

SearchBar: A Search-Centric Web History for Task Resumption and Information Re-finding

Dan Morris, Meredith Ringel Morris, Gina Venolia

Microsoft Research

Redmond, WA, USA

{dan, merrie, gina.venolia}@microsoft.com

ABSTRACT

Current user interfaces for Web search, including browsers and search engine sites, typically treat search as a transient activity. However, people often conduct complex, multi-query investigations that may span long durations and may be interrupted by other tasks. In this paper, we first present the results of a survey of users' search habits, which show that many search tasks span long periods of time. We then introduce SearchBar, a system for proactively and persistently storing query histories, browsing histories, and users' notes and ratings in an interrelated fashion. SearchBar supports multi-session investigations by assisting with task context resumption and information re-finding. We describe a user study comparing use of SearchBar to status-quo tools such as browser histories, and discuss our findings, which show that users find SearchBar valuable for task reacquisition. Our study also reveals the strategies employed by users of status-quo tools for handling multi-query, multi-session search tasks.

Author Keywords

Web search interfaces, Web history, persistent search.

ACM Classification Keywords

H3.3. Information search and retrieval: Search process.

INTRODUCTION

Web browsers and search engine sites are the primary tools people use to access the vast quantities of information available online. These tools typically treat information-seeking tasks as a transient activity, considering each of a user's search and browse actions as unrelated events rather than as part of a larger task. For example, if a user is searching for information on digital cameras, he will perform several low-level interactions with a Web browser and one or more search engines, *e.g.* typing the URL of a search engine site, querying for "digital camera ratings", clicking a result, hitting the back button, submitting a second query, etc. Although the browser views these actions as unrelated events, the user views them at a task

level, *i.e.*, "figure out what digital camera to purchase", and current search UIs don't provide affordances at this level.

Considering search and browse actions in isolation results in impoverished user interfaces that do not adequately assist users with common scenarios such as re-finding previously encountered Web pages, resuming an activity after an interruption, or conducting complex, multi-session investigations. In this paper, we present evidence that users often conduct multi-session Web investigations, along with evidence that such tasks are not adequately supported by current tools. We then present SearchBar, a system for proactively and persistently storing query histories, browsing histories, and users' notes and ratings in an interrelated fashion. We describe a study comparing SearchBar to status quo interfaces for Web information-seeking tasks, and report detailed findings on both the benefits of the SearchBar approach and on users' general strategies for handling complex, exploratory search tasks. We conclude with a discussion of related and future work.

Throughout this paper, we use the term *query* to refer specifically to keyword search using a general-purpose or vertical search engine, and the terms *search* and *investigation* to refer to the high-level task of finding content online using a variety of methods, such as querying, browsing, or navigating directly to known URLs.

MOTIVATION

In this section, we discuss research that highlights some of the shortcomings of current tools for information-seeking tasks, specifically re-finding information and managing interruptions during information-seeking.

Re-finding and Re-searching

Several studies have found that people frequently revisit previously viewed online content. For example, Aula *et al.*'s survey of experienced Web users [1] found that people use search engines more frequently than history or bookmarking tools to re-find online information, but often have difficulty remembering what queries they had used when they originally discovered the content in question. Jones *et al.* [11] conducted interviews about how people "keep found things found", and discovered that the most common strategy for re-finding information was to "do nothing", *i.e.* to not take explicit actions such as bookmarking or taking notes, but rather to assume (sometimes incorrectly) that one would be able to re-find

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2008, April 5–10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00.

needed content later on. Kellar *et al.* [14] interviewed Web users about content-monitoring habits, and identified information-gathering as a class of activity that often involved revisiting sites to monitor for changed information. They found that people relied on a combination of bookmarks, auto-completion, and search engines to return to the sites they monitored. Teevan *et al.* [25] conducted a study of Yahoo!’s query logs showing that nearly 40% of queries were attempts to re-find previously encountered results. Tauscher and Greenberg [24] studied users’ browser logs over a six-week period, and found that 58% of the pages visited were actually re-visits, while Cockburn and MacKenzie [5] reported an even higher rate of revisitation (81%). Wen [28] asked study participants to re-find pages that they had visited during a search task, and observed an overall success rate below 20%. Nearly all participants in Wen’s study tried to re-create their search paths to re-find information. A more recent study by Obendorf *et al.* [19] found that about 44% of pages visited were revisits. Their study revealed that users employed strategies such as re-searching and re-tracing to re-find content viewed a week or more in the past, but the authors noted that such strategies often frustrated users, due to problems remembering their original query or to expired link coloring information. They also note that browser history was nearly unused by participants, initiating only 0.2% of all actions despite the high rate of re-visitation.

