

A Field Study Characterizing Web-based Information Seeking Tasks

Melanie Kellar
Carolyn Watters
Michael Shepherd

Technical Report CS-2005-20

January 15, 2005

Faculty of Computer Science
6050 University Ave., Halifax, Nova Scotia, B3H 1W5, Canada

A Field Study Characterizing Web-based Information Seeking Tasks

Melanie Kellar, Carolyn Watters, and Michael Shepherd

Faculty of Computer Science, Dalhousie University, Halifax, NS, B3H 1W5, Canada

Email: {melanie, watters, shepherd@cs.dal.ca}

ABSTRACT

Previous studies have examined various aspects of user behaviour on the Web, including general information seeking patterns, search engine use, and revisitation habits. Little research has been conducted to study how users navigate and interact with their Web browser across different information seeking tasks. We have conducted a field study of 21 participants in which we logged detailed Web usage and asked participants to provide task categorizations of their Web usage based on the following categories: Fact Finding, Information Gathering, Browsing, and Transactions. We used implicit measures logged during each task session to provide usage measures, such as dwell time, number of pages viewed, and the use of specific browser navigation mechanisms. We also report on differences in how participants interacted with their Web browser across the range of information seeking tasks. Within each type of task we found several distinguishing characteristics. In particular, Information Gathering tasks were the most complex; participants spent more time completing this task, viewed more pages, and used the Web browser functions most heavily during this task. The results of this analysis have been used to provide implications for future support of information seeking on the Web, as well as direction for future research in this area.

INTRODUCTION

Previous research has contributed to a general understanding of the types of information seeking tasks in which users engage on the Web (Choo, Detlor, & Turnbull, 2000; Morrison, Pirolli, & Card, 2001; Sellen, Murphy, & Shaw, 2002). However, much of this research was conducted five to ten years ago. While commercial Web browser functionality (e.g., bookmarks, back button) has changed little over this time, the Web has changed dramatically. The number of available Web pages has grown exponentially since these early studies (from millions to several billions) and the uses of the Web have expanded dramatically to include wikis, blogs, and various Web-based applications (e.g., online photo sharing, banking, e-learning). Google has changed the way users search on the Web and the popularity of open source browsers, such as Firefox, has contributed to a wide variety of Web browser toolbars and plug-ins.

There is a large body of research examining how users navigate the Web (Catledge & Pitkow, 1995; Tauscher & Greenberg, 1997; Weinreich, Obendorf, Herder, & Mayer, 2006) as well as mechanisms for Web navigation (MacKay, Kellar, & Watters, 2005; Milic-Frayling, Sommerer, & Rodden, 2003; Moyle & Cockburn, 2003). However, these studies are typically conducted in the field without any understanding of the types of tasks undertaken by the user or in laboratory setting for a focused set of tasks. For instance, user revisitation patterns on the Web have been studied extensively (Cockburn & McKenzie, 2001; Herder, 2005; Tauscher & Greenberg, 1997) in a field setting, but without any understanding of how revisitation may be impacted by the underlying task type. How a user completes a revisitation task may be influenced by the user's task and intentions. For example, a user who revisits a Web page in order to re-find a previously found fact may need different navigational support than a user who is revisiting a Web page in order to monitor new information.

One gap that we have identified in the literature is the lack of research examining how users interact with their Web browsers within the context of task. Therefore, the goals of this research are threefold. First, we aim to gain an understanding of the types of Web browser functionalities (e.g., number of pages viewed, windows opened, use of Web browser navigation mechanisms) currently being used during information seeking tasks. Second, we aim to determine whether there are differences in the use of these functionalities across the different information seeking tasks. Third, we aim to provide an updated view of the types of tasks being performed on the Web. It is also important that this data is collected in as natural of a user environment as possible. An understanding of how users interact with their Web browsers across different information seeking tasks has the potential to provide valuable insight into how future tools should be designed.

This paper reports on a recent field study in which we studied differences in how users interact with their Web browser during the following information seeking tasks: Fact Finding, Information Gathering, Browsing, and Transactions. Participants were asked to annotate their Web usage with task information while using a custom Web browser that logged their interactions with the browser. The key contribution of this paper is 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. In the next section, we present an overview of the related literature followed by a description of the methodological and data collection techniques used during the field study. We then report general observations describing the characteristics of participants' task sessions as well as differences between the tasks according to 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 next provide a summary of the findings and discuss the implications of our results. Finally, we conclude with future directions for this research.

RELATED WORK

In this section we discuss related work exploring the relationship of task to user behaviour on the Web and methodologies for studying the information seeking behaviour of Web users.

Task and User Behaviour on the Web

There is a large body of theoretical research examining information seeking in both electronic and non-electronic environments (Belkin, 1980; Ellis, 1989; Kuhlthau, 1991; Marchionini, 1995; Wilson, 1997). However, 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. Cothey (2002) noted 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".

Several studies have examined general user behaviour on the Web. In one of the first studies of Web usage, Catledge and Pitkow (1995) classified usage strategies into three categories: serendipitous, general purpose, and searcher. Pitkow and Kehoe (1996) reported five main uses of the Web from the fourth Gvu WWW survey¹: browsing, entertainment, work, shopping, and other uses. They also noted that the activities had remained fairly consistent since the second study.

¹ http://www.cc.gatech.edu/gvu/user_surveys/

Only a few in-depth studies, however, 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 et al. (2000). They studied critical incidents 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 characterized as undirected viewing, conditioned viewing, informal search, and formal search. Morrison et al. (2001) studied significant Web actions through 2188 responses to the tenth GVU WWW user survey. 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 et al. (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 strong 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 enough 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).

While our research complements the previous research presented in Table 1, it differs several ways. First, we are providing an updated view of users' information seeking behaviour on the Web. The previous studies were published between 2000 and 2002 and the Web has changed substantially since this time. Second, we are incorporating a number of elements from each of the previous studies in this single study, including: a week-long observation of all participant behaviour in a natural environment, logging of fine-grain Web browser interactions, and detailed task information. Finally, we have conducted an in-depth analysis of the differences in Web browser interactions across information seeking task sessions.

