] had their own qualities or salient features. Each of them have their own place in browsing". The initial data collected during the study suggests that no single monitoring tool used in the experiment was appropriate for all monitoring tasks, which supports our hypothesis that task specific monitoring tools are beneficial.

8.5.1.1 Text Clip Tool

Participants strongly agreed that the functionality provided by the Text Clip tool was appropriate for the Fact Finding monitoring activities they completed. On the post-study questionnaire, 95% of participants agreed that the Text Clip tool was appropriate for the kind of information they were monitoring (changes to the student volunteer and eBay auction web page). During the post-study interviews, participants also mentioned a number of factual monitoring activities for which they could envision using the tool.

8.5.1.2 Page Updated Tool

Based on the study results, the suitability of the Page Updated tool for Maintenance monitoring tasks is unclear. The results of the post-study questionnaire are positive; 70% of participants agreed that the Page Updated tool was appropriate for the kind of information they were monitoring (i.e., changes to the tutorial web page). However, during the post-study interview, participants did not mention Maintenance tasks when they described tasks for

which they would find the Page Updated tool useful. Instead, they mentioned tasks synonymous with Browsing monitoring. It would appear that some participants did not appreciate the difference in the functionality provided by the Page Updated tool and the Enhanced Bookmarks. Two participants reported during the post-study interview that they did not have a strong preference between the Page Updated tool and the Enhanced Bookmarks.

These findings raise two questions. Does the design of the tool provide insufficient support for Maintenance monitoring tasks? Or, did our participants not engage in Maintenance monitoring tasks themselves, making them unable to appreciate the utility of the tool? While we do not know how many participants engaged in maintenance monitoring tasks as part of their usual monitoring activities, it was one of the less frequent monitoring activities reported in Chapter 7. Therefore, we would not expect that all participants would be familiar with this activity. It is clear that more research is needed to asses the suitability of the Page Updated tool for Maintenance monitoring activities.

8.5.1.3 Enhanced Bookmarks

The results of the study seem to indicate that the Enhanced Bookmarks tool provides appropriate support for Browsing monitoring activities. On the post-study questionnaire, 70% of participants agreed that the Enhanced Bookmarks tool was appropriate for the kind of information they were monitoring (e.g., blogs). During the post-study interviews, participants also mentioned a number of Browsing monitoring activities for which they could envision using the tool, such as sites that do not have RSS and news. Participants expressed some dissatisfaction with the Enhanced Bookmarks tool. Many of the issues raised by participants were usability issues, which are discussed in Section 8.5.2.

One interesting finding was that participants in this study seemed more positively disposed towards RSS than the participants who took part in our semi-structured interviews. Forty-five percent of the participants who took part in the laboratory study reported using some form of RSS, whereas only 10% of participants who took part in the semi-structured interviews from Chapter 7 reported using RSS. It is unclear if this is because RSS has

become more mainstream since the previous study[6], or if it is simply due to the population sampled. While the design of the Enhanced Bookmarks was somewhat based on our previous participants' aversion to RSS, the tool is still useful for the monitoring of web pages that are not suitable for RSS.

8.5.2 Usability and Utility of the Tools

8.5.2.1 Text Clip Tool

Participants were enthusiastic about the Text Clip tool and had several positive comments for the tool. Participant comments included "very easy to use", "I would find it useful", and "I was very impressed by the unobtrusive nature of the tool". Three participants even inquired if the tool was available for download. However, participants' enthusiasm for the tool could also be attributed to its novelty. Participants provided several suggestions to improve the utility of the Text Clip tool. This included the ability to provide a threshold for notifications of updated information and auditory cues in addition to the visual notification of updated information.

8.5.2.2 Page Updated Tool

Participants were not quite as enthusiastic about the Page Updated tool. In part, we think this is because many participants did not partake in Maintenance monitoring activities, so perhaps they did not appreciate the simple functionality provided by the tool. Future versions of the tool should investigate how to unobtrusively provide users with a representation of what has changed or is new on the web page. For example, one participant commented: "I didn't find the Page Updated tool very helpful, I still had to remember what week of the tutorial it was. I would want some sort of indication of what has changed." Many participants used the functionality provided by the Unix 'diff' program as an example of what they would have wanted.

---

[6] The interview were conducted in March and April 2006 while the laboratory study was conducted in December 2006 and January 2007.

8.5.2.3 Enhanced Bookmarks

Participants found the Enhanced Bookmarks tool useful, particularly for pages that did not offer RSS. Many participants would have liked some indication of what was updated or new on the page. For example, one participant stated that "highlighting changes on the page would be nice, draw attention to the new stuff". Participants also expressed concerns with the representation of the Enhanced Bookmarks. Participant comments included statements such as, "I'm curious as to what the bar measures exactly" and "Progress bar seemed to have somewhat arbitrary filling". While 100% of all participants reported that they liked the design of the user interface, more research is needed to understand how to visually represent the amount of change to a web page.

8.5.3 Reflections

In Section 8.2, we presented a table displaying the original findings and design recommendations from Chapter 7 as well as the tool functionality developed to support those findings. In Table 36, we reflect on the design recommendations and implemented tool functionality.

Table 36. For each monitoring activity (Fact Finding, Maintenance, Browsing), we reflect on the functionality developed based on the design recommendations.

Monitoring Activity	Implication and Recommendations from Chapter 7	Implemented Functionality	Reflections
Fact Finding	Fact Finding monitoring tools should only notify users when specific information on a web page has changed. In order to monitor factual information, this means that Fact Finding monitoring tools must allow users to: (1) Identify information to monitor. (2) Notify users when the information has changed. (3) Facilitate the delivery of the updated information.	(1) The tool allows users to highlight the information they would like to monitor. (2) A notification icon appears when the browser detects that the previously highlighted information has changed. (3) Users can view the updated information by clicking on the notification icon.	Fact Finding monitoring can be a complex activity and users may want to communicate that complexity to the tool. Therefore, tools could support users by allowing them to not only specify the information they would like to monitor, but also allow them to specify conditions, such as thresholds, end dates, etc.
	Notifications should display a discrete notification of updated content.	A notification icon alerts users to updated information but they must click on the icon to view the updated information.	Most participants found the notifications unobtrusive although some participants would have liked an auditory notification.
Maintenance	Users are usually interested in any changes to the web page; therefore, Maintenance monitoring tools should notify users when any information on the page has changed.	A notification icon appears when any information changes on a web page.	It is unclear whether our participants were familiar with Maintenance monitoring activities. However, many of our participants did express that they would have liked some indication of what had changed on the page. In order to be more effective, Maintenance monitoring tools could provide unobtrusive markers on the page, which could be toggled on/off, to indicate what has changed.
	It is important that the look and feel of the web page is not modified in such a way that the user does not see a true representation of the web page.	The tool does not modify the web page to indicate what has been modified on the web page.	
	Notifications should display a discrete notification of updated content.	A notification icon alerts users to updated information but they must click on the icon to view a message that the page has changed.	Most participants found the notifications unobtrusive although some participants would have liked an auditory notification.

Monitoring Activity	Implication and Recommendations from Chapter 7	Implemented Functionality	Reflections
Browsing	Due to the serendipitous nature of Browsing activities, it may be very difficult for monitoring tools to accurately predict what is of interest to a user. Therefore, instead of identifying content of interest, tools should provide users with an awareness of a web page's update activity.	Users are provided with an abstract representation of the percentage of the web page content that has changed since their previous visit.	Many participants had difficulty with the abstract representation of changes to the web page (i.e., the bars). More research is needed to better understand how this visualization could be more effectively delivered to users.
	Given the long term nature of monitoring activities and the common use of bookmarks, functionality to support Browsing monitoring could be provided through bookmarks.	The abstract representation (mentioned above) is integrated with the bookmark list.	Participants were positively disposed to the Enhanced Bookmarks. It is also important that this functionality is available with other bookmark formats, such as the Links toolbar. Some of our participants also reported using the LiveBookmarks feature within Firefox.
	With the exception of pages that updated infrequently, Browsing monitoring tools should not overload users through continual notifications of updates to web pages.	The tool is an awareness tool and does not issue any notifications to the user.	Participants did appreciate the awareness component of the tool.

8.6 Limitations

When designing this study, we considered the most appropriate methodology for evaluating prototype versions of the monitoring tools. Studying monitoring behaviour in a controlled setting is very difficult because the nature of monitoring requires continuous changes to web pages over elapsed periods of time. We chose to conduct a laboratory study since it enabled us to develop prototype versions of the tools and allowed us to ensure consistency between participant sessions using the instrument web browser.

While the tasks, web pages, and monitoring tool functionality was contrived, participants did appear to be engaged in the tasks during the study and many did not realize that the monitoring tools were not fully functioning. However, it is still important to acknowledge that participants did not have a vested interest in either the primary or secondary monitoring tasks and their perceptions of the tools may differ from self-motivated tasks. For instance, participants had trouble understanding the amount of change displayed by the Enhanced Bookmarks. However, with regular use, they may have come to understand exactly what is meant by "half a bar" on "web page x".

While the laboratory study provides an initial examination of the use of the tools, a more naturalistic study will help to fully understand the impact of these tools on users' monitoring behaviour. Therefore, a field study is needed to observe participants using the tools in a natural setting with self-motivated tasks.

8.7 Summary

In this chapter, we have presented three prototype task-specific monitoring tools and the results of a laboratory study conducted to examine: (1) whether task-specific monitoring tools are appropriate, (2) whether the three task specific monitoring tools are appropriate for the type of information being monitored, and (3) how we can improve the features of the monitoring tools. The results of this study appear to confirm the results of our previous study and suggest that task specific monitoring tools are useful and that no one monitoring tool can be used for all information monitoring tasks. Our evaluation of the Text Clip tool suggests that its functionality is appropriate for Fact Finding monitoring activities.

Participant responses seemed to indicate that the Enhanced Bookmarks were also appropriate for Browsing monitoring activities. Participants were somewhat less enthusiastic about the use of the Page Updated tool for Maintenance monitoring. There were several suggestions for improvements to the interface and functionality for all three monitoring tools. Based on participant responses, we have also reflected on the design implications developed in Chapter 7.