In summary, there is significant evidence that re-finding information constitutes a significant portion of Web use, and that current browser-based mechanisms that support this activity are poorly utilized. Obendorf *et al.* noted that “A time and task based search history could help to redo searches.” SearchBar is a realization of such a system.

Interruptions, Multi-tasking, and Multi-Session Tasks

In addition to re-finding previously encountered information, another user need not explicitly addressed by current Web interfaces is task resumption after an interruption. This interruption might be self-generated (*i.e.*, multi-tasking) or externally generated.

Previous Research on Interruptions and Search

Interruptions often cause people to forget the nature of a task they meant to complete [18]. Such interruptions are common – a diary study of knowledge workers found an average of 50 task shifts per person during a typical work week [6], and a log study found an average of 3.74 e-mail or IM-driven task shifts per hour, taking between 8 and 10 minutes on average for the person to return to the interrupted task [9]. However, previous work has not specifically explored how users manage interruptions during Web search tasks, as we do in our study.

Sellen *et al.*’s study of knowledge workers’ Web use [23] reported that 40% of the “information gathering” activities they observed (similar to Broder’s definition of “informational searches” [3]) were not completed in a single sitting, often due to external interruptions. Such “multi-session” information-seeking activities were also

Table 1. The 170 respondents who had resumed a Web search after a gap in time reported the longest gap they recalled between suspension and resumption.

Time Before Resumption	Respondents
Minutes	1.2%
Hours	24%
Days	51%
Weeks	13%
Months	10%
Years	1.2%

reported by Obendorf *et al.* [19], whose client-side log study of Web use found that an average of 15% of Web page revisits occurred after periods of a week or more; while this represented only a small portion of Web activity, the users in their study indicated that these long-term revisits were important but were difficult to achieve.

Survey on Search Habits

To learn more about multi-session Web investigations, we conducted our own formative study, surveying information workers at a large technology company about their Web search habits and needs. Our survey was sent to 740 people and was completed by 204 (28% response rate). 80% of respondents were male. Ages ranged from 21 – 61 years old (median = 36). Respondents had varied job roles within the company, such as software developers, researchers, managers, administrative assistants, attorneys, and student interns. All respondents self-identified as average (27%) or expert (74%) Web searchers, with none identifying as novice. All but one respondent reported using a search engine to find information online at least once per day. Our survey asked questions about a variety of Web use habits, many of which are beyond the scope of this paper. Here, we report only on portions of the survey relevant to the suspension, resumption, and/or re-execution of information-seeking tasks. We also highlight that the data presented in this section are self-reported, so we use these results as design guidelines but cannot verify their absolute accuracy.

We found that multi-session information-seeking tasks were a common experience among this demographic, with 83% of respondents answering “Yes” to the question “*Have you ever had a Web information-seeking task that has lasted for longer than a single session (i.e., a task which you put aside and resumed at a later time)?*”

We asked several follow-up questions to the 170 respondents who reported engaging in multi-session investigations. First, we asked them to report the longest gap they had experienced between suspending and resuming an investigation. The results in Table 1 show these responses, which indicate that time gaps before resumption can be lengthy – 75% indicated experiencing gaps of at least a full day, while 24% reported gaps of a week or more before resumption.

We then asked these respondents to describe the strategies

Table 2. Percent of respondents (n = 170) who mentioned using the following strategies when resuming a suspended investigation.

Strategy	Respondents
Active Storage / Active Retrieval	55%
Written or typed notes	30%
Browser bookmarks	26%
Save Web pages locally as files	5.3%
E-mail to self	2.4%
Online bookmarks/tagging	1.8%
Print Web pages	0.6%
Passive Storage / Active Retrieval	8.2%
Browser history	7.6%
Online query history	1.2%
Passive Storage / Passive Retrieval	49%
Memory	36%
Leave browser open	14%
Link coloring	2.4%
Autocomplete	2.4%

they use to refamiliarize themselves with the context of their suspended investigations upon resumption.