Task Session

In the study of user behaviour on the Web, a session is generally defined as a period of continuous Web usage. However, the specific definition of a session tends to vary across researchers and research disciplines. For instance, based on their client-side transaction logs, Catledge and Pitkow (1995) defined a session as period of continuous Web usage with no break in usage greater than 25.5 minutes. In their studies of search engine transaction logs, Jansen and Spink (2003) measured session duration from the time the first query was submitted to the search engine until the user quit the search engine. They reported that 52% of all sessions lasted less than 15 minutes. Montgomery and Faloutsos (2001) defined a session as a period of continuous Web usage, beginning when the user has not accessed the Web in the previous two hours. Grace-Martin and Gay (2001) used a ten minute of inactivity cut-off while Hawkey and Inkpen (2005b) used both a 10 and 30 minute cut-off. In this work, we use the term *task session* to represent a period of continuous Web usage, annotated with the same task information, with no break in usage greater than 25.5 minutes.

Information Seeking 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 specific user Web search behaviour 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. They reported that only 35% of all queries appeared to be of the type traditionally supported by search engines (e.g., directed and undirected search, advice seeking) while over 40% of the queries were non-informational, such as resource-based queries looking for products and services. This suggests that the primary focus of many commercial search engines may be misguided. Lee, Liu, and Cho (2005) have further extended this work to automatically classify Web search goals and were able to correctly categorize 90% of the search goals 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. (2005) 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".

Previous research has also examined the strategies users employ to conduct Web searches. Fidel & Efthimiadis (1999) studied user information seeking behaviour through interviews and direct observations. They reported that although there were common search strategies among the participants (e.g., search queries, rapidly scanned results), they also observed that individuals had developed their own personal search strategies. 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 Sracevic (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.

The field study described in this paper was conducted with experienced Web users and it is important to note that the information seeking strategies of users may be influenced by their level of experience. Aula, Jhaveri, and Kāki (2005) and Aula and Kāki (2003) studied the Web search strategies of expert users in order to better understand how to make the strategies of experienced users available to novice users. In a ten month longitudinal study Cothey (2002) examined the change in students' information seeking behaviour as they gained more experience over time. 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. Cooper (2001) found that session length for use of an electronic library catalogue increased over a sixteen month study period while the number of searches per session remained constant.

Much of the research examining users' search strategies on the Web has been conducted in the workplace. Rieh (2003) conducted one of the first studies examining Web searching behaviour in the home and found that users searched differently in the home than in previous research examining search strategies in the workplace. In that study, participants searched the Web more frequently, but for shorter periods of time and the types of searches conducted were much broader. During our field study we recruited laptop users in order to capture participants' Web usage *both* at home and at school/work; however, this study did not focus on differences in use between home and school/work but instead on capturing a comprehensive picture of web use.

Strategies for Studying Information Seeking Behaviour on the Web

A wide variety of methodologies have been employed to study information seeking behaviour on the Web and the three main strategies include laboratory experiments, sample surveys, and field studies. The choice of research strategy influences 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, 1995). 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. The interested reader can refer to Martzoukou (2005) for a more detailed review of information seeking research strategies. Within this research domain, we are beginning to see a wider variety of strategies employed, even within single studies.

Laboratory experiments have been used in a number of studies (Hölscher & Strube, 2000; Jhaveri & Rähä, 2005) and afford researchers a greater level of control. 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.

One major drawback of laboratory experiments is that they do not offer much 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, toolbars). Often the tasks are contrived, and task success can be influenced by several factors; including cognitive and problem-solving abilities and styles (Kim & Allen, 2002), domain knowledge (Hölscher & Strube, 2000), and Web experience (Cothey, 2002; Hölscher & Strube, 2000). One alternative, as used by Schiano, Stone, and Bectarte (2001), is to invite participants to perform a task they already needed to do.

Sample Surveys are often used (Aula et al., 2005; Morrison et al., 2001; Pitkow & Kehoe, 1996) because they can be administered to large and diverse population and can produce data with a high degree of generalizability. Unlike other research strategies, the method of data collection is relatively uncomplicated.

The downside of sample surveys 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.

Field studies are becoming an increasingly common research strategy for studying user behaviour on the Web (Choo et al., 2000; Sellen et al., 2002; Teevan et al., 2004). The primary strength of field studies is the increase in realism as participants are observed working in their own environment(s), with their own tools (e.g., bookmarks, history, choice of browser), and completing tasks that are motivated by the participant and not the researcher.

In general, field studies are conducted with a relatively small, homogenous set of participants, which can lessen the generalizability of results. Due to the natural environment in which field studies are conducted, data collection can be difficult and researchers must often accept a loss of precision and control. Researchers have employed a variety of method to observe their participants, such as transaction logs (client/proxy/server side), video/screen capture, and participant diaries.

In this research, we chose to conduct a field study with the aim of maximizing the realism of the research setting. This methodology does come at the expense of both generalizability and precision, but allowed us to study the information seeking habits of our participants in their own environments (home/work/school), engaging in their own tasks, and with access to their usual navigation mechanisms (e.g., bookmarks, history, auto-complete). While participants were asked to use a custom built Web browser, it mimicked Microsoft Internet Explorer (IE) as closely as possible and included all of their IE settings and usage history in order to try and minimize the effects of the study (and the logging tool) on their behaviour.

Implicit Measures

Implicit measures (Kelly & Teevan, 2003) consist of the collection of user behaviour traces that can be recorded without any intervention on the part of the user. 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, 2001; Morita & Shinoda, 1994; Oard & Kim, 2001). In our research, we are not using implicit measures to infer interest but rather we are using them to define task characteristics.

While implicit measures may be used on a per page basis, we have examined the implicit measures recorded over the course of a task session. The logged measures include dwell time, number of windows opened, number of pages loaded, the use of Web browser navigation mechanisms, time of day, the use of Google, the use of site specific searches, and Web browser functions. Previous researchers (Mat-Hassan & Levene, 2005; Seo & Zhang, 2000) have used implicit measures to explore information seeking behaviour on the Web; however, we are studying users' information seeking behavior over a wide range of information needs rather than within a single portal or dataset.

RESEARCH QUESTIONS

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

R1. What Web browser functionalities are currently being used during Web-based information seeking tasks? In particular, we are 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 Transactions.