Chapter 9

Conclusion

9.1 Summary

The research presented in this thesis is the result of three successive studies conducted to examine user behaviour on the Web in the context of task. The three studies presented provide new understanding of (1) the tasks users engage in on the Web, (2) how users interact with their web browsers to complete these tasks, and (3) how we can potentially better support users during these tasks.

In Chapter 2, we presented an overview of previous research that has studied information seeking on the Web and web-based monitoring. This review highlighted several gaps in the literature, including: discrepancies in the previous models of information seeking tasks on the Web, a weak understanding of web-based monitoring, and a lack of evaluation of the existing tools that support information seeking and web-based monitoring activities.

In Chapter 3, we provided an overview of the common research strategies used to study user behaviour on the Web. Many of the tradeoffs discussed in this chapter influenced our choice of methodology when designing the three studies that contributed to this thesis.

In Chapter 4, we described the design of a week long field study conducted with 21 participants. Over the course of the study, participants annotated their web usage with a task categorization (Fact Finding, Information Gathering, Browsing, Transactions, and Other)

and a short textual description of the task. This study was designed with two research goals in mind. The first goal was to develop a better understanding of types of tasks in which web users engage, while the second goal was to examine how users interact with their web browsers during these tasks.

In Chapter 5, we presented the first set of results from the field study. We examined how users interacted with their web browsers through the following implicit measures: dwell time, windows opened, pages loaded, use of web browser navigation tools, time of day, use of Google, use of site specific searches, and use of web browser functions. We found significant differences in these measures between different tasks, revealing several distinguishing characteristics for each task type. Fact Finding tasks were short-lived and search-based with also a heavy use of Google. Search queries within Fact Finding tasks were also the longest. Information Gathering tasks were the longest in duration and were also search-based with a heavy use of Google. Information Gathering tasks were characterized by a large number of pages viewed and the greatest use of browser functions (e.g., copy, print, save). Browsing tasks were short-lived, with a high level of revisitation. Transactions, consisting primarily of email, were short in duration and were the most often repeated task.

During the field study, participants used a task categorization that was developed through the pilot study and focus group presented in Chapter 4. Using the task data collected from participants over the course of the field study, we refined the categorization and developed the Web Information Classification. The classification consists of three web information goals: information seeking, information exchange, and information maintenance. Information seeking consist of tasks in which users are trying to change their state of knowledge and including Fact Finding, Information Gathering, and Browsing. Information exchange consists of online actions including Transactions and Communications. Information maintenance consists of visits to web pages to maintain information on the page.

In Chapter 6, we conducted a more detailed analysis of the use of web browser navigation mechanisms in the context of task, which was a relatively unstudied area. We identified three factors that play a role in the use of web browser navigation mechanism: whether a user is initiating a new task session or currently within a task session, the task type, and individual

differences. In both Chapter 5 and Chapter 6, we reflected on the results of our analyses and presented implications for the design and evaluation of new functionality to support users in their web information tasks.

In Chapter 7, we presented the results of a study designed to better understand the role of monitoring in the context of task. Based on the results of our previous study, we hypothesized that Monitoring was an activity within web information tasks, and not an independent task onto itself. We conducted a series of semi-structured interviews conducted with 40 participants from four sample populations. During the interviews, participants were asked to describe several aspects of their monitoring activities, such as the goal of the monitoring activity; the type of information being monitored; their use of navigation mechanisms; the duration and frequency of the activity; the use of user logins and search queries; follow-up activities; and the nature of the activity (e.g., work, school, or personal). The result of our study confirmed our hypothesis that monitoring is an activity that occurs within all web information tasks. The data collected yielded several recommendations for the design of task-specific monitoring tools.

In Chapter 8, we presented the design and development of three task-specific monitoring tools to support Fact Finding (Text Clip), Maintenance (Page Updated), and Browsing (Enhanced Bookmarks) monitoring activities. The design of these tools was based on the recommendations developed in Chapter 7. A laboratory experiment was conducted with 20 participants to evaluate the functionality of the task-specific monitoring tools. The results of the study determined that the use of task-specific monitoring tools is appropriate. In particular, participants were very positively disposed towards the Text Clip tool, which supports Fact Finding monitoring activities. Participants provided several suggestions for improvement to all three monitoring tools. Based on the study results and participant comments, we concluded this chapter with reflections upon the design recommendations developed in Chapter 7 .

9.2 Research Contributions

This thesis has made contributions to the field of Human-Computer Interaction (HCI), Information Science, and methodologies for studying user behaviour on the Web. In this

section, we describe the theoretical, applied, and methodological contributions of this research.

9.2.1 Theoretical Contributions

In this section we describe the two theoretical contributions, which include the development of the Web Information Classification and the characterization of web information tasks.

9.2.1.1 Web Information Classification

The Web Information Classification provides a high level categorization of web information tasks. This design of the classification was informed by both previous work (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002) and the research presented in this thesis. In Chapter 5, we present our initial version of the classification, which was based upon the results of the field study. In Chapter 7, we presented an updated Web Information Classification, which includes monitoring as an activity within all web information tasks.

Studies of the types of tasks in which users engage in on the Web provide great benefit to other researchers studying user behaviour on the Web. The Web Information Classification provides a task-based framework on which other researchers can build. These types of studies are often relied upon in many areas of research, such as the WWW, HCI, and Information Science research communities. As evidence, the three previous studies that heavily influenced our work (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002) have almost 200 citations combined by researchers studying a wide variety of research topics.

9.2.1.2 Characterizing Web Information Tasks

The results of our research have allowed us to better describe the characteristics of web information tasks and monitoring activities. The results of the field study presented in Chapter 5 provides insight into how users interact with their web browsers during web information tasks and also highlight the differences in these interactions between tasks. While the results of the semi-structured interviews in Chapter 7 are self-reports, they do

provide an initial characterization of how users employ their web browsers to complete their monitoring activities. Our results have shown that there are significant differences in how users interact with their web browser across tasks. While described further in Section 9.3.1 (Future Work), these differences may be useful in predicting the nature of the task in which a user is engaging.

9.2.2 Applied Contributions

In this section we describe the three sets of recommendations, which may be used to inform (1) the design of tools to support web information tasks, (2) the design and evaluation of web browser navigation mechanisms, and (3) the design of tools to support web-based monitoring activities.

9.2.2.1 Recommendations for Supporting Web Information Tasks

The characterization of web information tasks presented in Chapter 5 was used to develop recommendations and implications for the future design of tools to support web information tasks. The recommendations highlight the importance of supporting: dominant task attributes (e.g., search vs. revisitation); repeated tasks; and complex information seeking tasks, such as Information Gathering. We also highlight the lack of research examining how to improve the infrequently used history function and windows management within a task session.

9.2.2.2 Recommendations for the Design and Evaluation of Web Browser Navigation Mechanisms

In Chapter 6, we examined the impact of task and individual differences on the use of web browser navigation mechanisms. Based on these findings, we developed recommendations and implications for the design and evaluation of web browser navigation mechanisms. The recommendations highlight the importance of: the evaluation and interpretation of the use of web browser navigation mechanisms in the appropriate context; the importance of support task characteristics; and supporting individual differences.

9.2.2.3 Recommendations for Supporting Web-based Monitoring

In Chapter 7, we examined the role of monitoring within the Web Information Classification. Based on the findings from the semi-structured interviews, we developed a series of general and task-specific recommendations for the design of future web-based monitoring tools. The general recommendations aim to circumvent the typical reported barriers to the use of monitoring tools, which include accuracy of monitored information, loss of screen real estate, method of notification, tool functionality, and setup effort. The task-specific recommendations provide specific recommendations for supporting the characteristics of monitoring activities within Fact Finding, Information Gathering, Browsing, Communications, Transactions, and Maintenance. In Chapter 8, we also provide an evaluation of the functionality developed based on the recommendations for Fact Finding, Browsing, and Maintenance monitoring.

We expect that these three sets of recommendations and implications will be useful to both designers and researchers interested in better supporting web information tasks, web navigation, and web-based monitoring activities.

9.2.3 Methodological Contributions

The methodological approach used to conduct the field study provides a methodological contribution to the research community. We developed a custom-web browser that logs fine grained interactions that is available for use by other researchers in the field. As well, the description of our experiences in collecting task information and web browser interactions provides a model for other researchers to follow. We have reflected on our experiences in studying user behaviour on the Web in a natural setting in two workshop papers (Kellar, 2006; Kellar & Watters, 2005), which we expect may benefit other researchers conducting similar work.

9.3 Future Work

In this section we outline the five major areas of future work: (1) predicting task using implicit measures; (2) examining information seeking for mobile web users; (3) developing

integrated monitoring tools; (4) supporting efficient web use; and (5) monitoring peripheral information.

9.3.1 Predicting Task Using Implicit Measures

The automatic identification of a user's task has the potential to improve information filtering systems. Many information filtering systems rely on implicit measures of interest, whose effectiveness may be dependant upon the task at hand (Kellar, Watters, Duffy, & Shepherd, 2004; Kelly & Belkin, 2001). An understanding of the user's current task would allow the system to apply the most useful measures of interest. In Appendix L, we present a preliminary exploration of the usefulness of implicit measures to predict user task types using the data collected during the field study. While an aggregate model did not produce accurate enough results (accuracy = 53.8%), individual decision trees models have shown promise. The number of correctly classified tasks for each participant ranged from a high of 94.3% to a low of 43.6% (Kellar & Watters, 2006). We would like to further explore the use of machine learning techniques to improve the accuracy of the task predictions. We would also like to incorporate a categorization of the web pages visited, potentially according to web page genre.