The common strategies described by respondents can be distinguished along two axes: *initiative* (whether the user was *active* or *passive*) and *stage* (whether the strategy applies to the *storage* or *retrieval* of the task context). For example, an active approach to storage attempts to preserve the state, strategies, and/or results of a search task for later use, while a passive approach to storage involves no explicit user actions to save information during a search session. An active approach to retrieval seeks out mechanisms that have captured data (either explicitly or implicitly) that may help in re-establishing a search context, while a passive retrieval strategy uses only information that is automatically revealed by basic browser or search engine interfaces. We describe strategies by indicating the initiative type for each of those stages.

Active Storage/Active Retrieval strategies reported in our survey included saving the contents of a website (either into a local file or by printing), bookmarking a website (either locally in the browser or via an online bookmarking utility), and note-taking (either by hand, by typing into a document, via self-email, or via instant messages).

Examples of Passive Storage/Active Retrieval strategies include using browser history to find information from a previously executed search task, or visiting a search-engine-specific query history tracker to re-find a useful query.

Passive Storage/Passive Retrieval strategies reported by our survey population included relying on memory when resuming a search (and possibly starting over and repeating previously issued queries), leaving the browser open on the machine when suspending a task so its state was preserved

upon resumption, relying on the browser's autocomplete functionality to serve as a memory aid when typing queries, and relying on the browser's coloring of visited links to prompt memory about previously viewed Web sites.

Table 2 shows the percentage of people in our survey population who reported employing each of the above strategies (note that the numbers total to more than 100%, as some people listed several of these strategies). Note that no respondents described using an Active Storage/Passive Retrieval strategy, which is not surprising given that such a combination represents wasted storage effort.

In summary, previous work and our own survey indicate that information-seeking tasks are frequently interrupted, yet current Web browsers provide only partial support for task resumption. Furthermore, our data show that an effective tool for re-finding information should support a variety of re-access strategies.

SEARCHBAR

To improve the user experience of resuming suspended investigations and re-finding previously encountered online content, we created SearchBar (Figure 1), a system for capturing, storing, and presenting an annotatable, integrated search and browsing history within the context of a Web browser. SearchBar is implemented as a plugin for Internet Explorer. As such, it is able to monitor all navigation events, to recognize URLs corresponding to common search engines, and to track a user's browsing history subsequent to a search. More details on these features are provided below. Once enabled, SearchBar is visible in any new tabs or browser windows the user creates, and content is mirrored across all of these views. SearchBar can be hidden or revealed using a keyboard shortcut.

SearchBar presents visited Web pages hierarchically according to the search queries that led to those visits. A recent Pew Internet Project report [20] found that search engine use was the second most common online activity, after e-mail. The ubiquity of search in Web use, users' tendencies to repeat previously executed searches [25, 26], and the tendency of informational [3, 21] and exploratory [29] search tasks to involve multiple query-browse iterations inspired us to use queries as a fundamental organization metaphor in the SearchBar interface. Thus, SearchBar functions as a rich history mechanism, proactively capturing all queries issued and URLs visited subsequent to each query. The search-centric structure is manifested by showing all URLs browsed to subsequent to a query as child nodes of that query, in reverse-chronological order, thus revealing to users the process through which they reached a piece of information. Clicking on a URL or query navigates the browser to that page or those search results.

For the large portion of users who choose to employ only Passive Storage strategies, all queries and browsing are proactively captured and stored under SearchBar's "Default" topic. However, for those users who prefer to