R2. Are there differences in the patterns of use of Web browser functionality across Web-based information seeking tasks? We are interested in whether there are significant differences in how participants interact with their Web browser between Fact Finding, Information Gathering, Browsing, and Transaction.

The primary goals of this work are to gain a better understanding of (1) the types of Web browser functionalities (e.g., number of pages viewed, windows opened, use of Web browser navigation mechanisms) currently being used during information seeking tasks, (2) the differences in the use of these functionalities across the different information seeking tasks, and (3) the types of tasks users are currently performing on the Web. This new understanding may provide insight into how to better support Web-based information seeking. Therefore, we expect the results of this research can be used to guide the design of more effective tools to better support users in their information seeking tasks on the Web.

METHODOLOGY

In order to address our primary research questions (R1 and R2), two forms of data collection were needed. First, participants' Web usage and Web browser interactions were logged over the course of the field study using a custom-built Web browser. Second, participants were asked to use 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 first describes pilot study and focus group, followed by a description of the methodology and data collection strategies used in conducting the field study.

Pilot Study

A four day 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 five categories: Fact Finding, Information Gathering, Monitoring, Browsing, and Other. During a fifteen minute training session before the pilot, participants 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 study software (Web browser and electronic diaries) 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 et al., 2000; Morrison et al., 2001; Sellen et al., 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. Participants struggled with the task of Monitoring because it often lead 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 customized Web browser and logging tools as well as refine the training materials provided to participants.

Focus Group for Task Refinement

Ten participants from the Faculty of Computer Science at Dalhousie University (students and faculty) took part in an informal focus group, none of whom had taken part in the pilot study. We selected forty 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*, *Updating my blog*. The participants were asked to work together as a group to organize the 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.

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

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
<i>Looking for Specific Information</i>	Location of a conference workshop Finding percentage of the population that is left handed
<i>Passing Time/ Entertainment</i>	Random Surfing Just browsing EBay
<i>Transactions/ Communication</i>	Checking my email Online banking
<i>Information Gathering</i>	Trying to find a reviewer to review a conference paper Looking for references on a topic
<i>Routine/Hobby</i>	Reading my favourite comic Reading blogs
<i>Monitoring</i>	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 is an activity within information seeking tasks rather than an independent task onto itself. Therefore, we eliminated the monitoring category. We also 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 as they are both serendipitous in nature and lack specific goals. The resulting task categories, shown in Figure 2, are: 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.

FIG. 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 printer drivers 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 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 including 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 either 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.

Field Study

Twenty-one university 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.

Emailed recruitment notices were circulated 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. Also, 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 (CS) students are typically considered to be highly technical, all participants were experienced Web users. 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.

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 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 on a backup disk. Participants completed a final post-study questionnaire and were paid \$25 for their participation in the study.

Before we began our 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.

Data Collection

Over the course of the field study, three types of participant data were collected: qualitative task data, implicit measures data, and questionnaire data.

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, shown in Figure 3a, 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 in real-time using the task toolbar shown in Figure 3b, at the end of the day using the task diary shown in Figure 3c, or using a combination of both techniques.

FIG. 3. The custom web browser (a) was built to mimic IE and provided a task toolbar (b) for participants to record their task information in real-time. Participations could also use the task diary (c) to record their task information at the end of each day.

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.

TABLE 3. Task information log file.

Window_ID	Date_Time	Page_Title	URL	Task_Type	Task_Description
26229	8/14/2005 18:03:48.018	MSN.com	http://www.msn.com	Other	Homepage
26229	8/14/2005 18:04:12.273	Local Weather	http://www.weather.com/<snip>	Fact Finding	Weather
26229	8/14/2005 18:04:40:965	CNN.com	http://www.cnn.com	Just Browsing	News
26229	8/14/2005 18:05:41.812	Travelocity	http://www.travelocity.com	Info Gathering	Plan trip
26229	8/14/2005 18:05:48.572	Guides & Advice	http://www.dest.travelocity.com	Info Gathering	Plan trip

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. 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. The task diary also allowed all 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 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 3, in the following format: Window ID, Date & Time, Page Title, URL, Task Type, and Task Description.

Implicit Measures

The implicit measures data was collected by the custom Web browser, shown in Figure 3a, which was built to mimic IE and provided the same functionality, menus, and interface as IE. In our research we were primarily interested in implicit measures consisting of participants’ direct interactions with the web browser interface. Table 4 displays a classification of the implicit measures logged by the custom Web browser. This classification was partially based on Byrne, John, Wehrle, and Crow’s (1999) taskonomy of Web tasks, which was developed to better understand the range of tasks that Web browsers must support. Oard and Kim (2001) developed a classification of observable behaviours based on two dimensions: behaviour category (i.e., purpose of the category) and minimum scope of the object being manipulated. While this classification was developed in the context of implicit measures, it has a larger focus on documents and content, as opposed to Web browser interface interactions. Oard and Kim’s (2001) classification has also been further extended by Kelly and Teevan (2003) and Jansen and McNeese (2005).

TABLE 4. A classification of the implicit measures logged during the field study.

Document Complete Events†	Browser Function Events			
	File	Edit	View	Misc. Tools
Auto-Complete	New	Select All	Toggle Favorites	Highlight search terms
Back Button	Window	Find	Toggle history	Internet Options
Back Menu	Open	Copy‡	Stop	
Favorites	Save As	Paste‡	View Source	
Forward Button	Page Setup	Cut‡	Privacy Report	
Forward Menu	Print			
Google Toolbar	Print Preview			
History	Properties			
Home Button	Close			
Hyperlinks				
New Window				
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)				

The custom Web browser generated that provided a detailed summary of all user interactions within the browser during each Web session. Two main types of implicit measures were recorded: document complete events (i.e., use of Web browser navigation mechanisms) and browser function events. Document complete events were recorded as each Web page was loaded. A listing of all document complete events is shown in the leftmost column of Table 4 and included navigation mechanisms such as the back button, auto-complete, bookmarks, and history. Browser function 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.

For all logged implicit measure, the data collected consisted of the Window ID, Date & Time, Browser Event (document_complete or browser_function), and Description (of the event). Participants did not have access to this log file; therefore, the URLs of the pages loaded were omitted in case the corresponding URL was deleted using the task diary. Before analysis, timestamps were used to merge the two log files shown in Table 3 and Table 5.