9.3.2 Information Seeking for Mobile Web Users

The research presented in this thesis was focused on understanding the web activities of desktop users. While it is clear that there are fundamental differences between desktop and mobile web users, it is not apparent how those differences impact the type of support needed by mobile web users. We would like to extend this research by examining the types of tasks typically undertaken by mobile web users. In particular, we are interested in how the well the Web Information Classification models the web usage of mobile users. We expect that the information needs of mobile users may be very different from those of desktop users. Other considerations include the small screen size, which may impact how search interfaces and search results are presented to mobile users. GPS functionality can provides important contextual information, such as geographic information, which may be useful in improving information seeking on the Web for mobile users.

9.3.3 Integrated Monitoring Tools

In Chapter 8, we conducted a preliminary evaluation of three task-specific monitoring tools. While the results of the evaluation reinforced the appropriateness of task-specific monitoring tools, it is not our intent to simply build a series of "one off" web browser tools. The development of three separate tools was appropriate for the purpose of the study but a more integrated solution is desirable. For example, a Swiss Army Knife approach may be more useful, where several task-specific tools are integrated into a single system. Users could then choose the monitoring tool that best supports the information they would like to monitor. Future work should investigate how this can be achieved. Scope (van Dantzich, Robbins, Horvitz, & Czerwinski, 2002), which was presented in Chapter 2, is an example of monitoring system integrates and displays notifications from multiple sources.

9.3.4 Supporting Efficient Web Use

During the semi-structured interviews described in Chapter 7, a small number of participants (12.5%) reported that monitoring was a compulsive activity. We hypothesize that effective monitoring tools may be useful in mitigating this compulsive behaviour. In particular, tools that are accurate and reliable could reduce the amount of time users spend monitoring information on the Web. In addition, tools that deliver updated information to users without requiring them to navigate the Web may also reduce the amount of serendipitous browsing, leading to more efficient web usage.

Before we can study the impact of monitoring tools on the efficiency of web users, we must first develop metrics with which to measure this impact. For example, we hypothesize that the effectiveness of a monitoring tool can be measured in part, by the reduction in the number and length of web page visits to frequently monitored web pages. In order to fully understand the impact of these tools, we would like to develop strong and reliable metrics to study web-based monitoring.

9.3.5 Monitoring Peripheral Information

As the availability of dynamic information increases on the Web, it is important that we continue to study how to effectively support information monitoring. This thesis research

examined monitoring in a traditional display environment (e.g., single/dual monitor). One of the challenges of information monitoring is the fact that display space in traditional environments is at a premium; users cannot always afford to monitor information on a continual basis. Instead, monitored information is often retrieved as needed, meaning users may interrupt their task at hand in order to complete a monitor activity. This leads to distractions and provides an entry point for serendipitous web browsing.

Augmented surfaces (Rekimoto & Saitoh, 1999) offer novel opportunities for information monitoring by allowing users to extend their display space onto other surfaces in the environment such as tables and walls. Through a longitudinal study, we would like to evaluate in how the introduction of an awareness monitoring tool impact users' web monitoring activities, particularly in augmented display environments. It is expected that the introduction of augmented displays for information monitoring will allows users to be more efficient in their monitoring activities, particularly in workplace settings.

9.4 Conclusions

We believe there is great potential for the improvement of standard web browser tools and functionality. However, researchers must first be able to characterize the tasks in which users engage on the Web and understand how the current features of web browsers are used during different web tasks. In this thesis, we have presented the result of three successive studies that together provide insight into (1) the tasks in which users engage on the Web; (2) how users interact with their web browsers to accomplish these tasks; and (3) recommendations for how we can potentially better support users during these tasks.

References

Abrams, D., Baecker, R. and Chignell, M. (1998). Information Archiving with Bookmarks: Personal Web Space Construction and Organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '98)*, Los Angeles, CA, 41 - 48.

Amento, B., Terveen, L., Hunt, W. and Hix, D. (2000). Topicshop: Enhanced Support for Evaluating and Organizing Collections of Web Sites. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology*, San Diego, CA, 201-209.

Anick, P. (2003). User Terminological Feedback for Web Search Refinement - a Log-Based Study. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '03)*, Toronto, Canada, 88-95.

Apple Inc. (2007). Mac Os X Dashboard Web Clips: Retrieved September 27, 2006, from http://www.apple.com/macosx/

Atterer, R., Wnuk, M. and Schmidt, A. (2006). Knowing the User's Every Move: User Activity Tracking for Website Usability Evaluation and Implicit Interaction. In *Proceedings of the 15th International Conference on World Wide Web (WWW 2006)*, Edinburgh, Scotland, 203-212.

Aula, A., Jhaveri, N. and Käki, M. (2005). Information Search and Re-Access Strategies of Experienced Web Users. In *Proceedings of the 14th International World Wide Web Conference (WWW 2005)*, Chiba, Japan, 583 - 592.

Aula, A. and Käki, M. (2003). Understanding Expert Search Strategies for Designing User-Friendly Search Interfaces. In *Proceedings of IADIS International Conference WWW/Internet (ICWI 2003)*, Algarve, Portugal, 759-762.

Ayers, E. Z. and Stasko, J. T. (1995). Using Graphic History in Browsing the World Wide Web. In *Proceedings of the 4th International World Wide Web Conference*, Boston, MA.

Barford, P., Bestavros, A., Bradley, A. and Crovella, M. (1999). Changes in Web Client Access Patterns: Characteristics and Caching Implications. *World Wide Web*, 2(1-2): 15-28.

Beitzel, S. M., Jensen, E. C., Chowdhury, A., Grossman, D. and Frieder, O. (2004). Hourly Analysis of a Very Large Topically Categorized Web Query Log. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2004)*, Sheffield, United Kingdom, 321-328.

Bernard, H. R. (2000). *Social Research Methods: Qualitative and Quantitative Approaches*, Sage Publications.

Boardman, R. and Sasse, M. A. (2004). "Stuff Goes into the Computer and Doesn't Come Out": A Cross-Tool Study of Personal Information Management. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*, Vienna, Austria, 583-590.

Boyapati, V., Chevrier, K., Finkel, A., Glance, N., Pierce, T., Stockton, R. and Whitmer, C. (2002). ChangeDetector™: A Site-Level Monitoring Tool for the WWW. In *Proceedings of the 11th International Conference on World Wide Web (WWW 2002)*, 570-579.

Broder, A. (2002). A Taxonomy of Web Search. *SIGIR Forum*, 36(2): 3-10.

Byrne, M., John, B., Wehrle, N. and Crow, D. (1999). The Tangled Web We Wove: A Taskonomy of WWW Use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, Pittsburgh, PA, 544-551.

Byström, K. and Hansen, P. (2005). Conceptual Framework for Tasks in Information Studies. *Journal of the American Society for Information Science and Technology*, 56(10): 1050 - 1061.

Catledge, L. D. and Pitkow, J. E. (1995). Characterizing Browsing Strategies in the World-Wide Web. In *Proceedings of the Third International World-Wide Web Conference*, Darmstadt, Germany, 1065 - 1073.

Chakravarthy, S., Sanka, A. and Jacob, A. J. (2004). A Learning-Based Approach for Fetching Pages in WebVigiL. In *Proceedings of the ACM Symposium on Applied Computing (SAC 2004)*, Nicosia, Cyprus, 1725-1731.

ChangeDetect (2007). ChangeDetect Web Page Monitoring: Retrieved February 19, 2007, from http://www.changedetect.com/

Choo, C. W., Detlor, B. and Turnbull, D. (2000). Information Seeking on the Web: An Integrated Model of Browsing and Searching. *First Monday*. 5(2): Retrieved August 3, 2004, from http://www.firstmonday.org/issues/issues2005_2002/choo/index.html

Claypool, M., Le, P., Waseda, M. and Brown, D. (2001). Implicit Interest Indicators. In *Proceedings of the 6th International Conference on Intelligent User Interfaces (IUI 2001)*, Santa Fe, NM, 33-40.

Cockburn, A. and McKenzie, B. (2001). What Do Web Users Do? An Empirical Analysis of Web Use. *International Journal of Human-Computer Studies*, 54(6): 903-922.

Cockburn, A., McKenzie, B. and JasonSmith, M. (2002). Pushing Back: Evaluating a New Behavior for the Back and Forward Buttons in Web Browsers. *International Journal of Human-Computer Studies*, 57(5): 397-414.

Cothey, V. (2002). A Longitudinal Study of World Wide Web Users' Information-Searching Behavior. *Journal of the American Society for Information Science and Technology*, 53(2): 67-78.

Ellis, D. (1989). A Behavioural Approach to Information Retrieval System Design. *Journal of Documentation*, 45(3): 171-212.

Ellis, D., Cox, D. and Hall, K. (1993). A Comparison of the Information Seeking Patterns of Researchers in the Physical and Social Sciences. *Journal of Documentation*, 49(4): 356-369.

Fallows, D. (2006). Internet Users in Search of a Home. *Pew Internet & American Life Project*: Retrieved online February 27, 2007, from http://www.pewinternet.org/PPF/r/2196/report_display.asp

Fenstermacher, K. and Ginsburg, M. (2003). Client-Side Monitoring for Web Mining. *Journal of the American Society for Information Science and Technology*, 54(7): 625-637.

Fidel, R. and Efthimiadis, E. (1999). Web Searching Behavior of Aerospace Engineers. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '99)*, Berkley, CA, 319-320.

Flanagan, J. C. (1954). The Critical Incident Technique. *Psychological Bulletin*, 51(4): 327-358.

Ford, N., Wilson, T. D., Foster, A., Ellis, D. and Spink, A. (2002). Information Seeking and Mediated Searching. Part 4. Cognitive Styles in Information Seeking. *Journal of the American Society for Information Science and Technology*, 53(9): 728-735.

Google (2007). Personalizing Your Google Homepage: Retrieved February 19, 2007, from http://www.google.com/support/bin/topic.py?topic=1592

Grace-Martin, M. and Gay, G. (2001). Web Browsing, Mobile Computing and Academic Performance. *Educational Technology & Society*. 4(3): Retrieved February 19, 2006, from http://ifets.ieee.org/periodical/vol_2003_2001/grace_martin.html

Greenberg, S. and Boyle, M. (2006). Generating Custom Notification Histories by Tracking Visual Differences between Web Page Visits. In *Proceedings of Graphics Interface (GI 2006)*, Quebec, Canada, 227 - 234.