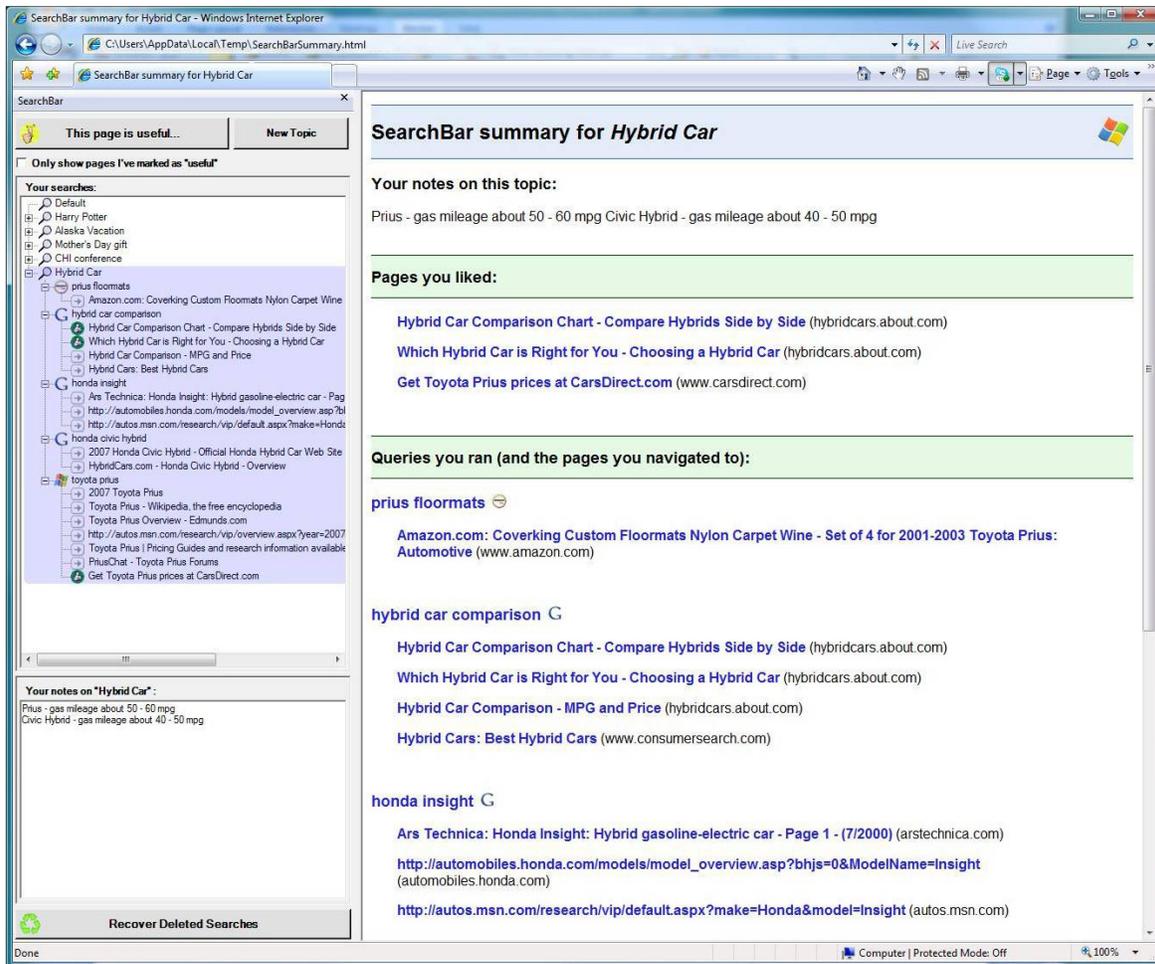


Figure 1: A screenshot showing SearchBar running within Internet Explorer. The SearchBar (left pane) includes a hierarchical history of recent search topics, queries, and results visited, along with the user’s notes on the current topic. The right pane is a standard Web browser, currently showing a summary of the active SearchBar topic.

employ Active Storage strategies, SearchBar offers the ability to create custom-named topics. All search engine queries and all navigation events subsequent to a query are stored under the currently-selected topic. The user can move URLs and queries among topics by dragging and dropping, if desired. Since our own survey and others’ prior work have reported that users often return to search tasks after gaps of a week or more, the topical organization can add value by helping such users more quickly resume their task contexts. Users may also delete topics, queries, or URLs from SearchBar via a right-click context menu. Deletion undo is provided via a recycle bin.

30% of our survey respondents reported taking notes as a way to help manage multi-session investigations. However, since these notes are stored in files or on paper, they can be difficult for users to retrieve when they eventually resume their searches. Hence, we added a note-taking region to the SearchBar itself, allowing users to keep their notes where they will be immediately visible when users resume their tasks. Notes are associated with the currently-selected topic and are automatically saved.

Users can also press the “thumbs up” button to indicate that a particular page is especially relevant to the current topic. This associates a “thumbs up” icon with the currently viewed Web page’s SearchBar entry, making the page easier to identify and return to quickly and also elevating the page to the top of the that topic’s summary (see below). Users can select a checkbox which filters the SearchBar to only display items that have been rated in this manner.