TABLE 5. Event log file.

Window_ID	Date_Time	Event	Description
26229	8/14/05 18:03:41.139	Browser_Function	Open_Session
26229	8/14/05 18:03:48.039	Document_Complete	First_Page
26229	8/14/05 18:04:12.404	Document_Complete	Bookmarks_Toolbar
26229	8/14/05 18:04:40.985	Document_Complete	Enter_Key
26229	8/14/05 18:05:25.659	Browser_Function	Paste_Combo
26229	8/14/05 18:05:41.832	Document_Complete	Go_Button
26229	8/14/05 18:05:48.582	Document_Complete	Clicked_Link
26229	8/14/05 18:05:52,337	Browser_Function	Close_Session

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.

Questionnaires

Participants completed three separate questionnaires over the course of the study. During the pre-study session, a demographic questionnaire was used to collect participants' demographic information and current Web usage. An inventory questionnaire of the Web browser navigation mechanisms used was also completed by participants during the pre-study session. Upon completion of the study, participants completed a post-study questionnaire which examined any difficulties they encountered during the study.

RESULTS

General Observations

In this section, we report general observations describing the characteristics of participants' task sessions. As previous 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. 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). 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 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. The breakdown of all task sessions across all participants is shown in Figure 4.

TABLE 6. The breakdown of tasks by participant.

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

Note. FF = Fact Finding; IG = Information Gathering; BR = Browsing; TR = Transactions; OT = Other.

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.

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.

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.

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.

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 maybe have been specific to our user population (i.e., mainly Computer Science students).

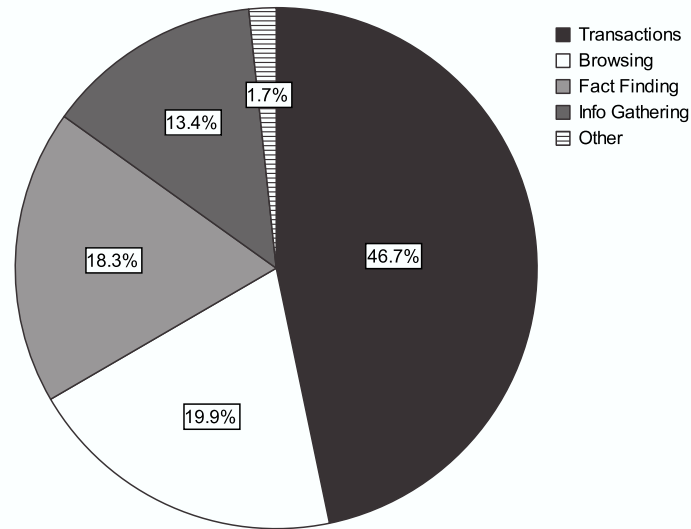


FIG 4. Breakdown of all tasks captured.

Repeated Tasks

We 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.

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.

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.

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.

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.

Implicit Measures

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 on this task and it was not included in any statistical analyses.

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 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 5. 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).

Significant differences were found for dwell time between task sessions (Kruskal-Wallis $H = 40.720$, $df = 3$, $p = 0.000$). 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 ($p=0.000$) and Transactions ($p = 0.000$) but not Browsing. Browsing task sessions were also significantly longer than both Fact Finding ($p = 0.000$) and Transactions ($p = 0.000$).

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.

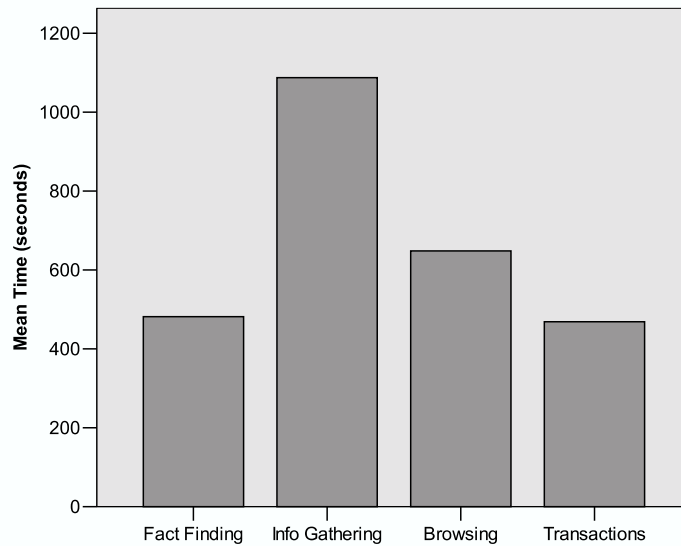


FIG. 5. The mean time (in seconds) spent completing each task.

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 Transactions task sessions may exhibit highly variable.

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 6 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 the Discussion section.

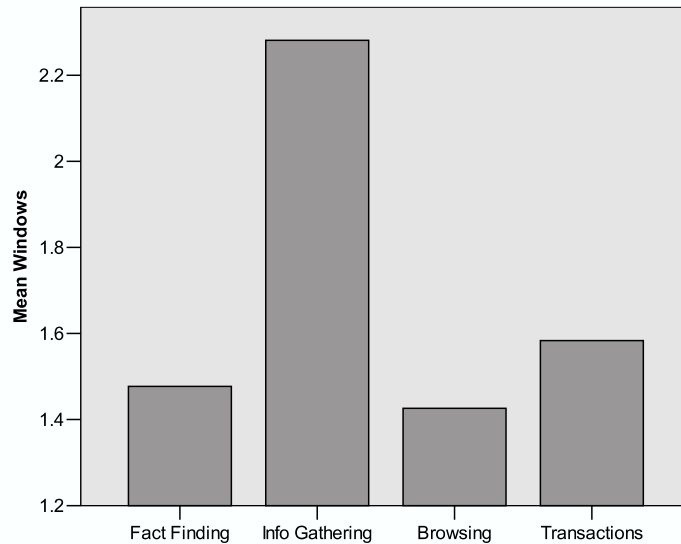


FIG. 6. The mean number of windows opened during each task.

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 7 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).

Significant differences were found for the number of pages loaded between task sessions (Kruskal-Wallis $H = 49.904$, $df = 3$, $p = 0.000$). 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 ($p = 0.000$), Browsing ($p = 0.000$) and Transactions ($p = 0.000$).