GVU's WWW User Surveys (2001). Retrieved August 13, 2005, from http://www.cc.gatech.edu/gvu/user_surveys/

Hawkey, K. and Inkpen, K. (2005a). Privacy Gradients: Exploring Ways to Manage Incidental Information During Co-Located Collaboration. In *Proceedings of the CHI '05 Extended Abstracts on Human Factors in Computing Systems*, Portland, OR, 1431 - 1434.

Hawkey, K. and Inkpen, K. M. (2006). Examining the Content and Privacy of Web Browsing Incidental Information. In *Proceedings of 15th International Conference on World Wide Web (WWW 2006)*, Edinburgh, Scotland, 123 - 132.

Hawkey, K. and Inkpen, K. M. (2005b). Web Browsing Today: The Impact of Changing Contexts on User Activity. In *Proceedings of the CHI '05 Extended Abstracts on Human Factors in Computing Systems*, Portland, OR, 1443-1446.

Herder, E. (2005). Characterizations of User Web Revisit Behavior. In *Proceedings of the Workshop on Adaptivity and User Modeling in Interactive Systems (ABIS 2005)*, Saarbrücken, Germany.

Herder, E. and Juvina, I. (2004). Discovery of Individual User Navigation Styles. In *Proceedings of the Workshop on Individual Differences in Adaptive Hypermedia (Adaptive Hypermedia 2004)*, Eindhoven, The Netherlands.

Hightower, R. R., Ring, L. T., Helfman, J. I., Bederson, B. B. and Hollan, J. D. (1998). Graphical Multiscale Web Histories: A Study of PadPrints. In *Proceedings of the 9th ACM Conference on Hypertext and Hypermedia*, 58-65.

Hoeber, O. and Yang, X. D. (2006). Interactive Web Information Retrieval Using WordBars. In *Proceedings of the International Conference on Web Intelligence (WI 2006)*, Hong Kong, China, 875-882.

Hölscher, C. and Strube, G. (2000). Web Search Behavior of Internet Experts and Newbies. In *Proceedings of the 9th International World Wide Web Conference on Computer Networks*, Amsterdam, The Netherlands, 337-346.

Hong, J. I., Heer, J., Waterson, S. and Landay, J. A. (2001). WebQuilt: A Proxy-Based Approach to Remote Web Usability Testing. *ACM Transactions on Information Systems*, 19(3): 263-285.

Horrigan, J. B. (2006). The Internet as a Resource for News and Information about Science. *PEW Internet & American Life Project*: Retrieved February 27, 2007, from http://www.pewinternet.org/PPF/r/2191/report_display.asp

Horrigan, J. B. and Morris, S. (2005). 13 Million Americans Made Donations Online after Hurricanes Katrina and Rita. *Pew Internet & American Life Project*: Retrieved online February 19, 2007, from http://www.pewinternet.org/PPF/r/2168/report_display.asp

Jansen, B. J. and McNeese, M. D. (2005). Evaluating the Effectiveness of and Patterns of Interactions with Automated Searching Assistance. *Journal of the American Society for Information Science and Technology*, 56(14): 1480-1503.

Jansen, B. J. and Spink, A. (2003). An Analysis of Web Documents Retrieved and Viewed. In *Proceedings of the 4th International Conference on Internet Computing*, Las Vegas, NV, 65-69.

Jansen, B. J., Spink, A., Bateman, J. and Saracevic, T. (2000). Real Life, Real Users, and Real Needs: A Study and Analysis of User Queries on the Web. *Information Processing and Management*, 36: 207-227.

Jansen, B. J., Spink, A. and Pedersen, J. (2005). A Temporal Comparison of AltaVista Web Searching. *Journal of the American Society for Information Science and Technology*, 56(6): 559-570.

Jhaveri, N. and Räihä, K.-J. (2005). The Advantages of a Cross-Session Web Workspace. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2005)*, Portland, OR, 1949-1952.

Jones, W., Dumais, S. and Bruce, H. (2002). Once Found, What Then?: A Study Of "Keeping" Behaviors in the Personal Use of Web Information. In *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2002)*, Philadelphia, PA, 391-402.

Kaasten, S., Greenberg, S. and Edwards, C. (2002). How People Recognize Previously Seen Web Pages from Titles, URLs and Thumbnails. In *Proceedings of Human Computer Interaction 2002*, London, UK, 247-265.

Kalbach, J. (2000). Designing for Information Foragers: A Behavioral Model for Information Seeking on the World Wide Web. *Internetworking*. 3(3): Retrieved September 12, 2006, from http://www.internettg.org/newsletter/dec2000/article_information_foragers.html

Kellar, M. (2006). Logging IE: Barriers to Studying User Behaviour on the Web. Workshop on *Logging Traces of Web Activity: The Mechanics of Data Collection* at the ACM International World Wide Web Conference (WWW 2006), Edinburgh, Scotland.

Kellar, M., Hawkey, K., Inkpen, K. M. and Watters, C. (In Press). Challenges of Capturing Natural Web-Based User Behaviours. In Use, in Situ: Extending Field Research Methods *International Journal of Human Computer Interaction.*

Kellar, M. and Watters, C. (2005). Studying User Behaviour on the Web: Methods and Challenges. Workshop on *Usage Analysis: Combining Logging and Qualitative Methods* at the ACM Conference on Human Factors in Computing Systems (CHI 2006). Portland, OR.

Kellar, M. and Watters, C. (2006). Using Web Browser Interactions to Predict Task. In *Proceedings of the ACM International World Wide Web Conference (WWW 2006)*, Edinburgh, Scotland, 843-844.

Kellar, M., Watters, C., Duffy, J. and Shepherd, M. (2004). Effect of Task on Time Spent Reading as an Implicit Measure of Interest. In *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2004)*, Providence, RI, 168-175.

Kellar, M., Watters, C. and Inkpen, K. M. (2007). An Exploration of Web-Based Monitoring: Implications for Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2007)*, To Appear.

Kellar, M., Watters, C. and Shepherd, M. (In Press). A Field Study Characterizing Web-Based Information-Seeking Tasks. *Journal of the American Society for Information Science and Technology.*

Kellar, M., Watters, C. and Shepherd, M. (2006a). A Goal-Based Classification of Web Information Tasks. In *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2006)*, Austin, TX, (CD-ROM Proceedings).

Kellar, M., Watters, C. and Shepherd, M. (2006b). The Impact of Task on the Usage of Web Browser Navigation Mechanisms. In *Proceedings of Graphics Interface 2006 (GI 2006)*, Quebec City, Canada, 235-242.

Kelly, D. and Belkin, N. J. (2004). Display Time as Implicit Feedback: Understanding Task Effects. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2004)*, Sheffield, United Kingdom, 377-384.

Kelly, D. and Belkin, N. J. (2001). Reading Time, Scrolling and Interaction: Exploring Implicit Sources of User Preferences for Relevance Feedback. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2001)*, New Orleans, LA, 408-409.

Kelly, D. and Teevan, J. (2003). Implicit Feedback for Inferring User Preference: A Bibliography. *SIGIR Forum*, 37(2): 18-28.

Kim, K.-S. and Allen, B. (2002). Cognitive and Task Influences on Web Searching Behavior. *Journal of the American Society for Information Science and Technology*, 53(2): 109-119.

Komlodi, A., Soergel, D. and Marchionini, G. (2006). Search Histories for User Support in User Interfaces. *Journal of the American Society for Information Science and Technology*, 57(6): 803 - 807.

Kuhlthau, C. (1991). Inside the Search Process: Information Seeking from the User's Perspective. *Journal of the American Society for Information Science and Technology*, 42(5): 361-371.

Lee, U., Liu, Z. and Cho, J. (2005). Automatic Identification of User Goals in Web Search. In *Proceedings of the 14th International Conference on World Wide Web*, Chiba, Japan, 391 - 400.

Lenhart, A. and Madden, M. (2007). Social Networking Websites and Teens. *Pew Internet & American Life Project*: Retrieved February 27, 2007, from http://www.pewinternet.org/PPF/r/2198/report_display.asp

Liesegang, S. J. M. (2006). Morning Coffee Firefox Add-Ons: Retrieved September 27, 2006 from https://addons.mozilla.org/firefox/2677

MacKay, B., Kellar, M. and Watters, C. (2005). An Evaluation of Landmarks for Re-Finding Information on the Web. In *Proceedings of the CHI '05 Extended Abstracts on Human Factors in Computing Systems*, Portland, OR, 1609 - 1612.

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

Marchionini, G. (1995). *Information Seeking in Electronic Environments.* New York, NY, Cambridge University Press.

Mat-Hassan, M. and Levene, M. (2005). Associating Search and Navigation Behavior through Log Analysis. *Journal of the American Society for Information Science and Technology*, 56(9): 913-934.

McCrickard, D. S., Chewar, C. M., Somervell, J. P. and Ndiwalana, A. (2003). A Model for Notification Systems Evaluation - Assessing User Goals for Multitasking Activity. *ACM Transactions on Computer-Human Interaction*, 10(4): 312-338.

McGrath, J. E. (1995). Methodology Matters: Doing Research in the Behavioral and Social Sciences. In J. G. R. Baeker, J. Grudin, W. Buxton, and S. Greenberg (Eds.), *Human-Computer Interaction: Toward the Year 2000*, Morgan Kaufmann Publishers Inc., 152-169.

Meho, L. and Tibbo, H. (2003). Modeling the Information-Seeking Behaviour of Social Scientists: Ellis' Study Revisisted. *Journal of the American Society for Information Science and Technology*, 54(6): 570-587.

Microsoft Developer Network (2007). MSDN Webbrowser Control Overview: Retrieved February 18, 2007, from http://msdn2002.microsoft.com/en-us/library/w2290k2023d.aspx

Milic-Frayling, N., Sommerer, R. and Rodden, K. (2003). Webscout: Support for Revisitation of Web Pages within a Navigation Session. In *Proceedings of Web Intelligence 2003 (WI 2003)*, Halifax, Canada, 689-693.