SearchBar captures queries from several major search engines (Yahoo!, Google, and Windows Live Search) as well as several popular vertical search services (e.g., Amazon.com, Wikipedia, CNET, and CitySearch). A branding icon next to each query indicates the source used for that search. SearchBar is extensible, allowing easy addition of additional search sources. This enables SearchBar to produce a more comprehensive history than engine-specific query histories, since it incorporates searches from the variety of sources that people utilize.

Clicking on any topic’s title displays an automatically-generated summary in the browser window, which shows any notes the user has taken on that topic and lists all the

pages the user has indicated as especially relevant, as well as giving a chronological history of all the queries issued and pages visited on that topic. All queries and pages in the summary and in the standard SearchBar display are hyperlinks that return the user to the associated query results or Web page.

EVALUATION

We conducted a study to evaluate the usability of SearchBar, and to explore whether it would be more effective in allowing users to manage, suspend, and resume multi-session, multi-query searches than status-quo tools. All of our study's tasks fall into Broder's "informational" task category [3]; this is consistent with our focus on complex, multi-session searches.

Demographics

We recruited 16 paid participants (9 female), none affiliated with our institution, ranging in age from their 20's to their 60's. Participants had a variety of occupations, and included an attorney, an editor, a business owner, a marketing specialist, and an insurance salesperson. Only two participants had technical professions.

Methodology

Our study had a between-subjects design, with eight participants each in the control and experimental conditions. In the control condition, participants completed the study tasks using Internet Explorer 7, and in the experimental condition participants used Internet Explorer 7 with the SearchBar plugin. The two conditions were otherwise identical. Each participant came to the lab for two sessions, each 90 minutes long, scheduled one week apart.

When participants arrived for the first session, they were seated at a dual-monitor (19" each) computer. Internet Explorer 7 (with the SearchBar plugin, in the experimental condition) was open on the left monitor, with the home page set to a search engine site, and a "Travel Itinerary" document was open on the right monitor.

Participants were told that they would play the role of the administrative assistant to the CEO of a fictional corporation. The CEO had requested that the assistant complete the itinerary for his upcoming business trip to Prague (none of the participants had ever visited Prague). The participant was given a printed copy of an e-mail from the CEO describing his itinerary requirements, such as the dates of his travel, his desire to visit three famous historic sites, see three theatre productions, etc. They were instructed to fill in the travel itinerary document, which initially contained only headings such as "museums", "theatre", etc., as per the requests in the CEO's e-mail.

Participants were also told that the CEO might interrupt them by sending instant messages (IMs) if he required immediate assistance. Participants were also told that they could use any available tools to assist them in managing their task, including pen and paper (provided), the Word, Excel, and Notepad applications (to which shortcuts were placed on the computer's desktop), or any features of

Internet Explorer (e.g., bookmarks, history). If participants were in the experimental condition, we also introduced them to SearchBar, allowing them to complete a tutorial and ask questions. Participants were not told anything about what would occur during the second week of the study.

After each participant had been working on the travel itinerary task for seven minutes, the experimenter sent an IM (playing the role of the CEO) asking for the name of a restaurant in Chicago that served New York Style pizza where he could take a client for lunch during a business trip. After the participant sent a response, the CEO told him to resume working on the itinerary. After the participant worked on the itinerary for seven more minutes, the CEO sent another IM, asking the assistant to recommend a digital camera he should purchase that met certain resolution and zoom requirements. Afterwards, the participant again resumed work on the itinerary for seven more minutes, when the CEO sent another IM requesting statistics about the incidence of the flu during 2005 and 2006, in order to help him decide whether to pay for employee flu shots. After this, the participant returned again to the itinerary task for one more minute, after which point the experimenter asked him to stop. To minimize user frustration and fatigue, if the participant took longer than seven minutes on any of the interrupting tasks, the CEO sent a follow-up IM asking for the best URL the participant had found so far, and then telling him or her to return to the travel itinerary task.

Participants then completed a questionnaire soliciting demographic information, information about their search habits, information about the session's tasks and the strategies employed by the participant, and feedback on the SearchBar tool (experimental condition only). After each participant left, the experimenter saved all state, including the browser's bookmarks and history, SearchBar data files, the itinerary document, and any additional digital or paper notes taken by the participant. Note that all participants used separate accounts on the study machine, so that the history, bookmarks, etc. in the Web browser reflected only their own actions and not those of other participants.