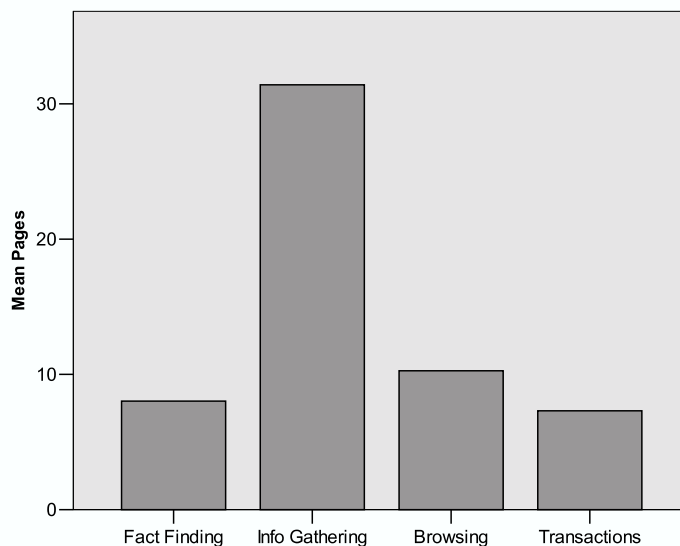


FIG. 7. The mean number of pages opened during each task.

Web Browser Navigation Mechanisms

The use of Web browser navigation mechanisms was logged during the field study (see Table 7 for a complete listing of navigation mechanisms logged). The custom Web browser differentiated between the use of the auto-complete function (Auto-Complete), selecting a URL from the drop-down address menu (Select URL), and typing a URL directly into the address bar (Typed-in URL); the use of these navigation mechanisms was logged as separate navigation events. Navigation classified as ‘Other’ consisted of navigation that was not explicitly captured, such as new windows launched by a Web page and clicked links that could not be detected (due to JavaScript, AJAX, etc.). 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. We observed that these events often occurred within Transactions (e.g., Web-based email and online applications). New Window 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 pop-up advertisements likely did not account for much of the new window usage. The custom Web browser provided participants with access to their IE auto-complete, bookmarks, history, and select-URL. This meant that participants had full access to their navigation history accrued before the study. In this analysis, we were interested in the most common methods of navigation when initiating a new task session. A more detailed exploration of the impact of task and individual differences on the use of Web browser navigation mechanisms can be found in Kellar, Watters, and Shepherd (2006b).

TABLE 7. The navigation mechanisms logged by the custom Web browser.

Auto-Complete	Forward Menu	New Window
Back Button	Google Toolbar	Other
Back Menu	History	Reload Button
Favorites	Home Button	Select URL
Forward Button	Hyperlinks	Typed-in URL

We observed several types of navigation mechanisms used to initiate new tasks: auto-complete, back button, clicked links, bookmarks, Google toolbar, URLs selected from the drop-down address menu, and typed-in URLs. Overall, typed-in URLs were the most common navigation mechanism used to initiate a new task session. The proportions of navigation mechanisms used across all tasks are shown in the bar chart displayed in Figure 8. For ease of readability, the navigation mechanisms with minimal use were not included in this figure.

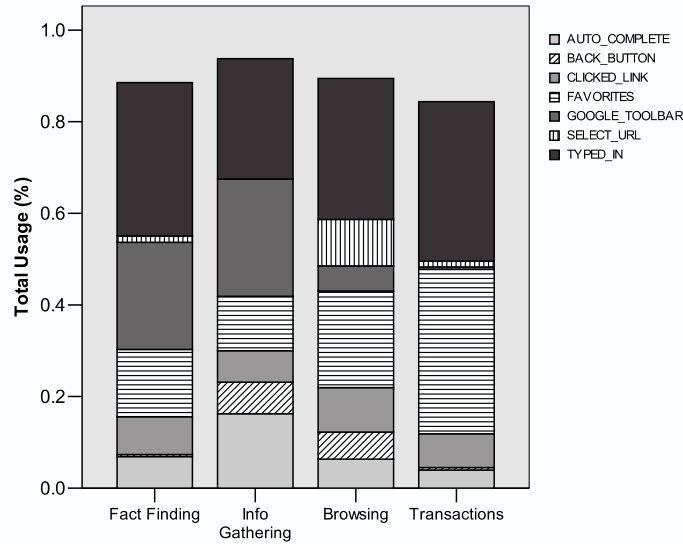


FIG. 8. The proportion of navigation tools used to initiate a new task.

Within Fact Finding task sessions, there were significant differences between the navigation mechanisms used [$\chi^2(9, N=218) = 233.101, p = 0.000$]. 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) showed that with the exception of the Google toolbar, typed-in URLs were used more often than all other navigation mechanisms ($p = 0.000$ for all). Similarly, the use of the Google toolbar was significantly higher than all other navigation mechanisms ($p = 0.000$ for all), with the exception of bookmarks.

Within Information Gathering task sessions, there were significant differences between the navigation mechanisms used [$\chi^2(7, N=160) = 78.800, p = 0.000$]. 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 more evenly distributed; 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.

There were significant differences within Browsing task sessions between the navigation mechanisms used [$\chi^2(10, N=237) = 216.878, p = 0.000$]. 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 ($p = 0.000$ for all).

Within Transactions, there were significant differences between the navigation mechanisms used [$\chi^2(10, N=557) = 1099.853, p = 0.000$]. 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 ($p = 0.000$ for all). Finally, task sessions labelled as Other were most commonly initiated using typed-in URLs (9/20 – 45%).

Time of Day

The time during which a task session was initiated was categorized 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 9 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.