Montgomery, A. L. and Faloutsos, C. (2001). Identifying Web Browsing Trends and Patterns. *IEEE Computer*, 34(7): 94-95.

Morita, M. and Shinoda, Y. (1994). Information Filtering Based on User Behavior Analysis and Best Match Text Retrieval. In *Proceedings of 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '94)*, Dublin, Ireland, 272-281.

Morrison, J. B., Pirolli, P. and Card, S. K. (2001). A Taxonomic Analysis of What World Wide Web Activities Significantly Impact People's Decisions and Actions. In *Proceedings of the CHI '01 Extended Abstracts on Human Factors in Computing Systems*, Seattle, WA, 163-164.

Moyle, M. and Cockburn, A. (2003). The Design and Evaluation of a Flick Gesture for 'Back' and 'Forward' in Web Browsers. In *Proceedings of the 4th Australasian Conference on User Interface (AUIC 2003)*, Adelaide, Australia, 39-46.

Mozdev (2004). Forecastfox: Retrieved February 1, 2006, from http://forecastfox.mozdev.org

Mozdev (2006). Page Update Checker: Retrieved February 1, 2006, from http://pageupdatechecker.mozdev.org/

Netvibes (2007). Retrieved February 18, 2007, from http://www.netvibes.com/

Newfield, D., Sethi, B. S. and Ryall, K. (1998). Scratchpad: Mechanisms for Better Navigation in Directed Web Searching. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology (UIST '98)*, San Francisco, CA, 1-8.

Noyes, J. and Baber, C. (1999). *User-Centred Design of Systems*, Springer.

Oard, D. W. and Kim, J. (2001). Modeling Information Content Using Observable Behavior. In *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2001)*, Washington, DC, 38-45.

Obendorf, H., Weinreich, H. and Hass, T. (2004). Automatic Support for Web User Studies with Scone and Tea. In *Proceedings of the CHI '04 Extended Abstracts on Human Factors in Computing Systems*, Vienna, Austria, 1135-1138.

OneStat.com (2004). Mozilla's Browsers Global Usage Share Is 7.35 Percent According to OneStat.com: Retrieved August 10, 2005, from http://www.onestat.com/html/aboutus_pressbox2034.html

Oppenheim, A. N. (1992). *Questionnaire Design, Interviewing, and Attitude Measurement*, Pinter.

Pitkow, J. E. and Kehoe, C. M. (1996). Emerging Trends in the WWW User Population. *Communications of the ACM*, 39(6): 106 - 108.

Plaue, C., Miller, T. and Stasko, J. (2004). InfoCanvas: A Highly Personalized, Elegant Awareness Display. In *Proceedings of Graphics Interface (GI 2004)*, London, Canada, 117-126.

Plaue, C. and Stasko, J. (2006). BlueGoo: Increasing Visual Interest of Peripheral Displays: Retrieved February 19, 2007, from http://www.gvu.gatech.edu/ii/bluegoo/

Reeder, R. W., Pirolli, P. and Card, S. K. (2001). WebEyeMapper and WebLogger: Tools for Analyzing Eye Tracking Data Collected in Web-Use Studies. In *Proceedings of the CHI '01 Extended Abstracts on Human Factors in Computing Systems*, Seattle, WA, 19-20.

Rekimoto, J. and Saitoh, M. (1999). Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99)*, Pittsburgh, PA, 378 - 385.

Rieh, S. Y. (2003). Investigating Web Searching Behaviour in Home Environments. In *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2003)*, Long Beach, CA, 255-264.

Rose, D. E. (2006). Reconciling Information-Seeking Behavior with Search User Interfaces for the Web. *Journal of the American Society for Information Science and Technology*, 57(6): 797-799.

Rose, D. E. and Levinson, D. (2004). Understanding User Goals in Web Search. In *Proceedings of 13th International Conference on World Wide Web (WWW 2004)*, New York, NY, 13-19.

Rozanski, H. D., Bollman, G. and Lipman, M. (2001). Seize the Occasion! The Seven-Segment System for Online Marketing. *Strategy & Competition*: Retrieved June 23, 2005, from http://www.strategy-business.com/press/article/19940?pg=19940

Schiano, D., Stone, M. and Bectarte, R. (2001). Search and the Subjective Web. In *Proceedings of the CHI '01 Extended Abstracts on Human Factors in Computing Systems*, 165-166.

schraefel, m. c., Zhu, Y., Modjeska, D., Wigdor, D. and Zhao, S. (2002). Hunter Gatherer: Interaction Support for the Creation and Management of within-Web-Page Collection. In *Proceedings of the 11th International Conference on World Wide Web (WWW 2002)*, Honolulu, HA, 172-181.

Sellen, A. J., Murphy, R. and Shaw, K. L. (2002). How Knowledge Workers Use the Web. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2002)*, Minneapolis, MN, 227-234.

Seo, Y.-W. and Zhang, B.-T. (2000). Learning User's Preferences by Analyzing Web-Browsing Behaviors. In *Proceedings of the 4th International Conference on Autonomous Agents*, Barcelona, Spain, 381-387.

Sparkleware (2006). SuperBot: Retrieved November 30, 2006, from http://www.sparkleware.com/superbot/index.html

Spink, A., Wolfram, D., Jansen, M. B. and Saracevic, T. (2001). Searching the Web: The Public and Their Queries. *Journal of the American Society for Information Science and Technology*, 52(3): 226-234.

Statistics Canada (2006). Canadian Internet Use Survey: Retrieved on February 27, 2006, from http://www.statcan.ca/Daily/English/060815/d060815b.htm

Tauscher, L. and Greenberg, S. (1997). How People Revisit Web Pages: Empirical Findings and Implications for the Design of History Systems. *International Journal of Human-Computer Studies*, 47(1): 97-137.

Teevan, J., Alvarado, C., Ackerman, M. S. and Karger, D. R. (2004). The Perfect Search Engine Is Not Enough: A Study of Orienteering Behavior in Directed Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2004)*, Vienna, Austria, 415-422.

Turnbull, D. (1998). WebTracker: A Tool for Understanding Web Use: Retrieved June 1, 2005, from http://www.ischool.utexas.edu/~donturn/research/webtracker/

van Dantzich, M., Robbins, D., Horvitz, E. and Czerwinski, M. (2002). Scope: Providing Awareness of Multiple Notifications at a Glance. In *Proceedings of the ACM Conference on Advanced Visual Interfaces (AVI 2002)*, Trento, Italy.

Weinreich, H., Obendorf, H., Herder, E. and Mayer, M. (2006). Off the Beaten Tracks: Exploring Three Aspects of Web Navigation. In *Proceedings of the 15th International Conference on World Wide Web (WWW 2006)*, Edinburgh, Scotland, 133-142.

Whittaker, S., Terveen, L. and Nardi, B. A. (2000). Let's Stop Pushing the Envelope and Start Addressing It: A Reference Task Agenda for HCI. *Human Computer Interaction*, 15(2-3): 75-106.

Yahoo! (2007). My Yahoo!: Retrieved February 19, 2007, from http://help.yahoo.com/l/us/yahoo/my/

Zhang, D., Zambrowicz, C., Zhou, H. and Roderer, N. (2004). User Information Seeking Behavior in a Medical Web Portal Environment: A Preliminary Study. *Journal of the American Society for Information Science and Technology*, 55(8): 670-684.

Appendix A: Field Study Demographic Questionnaire

Participant ID: _____

Questionnaire – Demographic

Please answer all questions as accurately and honestly as possible

1. Gender:
☐ Male
☐ Female

2. Age Group:
☐ 19 and under
☐ 20 – 29
☐ 30 – 39
☐ 40 – 49
☐ 50 – 59
☐ 60+

3. Academic Program
☐ Undergraduate
☐ Graduate
Department: _____

4. How many hours a week do you spend on the Web:
☐ Less than 5 hours
☐ 5 – 9 hours
☐ 10 – 19 hours
☐ 20 – 29 hours
☐ 30 – 39 hours
☐ 40+ hours

5. Do you have dedicated access to a computer (i.e., primary user):
☐ Yes ☐ No

→ If **Yes**, what kind of computer is it? (check all that apply if more than 1 dedicated computer)
☐ Desktop
☐ Laptop
☐ Tablet

6. How often do you use the following web browsers?

	Never	Rarely	Sometimes	Often	Always
Internet Explorer	☐	☐	☐	☐	☐
Netscape	☐	☐	☐	☐	☐
Mozilla	☐	☐	☐	☐	☐
Firefox	☐	☐	☐	☐	☐
Opera	☐	☐	☐	☐	☐
Safari	☐	☐	☐	☐	☐
Other:	☐	☐	☐	☐	☐

7. Please indicate the tools you often use to complete each of the following tasks:

	Fact Finding	Information Gathering	Just Browsing	Transactions
General search engines (e.g., Google)	☐	☐	☐	☐
Search engine within a web site	☐	☐	☐	☐
Bookmarks	☐	☐	☐	☐
History list	☐	☐	☐	☐
Enter sites from memory	☐	☐	☐	☐
Auto-complete	☐	☐	☐	☐
Follow links that were emailed to you	☐	☐	☐	☐
Others:	☐	☐	☐	☐

8. For each of the following tasks, please estimate the percentage of your web usage that is consumed by each task:

Fact Finding	_____%
Information Gathering	_____%
Just Browsing	_____%
Transaction	_____%

Thank You!

Appendix B: Field Study Navigation Mechanism Usage Questionnaire

Participant ID: _____

Questionnaire – Usage of Web Browser Tools

Please answer all questions as accurately and honestly as possible.