When participants arrived for the second session, one week later¹, the setup was the same as the first week's. Internet Explorer retained any settings they had accrued during the first session (bookmarks, history, SearchBar entries, etc.). Any digital notes they saved in other documents were still on the computer, and any paper notes they took were on the desk, along with the printout of the CEO's travel requests. However, the travel itinerary document was the blank template from the first session, containing none of the content they had added the first week.

Participants were told they would reprise their role as assistant to the CEO, who wanted his travel itinerary

¹ Three participants (all in the experimental condition) deviated from this one-week gap between sessions, returning after 6 days, 8 days, and 13 days respectively.

finished. Each participant was told that the computer was in the same state in which he had left it the previous week, except that a computer virus had destroyed the itinerary file. Participants were instructed to recreate the missing contents of the itinerary to the best of their ability, and, when they had finished doing so, to continue completing the itinerary by filling in the portions they had not had time to attend to the first week. Note that the itinerary task was refined through pilot testing to be long enough such that no participant completed the entire itinerary during either week of the study. Participants were also told that the CEO might still IM them if more pressing tasks arose. For the experimental condition, a printed copy of the SearchBar tutorial was available on the desk for reference.

Once again, the experimenter sent instant messages from the “CEO” to the participant after seven minute intervals of working on the travel itinerary task. This time, each of the three instant messages solicited information that depended on content found the first week: the name of a hotel located near the recommended pizza restaurant, the name of a store selling the recommended camera for a good price, and a comparison of the 2005-2006 flu incidence with that of prior years. Note that in each case, any specific information supplied by the participant to the experimenter the previous week, such as the camera model the participant had originally recommended, was not provided in the instant message; rather, we were interested in observing how easily participants could re-find this information.

As in the first session, after the participant responded to the final IM and had resumed work on the travel itinerary for one more minute, the experimenter stopped the session and distributed a second questionnaire, which solicited feedback on that session’s tasks and the strategies employed, and also on the SearchBar tool itself (experimental condition only).

Data gathered during the study included notes and observations from two experimenters who observed each session, questionnaires completed after each session, the itinerary documents from each of the two sessions, the transcripts of the instant messages, an enhanced browser history that included all navigations, window management events, and keyboard/mouse events, the browser’s bookmark lists, and any other digital or paper notes taken by participants. For participants in the experimental condition, we also logged the contents of the SearchBar and all interactions with the SearchBar, including passive interactions (insertion of new queries) and active interactions (clicking on results, creating topics, etc.).

RESULTS

Motivating Trends

At the end of the first session, we asked several questions about participants’ general search habits, which reconfirmed the motivations behind the creation of the

SearchBar prototype². The majority of participants reported that they have difficulty keeping track of information when conducting exploratory searches (median Likert response = 4.0), that they enter the same query terms multiple times (median = 4.0), and that they do not always finish an exploratory or informational search task in one sitting (median = 4.0).

We also asked our sixteen participants the same survey questions about multi-session searches that we asked as part of our larger survey (see Introduction section). The response pattern observed in our questionnaire is very similar to that observed in our survey, which had targeted a more technical demographic, suggesting that our survey findings generalize beyond technical workers to more diverse gender, occupational, and technical skill backgrounds. 94% of our participants reported having engaged in multi-session searches, and 80% of those who had performed multi-session searches reported resumption gaps of a day or more, with 40% recalling occasions when a week or more passed before they resumed a suspended search task.

General Task Strategies

We observed several general trends in search strategies that were common across the majority of participants. The use of the Web browser’s tabs to assist with task management was one salient trend. During the first session, all sixteen participants opened a new browser tab at least once when beginning a new task related to the CEO’s interruptions, and then returned to an old tab after finishing with the interrupting task.

The use of multiple search sources was another common trend we observed: 81% of participants used at least one vertical search engine (in addition to a general-purpose search engine) during the study, and 56% used two or more different general-purpose search engines. Overall, 88% of participants used multiple search engines at some point during the study.

Only two participants (both in the control condition) made bookmarks during the first week; these participants also accessed these bookmarks during the second session. However, only one of these participants made bookmarks relevant to the study tasks; the other used bookmarks only to mark two of his favorite search engine sites and to switch among those sites quickly.