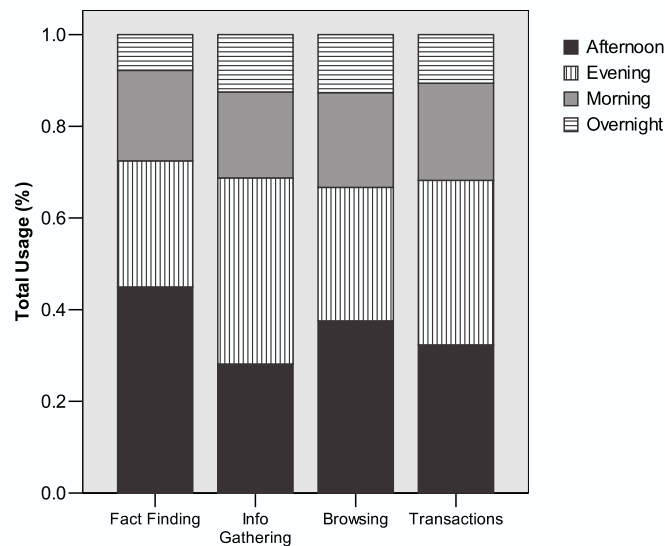


FIG 9. The proportion of tasks across time of day.

Within Fact Finding task sessions, there were significant differences between time of day [$\chi^2(3, N=218) = 63.505, p = 0.000$]. 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 ($p = 0.000$), evening ($p = 0.003$), and overnight ($p = 0.000$).

Within Information Gathering task sessions, there were significant differences between time of day [$\chi^2(3, N=160) = 28.750, p = 0.000$]. 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 [$\chi^2(3, N=237) = 32.755, p = 0.000$]. 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 ($p = 0.000$) but not the evening.

Within Transactions, there were significant differences between the navigation types [$\chi^2(3, N = 557) = 87.919, p = 0.000$]. 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 ($p = 0.000$) and overnight ($p = 0.000$).

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

Use of Google

The use of Google has become ubiquitous in today's Web environment. Aula, Jhaveri, and Kaki (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 the next section, Site-specific Searches), 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 sessions 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 heat sink fan"). The queries submitted during Information Gathering task sessions were more vague and tended to represent general topics rather than specific questions.

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.

Use of Browser Functions

Browser functions were logged as they were used and their use was associated with the task session being performed in the corresponding window. We were interested in how the use of these functions differed across task sessions. The following browser functions were logged during the field study: copy, paste, cut, find on page, print, save, and the creation of new bookmarks. 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 [$\chi^2(3, N =67) = 24.164, p = 0.000$].

TABLE 8. The use of browser functions within tasks.

Function		FF	IG	BR	TR
		(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%	32/160 20%	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%

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).

DISCUSSION

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 it is most likely to be specific to the population sampled in this study.

TABLE 9. General task characterization.

Fact Finding	Information Gathering
<ul style="list-style-type: none"> • Shorter duration • Small number of pages viewed • Large search component • Relatively longer queries • Little use of browser functions • Typed in URLs, Google Toolbar, Bookmarks 	<ul style="list-style-type: none"> • 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
<ul style="list-style-type: none"> • Shorter duration • Small number of pages viewed • Often repeated • Little use of browser functions • Typed-in URLs, Bookmarks, Select-URL 	<ul style="list-style-type: none"> • 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, lasted eight minutes on average, and recorded an average of 1.5 windows opened during a task session. 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 related and personal tasks. We observed a relatively small number of pages viewed during Fact Finding task sessions. 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. On average, 2.3 windows were opened per tasks session. We observed a relatively high number of pages loaded during Information Gathering task sessions. Many Information Gathering tasks were related to participants' course or research work. Over half of all Information Gathering tasks were repeated at least once. 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 and an average of 1.4 windows were opened during the task session. We observed a relatively small number of pages viewed during Browsing task sessions. 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 Transactions task session was close to eight minutes. Transactions were the most often repeated task, with 95% of all Transactions 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. In Kellar, Watters, and Shepherd, (2006a), we present a classification of Web information tasks, consisting of three main information goals: Information Seeking, Information Exchange, and Information Maintenance. We found that while all tasks categorized as Transactions by our participants shared the same goal (Information Exchange), there was a clear distinction two types of tasks. Tasks with a communication component were more strictly defined as Communications (e.g., email, Web publishing) while those tasks were based on the exchange on information through online actions were categorized as Transactions (e.g., banking, online shopping). Further research is needed to more carefully characterize these two task types.

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, where participants interacted much more with the Web browser, loaded significantly more Web pages and used more browser functions, for a longer period of time.

When we compared our results with the previous research (Choo et al., 2000; Morrison et al., 2001; Sellen et al., 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.

Implications and Directions for Future Research

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. These results provide insight into those tasks that need stronger support in Web browsers 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.

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 but that 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	

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 engaging in repeated tasks in order to monitoring new information, re-find previously found information, and engage in 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 Transactions task sessions appeared to be primarily due to monitoring of particular Web sites. Each of these different types of tasks require different support. While Web browsers and information systems may not be able to reliably anticipate what type of static information users may want to re-find during subsequent Fact Finding task sessions, improved history functions (as discussed below) 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. 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.

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. Providing a visual representation or traces of all Web browser interactions may help users to work more efficiently during a task session as well as during future task session.

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 et al., 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 session, 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 and more research is needed into how to provide more effective representation of previous visited pages. As reported in a more detailed analysis of the use of Web browser mechanism in this study (Kellar et al., 2006b), a history function that better supports individual differences among users, such as filing, searching, and typed commands, 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.

Windows 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 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 strategy (Hawkey & Inkpen, 2005b; Weinreich et al., 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.

Methodological Decisions and Study Limitations

It is important that we also acknowledge the limitations of this study. We used a convenience sample population consisting of university students, meaning that we can not expect that our results will generalize to a more general population. 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 10 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 examining user behaviour on the Web (Hawkey & Inkpen, 2005a, 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.

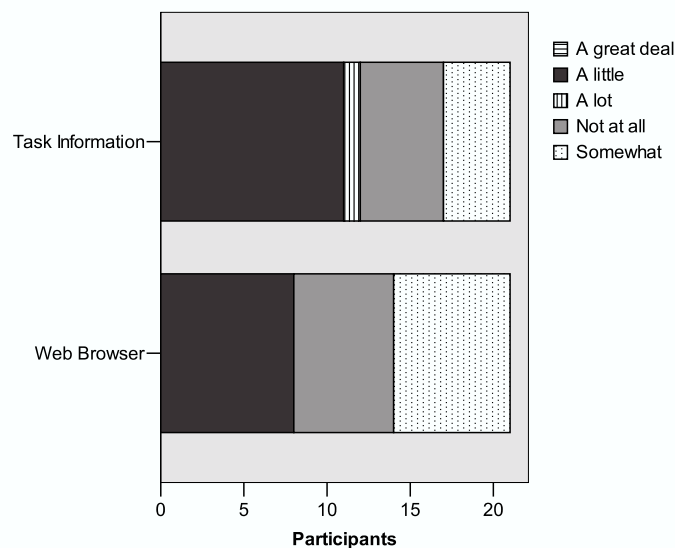


FIG 10. Participant responses to questions asking if having to record their task information and use a custom web browser impacted how they normally work on the Web.