1. **How often do you use *bookmarks (Favorites)* to navigate the Web?**

Never ☐ **Rarely** ☐ **Sometimes** ☐ **Often** ☐ **Always** ☐

Why? _____

_____

2. **How often do you use the *history* function to navigate the Web?**

Never ☐ **Rarely** ☐ **Sometimes** ☐ **Often** ☐ **Always** ☐

Why? _____

_____

3. **How often do you use the *auto-complete* function to navigate the Web?**

Never ☐ **Rarely** ☐ **Sometimes** ☐ **Often** ☐ **Always** ☐

Why? _____

_____

4. **How often do you use a *search toolbar* to navigate the Web? (E.g., Google or Yahoo toolbar)**

Never ☐ **Rarely** ☐ **Sometimes** ☐ **Often** ☐ **Always** ☐

Which one? _____

Why? _____

_____

5. **How often do you use the *back* and *forward buttons* to navigate the Web?**

Never	Rarely	Sometimes	Often	Always
☐	☐	☐	☐	☐

Why?_____

_____

6. **How often do you use the *back* and *forward menus* navigate the Web?**

Never	Rarely	Sometimes	Often	Always
☐	☐	☐	☐	☐

Why?_____

_____

7. **If you use any *other* browser tools that were not listed above to navigate the Web, please complete the following section:**

i.) Name of web browser tool: _____

How often do you use this tool to navigate the Web?

Never	Rarely	Sometimes	Often	Always
☐	☐	☐	☐	☐

Why?_____

_____

ii.) Name of web browser tool: _____

How often do you use this tool to navigate the Web?

Never	Rarely	Sometimes	Often	Always
☐	☐	☐	☐	☐

Why?_____

_____

Thank You!

Appendix C: Field Study Post Study Questionnaire

Participant ID: _____

Questionnaire – Post-Study

Please answer all questions as accurately and honestly as possible.

1. **For each of the following tasks, please indicate how difficult each was to identify:**

	Very Difficult	Difficult	Neutral	Easy	Very Easy
Fact Finding	☐	☐	☐	☐	☐
Information Gathering	☐	☐	☐	☐	☐
Just Browsing	☐	☐	☐	☐	☐
Transaction	☐	☐	☐	☐	☐

Comments: _____

_____

_____

2. **In general, do you think your task descriptions (typed in text) were accurate?**

Strongly Disagree	Disagree	Undecided	Agree	Strongly Agree
☐	☐	☐	☐	☐

Comments: _____

_____

_____

3. **In general, do you think your task categorizations (buttons) were accurate?**

Strongly Disagree	Disagree	Undecided	Agree	Strongly Agree
☐	☐	☐	☐	☐

Comments: _____

_____

_____

4. Which pairs of tasks were most difficult to distinguish between?

_____

_____

_____

5. What kinds of web activity did you assign as *Other*?

_____

_____

_____

6. While working on one web task, did you find you switched to other web tasks?

Never	Rarely	Sometimes	Very Often	Always
☐	☐	☐	☐	☐

Comments: _____

_____

_____

7. Did the web browser used in this study change the way you usually work on the Web?

Not at all	A little	Somewhat	A lot	A Great Deal
☐	☐	☐	☐	☐

Please explain: _____

_____

_____

Did having to document the types of tasks in which you were engaging change the way you usually work on the Web?

Not at all	A little	Somewhat	A lot	A Great Deal
☐	☐	☐	☐	☐

Please explain:_____

_____

_____

8. **Was the functionality provided by the browser used in this study different from your usual browser?**

Not at all	A little	Somewhat	A lot	A Great Deal
☐	☐	☐	☐	☐

Please explain:_____

_____

_____

9. **How often did you use the task diary to assign tasks?**

Never	Rarely	Sometimes	Very Often	Always
☐	☐	☐	☐	☐

Why?_____

_____

10. **Other Comments:**

_____

_____

_____

Thank You!

Appendix D: Semi-structured Interviews Student Demographic Questionnaire

Demographic Questionnaire **Participant ID:** _____

Please answer all questions as accurately and honestly as possible

1. Gender:

☐ Male
☐ Female

2. Age Group:

☐ 20 and under
☐ 21 – 30
☐ 31 – 40
☐ 41 – 50
☐ 51 – 60
☐ Over 60

3. Degree Program: _____

Year in program: _____

4. How comfortable are you using computers?

Very Uncomfortable	Uncomfortable	Somewhat Comfortable	Comfortable	Very Comfortable
1	2	3	4	5

5. How comfortable are you using the Internet?

Very Uncomfortable	Uncomfortable	Somewhat Comfortable	Comfortable	Very Comfortable
1	2	3	4	5

6. How often do you use computers *during the week*?

Never	Once	A few times	Daily	Several times
1	2	3	4	5

7. How often do you use computers *during the weekend*?

Never	Once	Daily	Several times
1	2	3	4

8. How often do you use the Internet *during the week*?

Never	Once	A few times	Daily	Several times
1	2	3	4	5

9. How often do you use the Internet *during the weekend*?

Never	Once	Daily	Several times
1	2	3	4

10. What is your primary web browser? (Check one)

- ☐ Firefox
- ☐ Internet Explorer
- ☐ Mozilla
- ☐ Netscape
- ☐ Opera
- ☐ Safari
- ☐ Other _____

11. What other web browsers do you use on an occasional basis? (Check all that apply)

- ☐ Firefox
- ☐ Internet Explorer
- ☐ Mozilla
- ☐ Netscape
- ☐ Opera
- ☐ Safari
- ☐ Other _____

Appendix E: Semi-structured Interviews Worker Demographic Questionnaire

Demographic Questionnaire **Participant ID:** _____

Please answer all questions as accurately and honestly as possible

1. **Gender:**

☐ Male
☐ Female

2. **Age Group:**

☐ 20 and under
☐ 21 – 30
☐ 31 – 40
☐ 41 – 50
☐ 51 – 60
☐ Over 60

3. **Job Title:** _____

4. **How comfortable are you using computers?**

Very Uncomfortable	Uncomfortable	Somewhat Comfortable	Comfortable	Very Comfortable
1	2	3	4	5

5. **How comfortable are you using the Internet?**

Very Uncomfortable	Uncomfortable	Somewhat Comfortable	Comfortable	Very Comfortable
1	2	3	4	5

6. **How often do you use computers *during the week*?**

Never	Once	A few times	Daily	Several times a day
1	2	3	4	5

7. How often do you use computers *during the weekend*?

Never	Once	Daily	Several times a weekend
1	2	3	4

8. How often do you use the Internet *during the week*?

Never	Once	A few times	Daily	Several times a day
1	2	3	4	5

9. How often do you use the Internet *during the weekend*?

Never	Once	Daily	Several times a weekend
1	2	3	4

10. What is your primary web browser? (Check one)

- ☐ Firefox
- ☐ Internet Explorer
- ☐ Mozilla
- ☐ Netscape
- ☐ Opera
- ☐ Safari
- ☐ Other _____

11. What other web browsers do you use on an occasional basis? (Check all that apply)

- ☐ Firefox
- ☐ Internet Explorer
- ☐ Mozilla
- ☐ Netscape
- ☐ Opera
- ☐ Safari
- ☐ Other _____

Appendix F: Semi-structured Interviews Interview Guide

Participant ID: _____

Date: _____

Hi, Thank you for taking the time to participate in my study.

As part of my PhD thesis work, I am exploring user behaviour on the Web. In particular, I am interested in the activity of monitoring. So monitoring is an activity that occurs when a user revisits particular web pages to look for things that have changed, such as new or updated information. Can you think of any websites that you visit for to check for updated information?

Monitoring may take place for just a set period of time, say for a couple of weeks until a particular event occurs, or it could take place indefinitely. It could also take place several times a day, once a week, or even on a monthly basis. There is no really no set time limit. So today, I am going to be asking you questions about some of the monitoring activities that you engage in.

So to start I'll be asking you some general questions about your web usage and then I will ask some more specific questions about your monitoring activities both work-related and personal. You may find that some of the questions are repetitive but it is important for me to learn about different tasks that you complete.

I'll be asking you questions from an interview guide, and I use this so that I ask everyone the same questions but we can also deviate from it during the interview. I'll also be taking notes as you answer but would also like to tape record this session. The recordings will only be used to help me transcribe parts of the interview I wasn't able to take notes fast enough for.

Do you have any questions before we begin?

1. Roughly how much total time do you think you spend on the Web each day? How much is actively interacting with the Web and how much is in the background?
2. How much of that is at work do you think? What about at home?
3. Where do you usually work at school?
4. Can you describe some of the activities you use the Web for?
5. How many computers do you use? Do you use the same computer at work as at home?

Work Task
6. Can you think of a work related monitoring task? How would you describe the task? Task: Goal: Home or Work?
7. How long have you been doing this task for? How long do you think you'll continue to do it? Or is it a permanent
8. How often do you do this task? How long does it take? Is it part of a routine? Do you often complete this task at the same time each day?

9. Can you show me on your computer the steps you take to complete this task? So, what are the steps you take to get to the web page and how do you go about find the information?
10. So what kind of information are you looking for when you come to this page? Is it something specific? Or more general?
11. Once you find the information that you are looking for, do you take any further actions? Do you make a decision? Take further action?
12. Do you do anything to save the information? So for example, print it, save it, email it, etc.?
13. Do you find it is a lot of work to find the information?
14. Can you think of anything that would make this task easier for you?

Personal Task
15. Ok, we're going to switch gears now and if you don't mind I'd like to ask you about your personal, so by that I mean non-work-related monitoring activities. Can you think of a personal monitoring task that you do? Task: Goal: Home or Work?
16. How long have you been doing this task for? How long do you think you'll continue to do it? Or is it a permanent
17. How often do you do this task? How long does it take? Is it part of a routine? Do you often complete this task at the same time each day?

18. Can you show me on your computer the steps you take to complete this task? So, what are the steps you take to get to the web page and how do you go about find the information?
19. So what kind of information are you looking for when you come to this page? Is it something specific? Or more general?
20. Once you find the information that you are looking for, do you take any further actions? Do you make a decision? Take further action?
21. Do you do anything to save the information? So for example, print it, save it, email it, etc.?
22. Do you find it is a lot of work to find the information?
23. Can you think of anything that would make this task easier for you?