Five participants (all in the control condition) used the browser’s history during the second session of the study. Of

² Because our data is not normally distributed, all analyses use non-parametric tests. We use the Mann-Whitney *U* test for comparisons between the control and experimental groups, and the Wilcoxon test for comparisons of the experimental participants’ week 1 versus week 2 data. Likert scale responses are on a scale of 1 to 5, where 1 = strongly disagree and 5 = strongly agree.

the five people who used history, three had difficulty finding it (two had to use the browser's "help" facility to locate the history, and one requested assistance from the experimenters in locating it).

Impact on Task Resumption

Participants' feedback indicated that SearchBar made it easier to re-find information during the second session, with less redundant work. Participants using SearchBar reported significantly more agreement with the statement "It was easy to faithfully reproduce deleted portions of the travel itinerary" (median = 4.0) than those in the control condition (median = 2.5) ($z = -1.93, p = .05$). SearchBar users also reported significantly less agreement with the statement "I repeated a lot of work that I had done last week in order to re-find relevant information" (median = 2.0) than those in the control condition (median = 4.0) ($z = -3.07, p < .01$). Also, for the "flu vaccine" task, which participants in both conditions rated as the most difficult of the tasks they performed, participants using SearchBar agreed significantly more that "It was easy to find information that I encountered last week about the flu vaccine" (median = 4.5) than those in the control condition (median = 2.0) ($z = -2.28, p < .03$).

Analysis of the experimental group's log files support participants' perceptions that SearchBar assisted them in avoiding redundant work. Participants in the experimental condition clicked on queries and pages in the SearchBar to achieve 31.7% of re-querying actions and 31.5% of re-navigation actions, respectively. The value of SearchBar seemed more apparent to users during the second week of the study, when they relied on the tool for re-navigating significantly more ($\mu = 42.2\%$ of re-navigations) than during the first week ($\mu = 23.5\%$ of re-navigations) ($z = 2.52, p < .02$) (mean numbers of re-navigations were 37 and 41.5, respectively). The remainder of re-navigation events were initiated by other in-browser mechanisms, such as auto-completion of URLs or following previously-followed links. Similarly, in the first week of the study, all sixteen participants opened a new browser tab to initiate a new search task at least once during the course of the session. However, in the second week this behavior was primarily observed among control-condition participants: six control-condition participants and only one experimental-condition participant ($z = -2.44, p < .02$). We hypothesize that this difference can be attributed to the availability of SearchBar, namely that the persistent state provided by SearchBar enabled experimental-condition participants to multi-task without opening additional tabs (which was a common strategy among control-condition participants for maintaining persistent state).

In both conditions, participants performed similarly in terms of their abilities to faithfully reproduce the contents of the deleted travel itinerary, accurately reproducing, on average, 74.7% of their entries (4.9 entries) from week 1 in the control group and 74.0% in the experimental (4.5

entries).³ However, participants using SearchBar made more progress on adding new content to the itinerary in addition to completing the reproduction of week one's entries⁴; participants in the control condition completed a mean of 2.0 extra itinerary items, while SearchBar participants completed a mean of 3.75 extra entries. Note that, while this trend suggests that SearchBar helped participants re-find content more effectively than status-quo tools, this difference is not statistically significant. However, participants' abilities to successfully re-find information in the control condition was likely artificially inflated by the study design – because the participants' computer accounts were used only during the study sessions themselves, their browser histories were relatively short and uncluttered with pages from other intervening tasks, making the history a much simpler and, likely, more usable re-finding tool than it would be in an ecologically valid setting.

Experimental participants' questionnaire responses after the second session (Table 3) show that they agreed strongly that SearchBar helped them remember what they had accomplished during the first week and made it easier to re-find relevant information about the CEO's interrupting requests.

SearchBar Usability

Table 3 shows experimental participants' Likert-scale ratings of the usability of SearchBar's features, after both the first and second session. Participants gave high overall ratings to SearchBar's usability, indicating agreement that it was easy to learn and remember how to use SearchBar. Although participants rated SearchBar as only moderately useful during the first week (median = 3.5), when they were able to use open tabs to maintain the context of interrupted tasks, the utility rating increased dramatically the second week (median = 5) ($z = -1.91, p = .05$), indicating that SearchBar was highly valued for resuming longer-term suspended tasks where it was not possible to maintain context via unchanging browser state.