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.

CONCLUSIONS AND FUTURE WORK

We have presented the results of a study where we examined how participants interacted differently with their Web browsers across information seeking tasks. Within each type of task (Fact Finding, Information Gathering, Browsing, 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 and 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.

We have two main areas to explore as part of our future work. First we will explore whether the implicit measures collected could be used to predict the type of task in which a user is engaging. A preliminary exploration of the usefulness of implicit measures to predict users' tasks using individual decision trees models has shown promise, with prediction accuracy ranging from 44% to 94% (Kellar & Watters, 2006). We would like to further explore other machine learning techniques to improve the accuracy of the task predictions based on implicit measures. If successful, predictive models would be useful in providing proactive support for information seeking tasks.

As this study was not designed to formally study monitoring behaviour, the second piece of future work will focus on Web-based monitoring. While we were able to gather informal observations about monitoring, we were not able to study it in a precise manner. However, it has become clear through this study that monitoring is a frequent Web activity. We would like to further study whether monitoring is in fact an independent information seeking task or simply an activity across all Web information tasks. Therefore, a more focused study will be conducted to better understand the types of information that users monitor as well as how monitoring tasks are carried out within the Web browser.

ACKNOWLEDGEMENTS

This research was funded by the Natural Sciences and Engineering Research Council of Canada. We would like to thank the anonymous reviewers for their insightful and helpful comments, members of the EDGE and WIFL research labs, and all study participants for their time and effort.

REFERENCES

- Atterer, R., Wnuk, M., & Schmidt, A. (2006). Knowing the User's Every Move: User Activity Tracking for Website Usability Evaluation and Implicit Interaction (pp. 203-212). Proceedings of the 15th International Conference on World Wide Web (WWW 2006), Edinburgh, Scotland.
- Aula, A., Jhaveri, N., & Kaki, M. (2005). Information Search and Re-Access Strategies of Experienced Web Users (pp. 583 - 592). Proceedings of the 14th International Conference on World Wide Web (WWW 2005), Chiba, Japan.

- Aula, A., & Kaki, M. (2003). Understanding Expert Search Strategies for Designing User-Friendly Search Interfaces (pp. 759-762). Proceedings of the IADIS International Conference WWW/Internet (ICWI 2003), Algarve, Portugal.
- Ayers, E.Z., & Stasko, J.T. (1995). Using Graphic History in Browsing the World Wide Web. Proceedings of the 4th International Conference on World Wide Web Boston, MA.
- Beitzel, S.M., Jensen, E.C., Chowdhury, A., Grossman, D., & Frieder, O. (2004). Hourly Analysis of a Very Large Topically Categorized Web Query Log (pp. 321-328). Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2004), Sheffield, United Kingdom.
- Belkin, N.J. (1980). Anomalous States of Knowledge as a Basis for Information Retrieval. *The Canadian Journal of Information Science*, 5, 133-143.
- Broder, A. (2002). A Taxonomy of Web Search. *SIGIR Forum*, 36(2), 3-10.
- Byrne, M.D., John, B.E., Wehrle, N.S., & Crow, D. (1999). The Tangled Web We Wove: A Taxonomy of WWW Use (pp. 544-551). Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '99), Pittsburgh, PA.
- Catledge, L.D., & Pitkow, J.E. (1995). Characterizing Browsing Strategies in the World-Wide Web (pp. 1065 - 1073). Proceedings of the Third International World-Wide Web Conference, Darmstadt, Germany.
- Choo, C.W., Detlor, B., & Turnbull, D. (2000). Information Seeking on the Web: An Integrated Model of Browsing and Searching. *First Monday*, 5(2), Online at http://firstmonday.org/issues/issues5_2/choo/index.html.
- Claypool, M., Le, P., Waseda, M., & Brown, D. (2001). Implicit Interest Indicators (pp. 33-40). Proceedings of the 6th International Conference on Intelligent User Interfaces (IUI 2001), Santa Fe, NM.
- Cockburn, A., & McKenzie, B. (2001). What Do Web Users Do? An Empirical Analysis of Web Use. *International Journal of Human-Computer Studies*, 54(6), 903-922.
- Cooper, M.D. (2001). Usage Patterns of a Web-Based Library Catalog. *Journal of the American Society for Information Science and Technology*, 52(2), 137-148.
- 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.
- Fidel, R., & Efthimiadis, E. (1999). Web Searching Behavior of Aerospace Engineers (pp. 319-320). Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '99), Berkeley, CA.
- Grace-Martin, M., & Gay, G. (2001). Web Browsing, Mobile Computing and Academic Performance. *Educational Technology & Society*, 4(3), 95-107.
- Hawkey, K., & Inkpen, K. (2005a). Privacy Gradients: Exploring Ways to Manage Incidental Information During Co-located Collaboration. (pp. 1431 - 1434). Proceedings of the CHI '05 Extended Abstracts on Human Factors in Computing Systems, Portland, OR.
- Hawkey, K., & Inkpen, K. (2005b). Web Browsing Today: The Impact of Changing Contexts on User Activity (pp. 1443 - 1446). Proceedings of the CHI '05 Extended Abstracts on Human Factors in Computing Systems, Portland, OR.