Now, I just have a few more general questions. When you answer these questions, it can be from either the perspective of work-related or personal monitoring tasks.
24. How often do you think you pick up a new monitoring task? Daily Weekly Monthly Every Couple of Months Yearly Never
25. Can you describe any monitoring tasks that you might do for a short period of time? Sometimes this will be something that happens on an annual basis, so for example, the 2 weeks leading up to the super bowl, or it could be just a one off. So for example, if I'm waiting to see if school is going to be cancelled because of a snowstorm? How do you monitor this information?
26. Do you ever receive emails that send you updated information? Do you then check the webpage? Or does it save you from having to check it?
27. Do you use any other tools that help provide you with updated or new information, such as RSS feeds?

The very last thing I would like to do is introduce a couple of monitoring tools to you and get your feedback. They are all plug-ins that can be downloaded for Firefox.

28. The first is a tool called Forecast Fox. It allows you to see the current weather, as well as the forecast for the next 2 days. What do you think about it? [show short video clip]

29. Could you see yourself using a tool like this?

30. Why or Why not?

31. The second is a tool called Page Update Checker. What it does is, you provide it with websites URLS of pages for which you would like to be notified of any changes to the pages. It pops up this little icon here when there is a change in one of your pages. You just click on the page and then it loads the updated page for you. What do you think about it? [show short video clip]

32. Could you see yourself using a tool like this?

33. Why or Why not?

34. Do you have any comments or questions before we wrap up?

End Time: _____

Recorder: _____

Appendix G: Semi-structured Interviews Categorization of Monitoring Activities

During the semi-structured interviews, participants were asked to describe three personal and three school/work related monitoring activities. Table 37 shows the full list of reported monitoring activities, categorized by task type.

Table 37. All reported monitoring activities, categorized by task type.

	Browsing	Communications	Fact Finding	Info. Gathering	Maintenance	Transactions	Total
Academic Programmes	0	0	0	1	0	0	**1**
Assignments/Notes/Grades	0	0	8	0	0	0	**8**
Balance	0	0	0	0	0	11	**11**
Blogs/Wikis	9	0	0	0	0	0	**9**
Books/Movies/Music	6	0	0	1	0	0	**7**
Bulletin Board	5	1	0	2	1	0	**9**
Comics	2	0	0	0	0	0	**2**
Commodities	0	0	1	2	0	0	**3**
Conference Details	0	0	4	0	0	0	**4**
Course Offerings	0	0	2	1	0	0	**3**
Department Page	4	0	0	0	3	0	**7**
Email	0	10	0	0	0	0	**10**
Employment	1	0	0	2	0	0	**3**
Files	0	0	2	0	0	1	**3**
Horoscope	1	0	0	0	0	0	**1**
Housing	0	0	0	1	0	0	**1**
News	22	0	0	0	0	0	**22**
Online Community	3	0	0	0	0	0	**3**
Page Maintenance	0	0	0	0	1	0	**1**
Project Plans	0	0	0	0	1	0	**1**
Purchase	4	0	3	0	0	2	**11**
Research Literature	3	0	0	7	0	0	**10**
Schedules	0	0	2	0	0	0	**2**
Sports	8	0	5	1	0	0	**16**
Status	0	0	5	0	1	1	**8**
Ticket System	4	0	0	0	0	0	**4**
Weather	0	0	7	0	0	0	**7**
Web Cam	1	0	0	0	0	0	**1**
Web Stats	1	0	3	0	0	1	**5**
TOTAL	**74**	**11**	**42**	**18**	**7**	**16**	**173**

Appendix H: Laboratory Study Coding Sheet

Participant ID:
Part I: Bookmarks
Part II: Primary and Secondary **CHI 1:** **eBay 1:**
Part III: Enhanced Bookmarks
Part IV: **DAL:** **CHI2:** **CHI3:**

Appendix I: Laboratory Study Demographic Questionnaire

Participant ID: _____

Please answer all questions as accurately and honestly as possible

1. Sex:

☐ Male
☐ Female

2. Age Group:

☐ 20 and under
☐ 21 – 30
☐ 31 – 40
☐ 41 – 50
☐ 51 – 60
☐ Over 60

3. Programme:

☐ Undergraduate
☐ Graduate

Degree: _____

Year: _____

4. How many hours a week do you spend on the Web:

☐ Less than 5 hours
☐ 5 – 9 hours
☐ 10 – 19 hours
☐ 20 – 29 hours
☐ 30 – 39 hours
☐ Over 40 hours

5. What is your preferred web browser?

☐ Internet Explorer
☐ Netscape
☐ Mozilla
☐ Firefox
☐ Opera
☐ Safari
☐ Other: _____

Other Side →

6. Do you currently use any tools that help you monitor information on the Web?

☐ Yes
☐ No

If Yes, please describe the tools you use (e.g., RSS, weather widgets, dashboard, etc.):

Appendix J: Laboratory Study Post-study Questionnaires

Participant ID: _____

Please answer the following questions with respect to this tool:

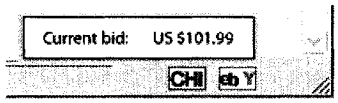

	1=Disagree						7=Agree
The tool effectively supports rapid reaction to important information:	1	2	3	4	5	6	7
Appropriate reactions to the tool are obvious and intuitive:	1	2	3	4	5	6	7
The tool instilled confidence that all important items were seen:	1	2	3	4	5	6	7
The tool provides an overall sense of the information:	1	2	3	4	5	6	7
The tool supports easy understanding of how information changes over time:	1	2	3	4	5	6	7
The tool provides these overall understandings of information in a non-intrusive manner:	1	2	3	4	5	6	7
Important new information can be quickly recognized and accessed and it is possible to stay aware of all information without loosing your place in the document:	1	2	3	4	5	6	7
The tool functionality was appropriate for the kind of information I was monitoring:	1	2	3	4	5	6	7
Overall, I liked the functionality provided by the monitoring tool:	1	2	3	4	5	6	7

	1=Disagree						7=Agree
Overall, I liked the design of the user interface (i.e., look and feel):	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Overall, the tool was easy to use:	**1**	**2**	**3**	**4**	**5**	**6**	**7**
I would like to use a similar tool within my own web usage:	**1**	**2**	**3**	**4**	**5**	**6**	**7**

Comments:

Participant ID: _____

Please answer the following questions with respect to this tool:

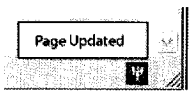

	1=Disagree						7=Agree
The tool effectively supports rapid reaction to important information:	1	2	3	4	5	6	7
Appropriate reactions to the tool are obvious and intuitive:	1	2	3	4	5	6	7
The tool instilled confidence that all important items were seen:	1	2	3	4	5	6	7
The tool provides an overall sense of the information:	1	2	3	4	5	6	7
The tool supports easy understanding of how information changes over time:	1	2	3	4	5	6	7
The tool provides these overall understandings of information in a non-intrusive manner:	1	2	3	4	5	6	7
Important new information can be quickly recognized and accessed and it is possible to stay aware of all information without loosing your place in the document:	1	2	3	4	5	6	7
The tool functionality was appropriate for the kind of information I was monitoring:	1	2	3	4	5	6	7
Overall, I liked the functionality provided by the monitoring tool:	1	2	3	4	5	6	7

	1=Disagree						7=Agree
Overall, I liked the design of the user interface (i.e., look and feel):	**1**	**2**	**3**	**4**	**5**	**6**	**7**
Overall, the tool was easy to use:	**1**	**2**	**3**	**4**	**5**	**6**	**7**
I would like to use a similar tool within my own web usage:	**1**	**2**	**3**	**4**	**5**	**6**	**7**

Comments:

Participant ID: _____

Please answer the following questions with respect to this tool:

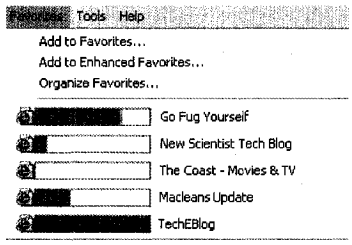

	1=Disagree						7=Agree
The tool effectively supports rapid reaction to important information:	1	2	3	4	5	6	7
Appropriate reactions to the tool are obvious and intuitive:	1	2	3	4	5	6	7
The tool instilled confidence that all important items were seen:	1	2	3	4	5	6	7
The tool provides an overall sense of the information:	1	2	3	4	5	6	7
The tool supports easy understanding of how information changes over time:	1	2	3	4	5	6	7
The tool provides these overall understandings of information in a non-intrusive manner:	1	2	3	4	5	6	7
The tool functionality was appropriate for the kind of information I was monitoring:	1	2	3	4	5	6	7
Overall, I liked the functionality provided by the monitoring tool:	1	2	3	4	5	6	7
Overall, I liked the design of the user interface (i.e., look and feel):	1	2	3	4	5	6	7

	1=Disagree						7=Agree
Overall, the interface was easy to use:	**1**	**2**	**3**	**4**	**5**	**6**	**7**
I would like to use a similar tool within my own web usage:	**1**	**2**	**3**	**4**	**5**	**6**	**7**

Comments:

Appendix K: Laboratory Study Interview Guide

Participant ID: _____

Interview Guide

PART I:

1. **So what was your final position on the Student Volunteer List?**

   Did you make it on to the accepted list?

2. **Were there any new bids on your guitar hero auction?**

   How much was your eBay auction at?

   Did you make your money back?

3. **Was the Tutorial page updated?**

   Do your TAs need to make any edits?

   How many?

Part II:

4. **When using the regular bookmarks, how did you choose which site to browse first?**

5. **When using the enhanced bookmarks, how did you choose which site to browse first?**

6. **For each of the tools, was there anything in particular that you liked or did not like about the interface or functionality?**

7. **Do you use any similar tools in real life? If so, what for?**

8. Could you envision using these tools in your own web usage? If so, what for?

9. For the two notification tools, would you have preferred to hover on the icon to see the updated information? Or do you prefer clicking on it?