- Hawkey, K., & Inkpen, K. (2006). Examining the Content and Privacy of Web Browsing Incidental Information (pp. 123 - 132). Proceedings of the 15th International Conference on World Wide Web (WWW 2006), Edinburgh, Scotland.
- Herder, E. (2005). Characterizations of User Web Revisit Behavior. Proceedings of the Workshop on Adaptivity and User Modeling in Interactive Systems (ABIS 2005), Saarbrücken, Germany.
- Hölscher, C., & Strube, G. (2000). Web Search Behavior of Internet Experts and Newbies (pp. 337-346). Proceedings of the 10th International Conference on World Wide Web (WWW 2000), Amsterdam, The Netherlands.
- Jansen, B.J., & Spink, A. (2003). An Analysis of Web Documents Retrieved and Viewed (pp. 65-69). Proceedings of the 4th International Conference on Internet Computing, Las Vegas, NV.
- 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., Spink, A., & 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., & Rähkä, K.-J. (2005). The Advantages of a Cross-Session Web Workspace (pp. 1949-1952). Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2005), Portland, OR.
- Kaasten, S., Greenberg, S., & Edwards, C. (2002). How People Recognize Previously Seen Web Pages from Titles, Urls and Thumbnails (pp. 247-265). Proceedings of the Human Computer Interaction 2002, London, UK.
- Kellar, M., & Watters, C. (2006). Using Web Browser Interactions to Predict Task (pp. 843-844). Proceedings of the 15th International Conference on World Wide Web (WWW 2006), Edinburgh, Scotland.
- Kellar, M., Watters, C., & Shepherd, M. (2006a). A Goal-Based Classification of Web Information Tasks (To Appear). Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2006), Austin, TX.
- Kellar, M., Watters, C., & Shepherd, M. (2006b). The Impact of Task on the Usage of Web Browser Navigation Mechanisms (pp. 235-242). Proceedings of Graphics Interface (GI 2006), Quebec City, Canada.
- Kelly, D., & Belkin, N.J. (2001). Reading Time, Scrolling and Interaction: Exploring Implicit Sources of User Preference for Relevance Feedback (pp. 408-409). Proceedings of the 24th Annual International Conference on Research and Development in Information Retrieval (SIGIR 2001), New Orleans, LA.
- Kelly, D., & Teevan, J. (2003). Implicit Feedback for Inferring User Preference: A Bibliography. *SIGIR Forum*, 37(2), 18-28.
- Kim, K.-S., & 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.
- 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., & Cho, J. (2005). Automatic Identification of User Goals in Web Search (pp. 391-400). Proceedings of the 14th International Conference on World Wide Web (WWW 2005), Chiba, Japan.
- MacKay, B., Kellar, M., & Watters, C. (2005). An Evaluation of Landmarks for Re-Finding Information on the Web (pp. 1609 - 1612). Proceedings of the CHI '05 Extended Abstracts on Human Factors in Computing Systems Portland, OR.

- Marchionini, G. (1995). *Information Seeking in Electronic Environments*. New York, NY: Cambridge University Press.
- Martzoukou, K. (2005). A Review of Web Information Seeking Research: Considerations of Method and Foci of Interest. *Information Research*, 10(2), Available at <http://InformationR.net/ir/10-12/paper215.html>.
- Mat-Hassan, M., & 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.
- McGrath, J.E. (1995). Methodology Matters: Doing Research in the Behavioral and Social Sciences. In J.G. R. Baeker, W. Buxton, and S. Greenberg (Ed.), *Human-Computer Interaction: Toward the Year 2000* (pp. 152-169).
- Milic-Frayling, N., Sommerer, R., & Rodden, K. (2003). Webscout: Support for Revisitation of Web Pages within a Navigation Session (pp. 689-693). *Proceedings of the 2003 IEEE/WIC International Conference on Web Intelligence (WI '03)*, Halifax, Canada.
- Montgomery, A.L., & Faloutsos, C. (2001). Identifying Web Browsing Trends and Patterns. *IEEE Computer*, 34(7), 94-95.
- Morita, M., & Shinoda, Y. (1994). Information Filtering Based on User Behavior Analysis and Best Match Text Retrieval (pp. 272-281). *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '94)*, Dublin, Ireland.
- Morrison, J.B., Pirolli, P., & Card, S.K. (2001). A Taxonomic Analysis of What World Wide Web Activities Significantly Impact People's Decisions and Actions (pp. 163-164). *Proceedings of the CHI '01 Extended Abstracts on Human Factors in Computing Systems*, Seattle, WA.
- Moyle, M., & Cockburn, A. (2003). The Design and Evaluation of a Flick Gesture for 'Back' and 'Forward' in Web Browsers (pp. 39-46). *Proceedings of the 4th Australasian Conference on User Interface (AUIC 2003)*, Adelaide, Australia.
- Oard, D.W., & Kim, J. (2001). Modeling Information Content Using Observable Behavior (pp. 38-45). *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2001)*, Washington, DC.
- Pitkow, J.E., & Kehoe, C.M. (1996). Emerging Trends in the Www User Population. *Communications of the ACM*, 39(6), 106 - 108.
- Rieh, S.Y. (2003). Investigating Web Searching Behaviour in Home Environments (pp. 255-264). *Proceedings of the Annual Meeting of the American Society for Information Science and Technology (ASIS&T 2003)*, Long Beach, CA.
- Rose, D.E., & Levinson, D. (2004). Understanding User Goals in Web Search (pp. 13-19). *Proceedings of the 13th International Conference on World Wide Web (WWW 2004)*, New York, NY.
- Rozanski, H.D., Bollman, G., & Lipman, M. (2001). Seize the Occasion! The Seven-Segment System for Online Marketing. *Strategy & Competition*, Online at: <http://www.strategy-business.com/press/article/19940?pg=19940>.
- Schiano, D.J., Stone, M., & Bectarte, R. (2001). Search and the Subjective Web (pp. 165-166). *Proceedings of the CHI '01 Extended Abstracts on Human Factors in Computing Systems*, Seattle, WA.
- Sellen, A.J., Murphy, R., & Shaw, K.L. (2002). How Knowledge Workers Use the Web (pp. 227-234). *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2002)*, Minneapolis, MN.
- Seo, Y.-W., & Zhang, B.-T. (2000). Learning User's Preferences by Analyzing Web-Browsing Behaviors (pp. 381-387). *Proceedings of the International Conference on Autonomous Agents 2000*, Barcelona, Spain.

Spink, A., Wolfram, D., Jansen, M.B., & 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.

Tauscher, L., & 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, 97-137.

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

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

Wilson, T.D. (1997). Information Behaviour: An Interdisciplinary Perspective. *Information Processing and Management*, 33(4), 551-572.