10. Which tool did you like best?

So if I asked you to rank the tools in order of preference, what would your ranking be? (ties are allowed)

FF _____

MA _____

BR _____

11. For the student volunteer list, if given the choice, which tool would you use? Why?

FF MA BR None

12. For the eBay auction, if given the choice, which tool would you use? Why?

FF MA BR None

13. For the tutorial page, if given the choice, which tool would you use? Why?

FF MA BR None

14. For the pages you were just browsing, if given the choice, which tool would you use? Why?

FF MA BR None

15. Have you ever visited any the web pages used in this study before? How Often

- ☐ eBay
- ☐ CHI SV
- ☐ Go Fug Yourself
- ☐ TechEBlog
- ☐ New Scientist Tech Blog
- ☐ The Coast
- ☐ MacLeans blog

16. Any other comments?

Appendix L: Predicting Task

Using Web Browser Interactions to Predict Task

Melanie Kellar and Carolyn Watters
Faculty of Computer Science, Dalhousie University
Halifax, Nova Scotia, Canada
{kellar, watters}@cs.dal.ca

ABSTRACT
The automatic identification of a user's task has the potential to improve information filtering systems that rely on implicit measures of interest and whose effectiveness may be dependant upon the task at hand. Knowledge of a user's current task type would allow information filtering systems to apply the most useful measures of user interest. We recently conducted a field study in which we logged all participants' interactions with their web browsers and asked participants to categorize their web usage according to a high-level task schema. Using the data collected during this study, we have conducted a preliminary exploration of the usefulness of logged web browser interactions to predict users' tasks. The results of this initial analysis suggest that individual models of users' web browser interactions may be useful in predicting task type.

Categories and Subject Descriptors
I.5.1 **[Pattern Recognition]**: Models - *Structural*. H.5.4. **[Information Interfaces and Presentation]**: Hypertext/ Hypermedia - *Navigation*.

General Terms: Algorithms, Experimentation, Human Factors.

Keywords: Web, task, field study, information filtering, task prediction, decision tree.

1. INTRODUCTION
The automatic categorization of a user's web activity could serve useful in developing more intuitive and effective information filtering systems. Many information filtering systems rely on implicit measures to infer user interest, such as dwell time, mouse and keyboard activity, and document interactions (e.g., copy, cut, save, print) [7]. However, we do not yet have a solid understanding of the role task plays in the effectiveness of implicit measures for information filtering. While the evidence is somewhat incongruous, there is evidence to suggest that the effectiveness of an implicit measure may be dependant upon the task at hand. For example, previous research has found through laboratory experiments that the usefulness of dwell time was influenced by the type of task [5; 2]. However, a later field study found that task did not significantly impact the usefulness of dwell time as an implicit measure of interest [4]. We believe that once it becomes clear which implicit measures are most effective for a task type, the ability to automatically infer a user's task would allow us to apply the most appropriate measures of user interest, thereby improving the effectiveness of information filtering systems.

WWW 2006, May 23–26, 2006, Edinburgh, Scotland.
ACM 1-59593-323-9/06/0005.

While previous research has examined different aspects of automatic task prediction, such as the identification of goals [6], active tasks [8], and focused task types [1], we are interested in identifying task types from a high-level perspective across all user tasks. In a recently conducted field study [3], we observed that participants displayed significant differences in how they interacted with their web browser across different information seeking tasks and online transactions. We are looking to capitalize on these differences by automatically categorizing users' tasks according to a high-level schema. In this paper, we report on preliminary work using traces of users' web browser interactions, collected in a realistic setting, to automatically identify high-level tasks.

2. METHODOLOGY
The data was collected during a one week field study conducted with 21 participants. The goal of the study was to gain an understanding of the characteristics of information seeking tasks on the Web and how the features of current web browsers are being used to complete these tasks. Over the course of the study, participants were asked to use a custom web browser that logged all of their interactions with the web browser, including URLs visited, navigation mechanisms used (e.g., back button, bookmarks, clicked links, auto-complete), and use of browser functions (e.g., copy, paste, find).

Participants were also asked to provide a categorization and brief textual description for all their web usage over the course of the study. This could be done either in real time using a task toolbar embedded in the custom web browser or at the end of each day using a task diary. Participants categorized their web usage according to the following schema: *Fact Finding* (looking for specific information, such as an email address or a sports score), *Information Gathering* (collecting information, often from multiple sources to make a decision, write a report, etc.), *Browsing* (serendipitous web navigation for the purpose of entertainment or to see what is new), and *Transactions* (online actions such as email or banking). We also included the category Other for tasks which did not fit within the given categories.

3. DATA
Overall, participants recorded 1192 tasks involving 13,498 web pages over the week long study. Each task was associated with a number of data features logged by the web browser during the field study. The features included in our model are: task duration (calculated by summing the web page dwell time across the task), number of pages loaded, number of windows opened, type of web browser navigation mechanism used to initiate the task (e.g., bookmarks, clicked link, typed-in URL), use of browser functions (e.g., cut, copy, paste, save), and use of Google.

4. RESULTS

4.1 Task Prediction

We observed strong differences in how participants interacted with their web browsers during different tasks and were interested in whether logged interactions could be used to predict the category of task, from a high-level perspective. Using decision trees, we constructed several preliminary models using the Weka C4.5 decision tree package. All models were trained using 10-fold cross validation. Only a small number of tasks were labelled as Other (1.7% - 20/1192); therefore, we did not include this data in the task prediction.

An initial decision tree was built using 1172 tasks from all 21 participants. This aggregate model correctly classified 53.8% of the task instances. The results of this classification suggest that an aggregate model is not appropriate and we suspect it is because user behaviour on the Web is simply too individual. For instance, we observed that the use of web browser navigation mechanisms was highly dependant upon the individual [3]. We also observed that there was a high degree of variability across the participants for measures such as dwell time, number of pages viewed, and number of windows loaded.

In response to these results, we constructed individual models for each participant who had logged more than 30 tasks (17/21 participants). The percentage of positive classification is shown in Figure 1 for each participant. The number of correctly classified tasks by participant ranged from a high of 94.3% to a low of 43.6%. We expect that more accurate models could be produced with a larger data set, as well as more sophisticated feature selection and machine learning techniques.

4.2 Task Initiation

There are several aspects of a user's task that an automatic task categorization system must detect, such as task initiation, task type, task switching, and task termination. We observed a strong identifier of new task initiation through the use of a particular set of web browser navigation mechanisms. This consists of New Task Session (NTS) navigation mechanisms [3], which were typically employed by participants either when initiating a new task session or when changing navigation strategies within a session. The NTS navigation mechanisms include the auto-complete function, bookmarks, the Google toolbar, the home button, selecting a URL from the drop-down address bar, and typed-in URLs. We observed 1314 instances of NTS navigation mechanisms across all participants and tasks. Of those instances, 70.3% (924/1314) occurred when participants were initiating a new task. This suggests that the use of the web browser navigation mechanisms may play a useful role in detecting the beginning of new task sessions.

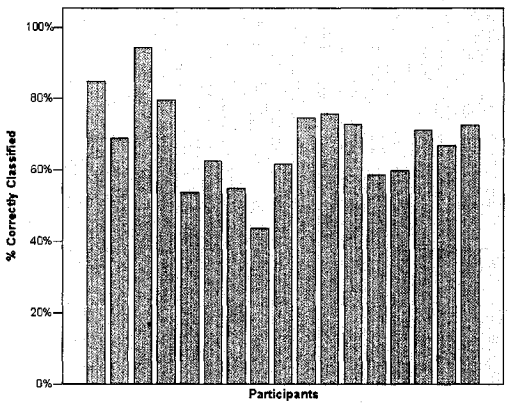

Figure 1. The percentage of correctly classified tasks, by participant

5. CONCLUSIONS

We have presented initial work in automatically classifying web-based tasks based on user interactions with their web browser. The data used in our predictive models represents a realistic picture of how users interact on the Web and was collected over the course of a week long field study. Moreover, the task categorization represents all high-level tasks in which users engage on the Web, not simply a subset of randomly chosen tasks. The results of this work suggest that more accurate task prediction is possible when individual models are used, in comparison with aggregate models of user behaviour.

As part of our future work, we would like to further explore machine learning techniques to improve the accuracy of the task predictions. As well, we would like to incorporate a categorization of the web pages visited, potentially according to web page genre.

6. ACKNOWLEDGMENTS

This research was funded by the Natural Science and Engineering Research Council of Canada (NSERC).

7. REFERENCES

[1] Iqbal, S. T. and Bailey, B. P. (2004). Using Eye Gaze Patterns to Identify User Tasks. In Proceedings of the Grace Hopper Celebration of Women in Computing.

[2] Kellar, M., Watters, C., Duffy, J. and Shepherd, M. (2004). Effect of Task on Time Spent Reading as an Implicit Measure of Interest. In Proceedings of ASIS&T 2004 Annual Meeting, Providence, RI, 168-175.

[3] Kellar, M., Watters, C. and Shepherd, M. (2006). The Impact of Task on the Usage of Web Browser Navigation Mechanisms. In Proceedings of Graphics Interface 2006, Quebec City, Canada, To Appear.

[4] Kelly, D. and Belkin, N. (2004). Display Time as Implicit Feedback: Understanding Task Effects. In Proceedings of SIGIR 2004, Sheffield, UK, 377-384.

[5] Kelly, D. and Belkin, N. J. (2001). Reading Time, Scrolling and Interaction: Exploring Implicit Sources of User Preference for Relevance Feedback. In Proceedings of SIGIR 2001, New Orleans, LA, 408-409.

[6] Lee, U., Liu, Z. and Cho, J. (2005). Automatic Identification of User Goals in Web Search. In Proceedings of WWW 2005, Chiba, Japan, 391-400.

[7] Oard, D. W. and Kim, J. (2001). Modeling Information Content Using Observable Behavior. In Proceedings of ASIS&T 2001 Annual Meeting, Washington, DC.

[8] Oliver, N., Smith, G., Thakkar, C. and Surendran, A. C. (2006). Swish: Semantic Analysis of Window Titles and Switching History. In Proceedings of IUI '06, Sydney, Australia, 194-201.