Participants did not mind the screen real-estate devoted to SearchBar, indicating disagreement with the statement that it took up too much space in the browser window, even during the first week when it was less heavily utilized. Participants also indicated that they would use SearchBar both at home and at work, indicating on average that they would likely use it between one and multiple times per day in each setting. The ability to utilize topics as an

³ Note that one control-group participant's week 1 itinerary was corrupted; only the remaining 15 participants' data was considered in the comparison of week 1 and 2 itineraries.

⁴ We use itinerary items completed as a proxy for efficiency, since timing data was difficult to accurately measure because several participants completed portions of the itinerary out of order, mixing their re-creation of old content with their addition of new content.

Table 3. Experimental group’s median Likert-scale ratings of the SearchBar tool after the first and second sessions.

	Week 1	Week 2
It was easy to learn how to use SearchBar.	4.5	N/A
The organization by topics was helpful.	4	4.5
The per-topic summaries were helpful.	4	4
SearchBar was useful for accomplishing today’s tasks.	3.5	5
The “thumbs up” feature was helpful.	3.5	3.5
The ability to take notes using SearchBar was helpful.	3	3
SearchBar took up too much space in my browser window.	2	2
SearchBar was confusing to use.	1.5	1.5
It was difficult to remember to create a new topic.	4	4
It was easy to remember how to use SearchBar.	N/A	5
SearchBar helped me remember what I accomplished last week.	N/A	5
SearchBar made it easy to re-find information about cameras, restaurants, and flu vaccines that I had encountered last week.	N/A	4.5

organizational structure, the thumbs up feature, and the summaries were seen as the most useful aspects of the tool; the note-taking feature was less popular. Log file analysis shows that participants indeed used the “thumbs up” feature frequently, marking a mean of 8.25 pages each with a thumbs up during the study.

However, some participants stated that it was difficult to remember to create a new topic before beginning a search task. Log file analysis shows that participants did take the time to create topics, creating a mean of 2.0 topics each during the study. However, this average number of topics is lower than the number of tasks completed, reflecting the fact that participants used other methods, primarily the maintenance of multiple tabs in their browsers, for separating tasks during short-term interruptions and multi-tasking. Two of the eight experimental participants created six topics each, while three did not create any topics; this variance is not surprising given our survey findings showing that people differ in whether they use active or passive storage strategies during Web investigations. However, even those who utilized a more passive approach and didn’t take time to create their own topics still benefited from SearchBar’s proactive capture and automatically-imposed query/navigation structure.

DISCUSSION

The results presented in the previous section confirm that SearchBar is intuitive and easy to use. Although they were not explicitly required to use SearchBar during our study, participants made extensive use of SearchBar for re-executing queries and re-visiting Web pages and appeared

to integrate SearchBar into their workflow thoroughly and immediately. Questionnaire results also support this notion; participants indicated that SearchBar was easy to learn and easy to remember to use a week later. Perhaps most importantly, participants indicated that SearchBar was useful in accomplishing search tasks, particularly when resuming a task after a long interruption.

One of the major drawbacks of SearchBar is the need to create topics manually. Questionnaire results confirmed that participants found it difficult to remember to create a new topic when initiating a new task, although several participants did use this feature. The need to create new topics before beginning a search task is reduced significantly by the ability to create topics after the fact and move existing queries into the new topic; this is likely to be particularly common in long-term, real-world use, where users may begin a search task expecting to use only a small set of queries and later see that task evolve into a larger investigation. Note that topic creation is not always necessary: for smaller tasks, the automatic insertion of search queries into the “Default” topic often suffices, and will still provide SearchBar’s query-centric Web history and representation of query iteration. Automatic creation of new topics based on lexical, semantic, or temporal relationships among queries remains as future work.

Furthermore, SearchBar presently tracks only browsing sequences which begin with a search query, when in fact other browsing mechanisms, such as directly typing URLs into a browser, could also be used to continue a search, even without an intervening query. More sophisticated mechanisms for associating continued browsing with recent queries thus remain an important topic of future work; such mechanisms might be based on semantic similarity between a page and a recent query or temporal proximity between a navigation event and a recent query.

A longitudinal study assessing SearchBar’s scalability over time was beyond the scope of the present evaluation. The history and bookmarking systems in current browsers, for example, typically become cluttered as a result of long-term use. However, we aimed to make SearchBar more scalable than current history tools by:

- (a) Using a hierarchical model centered around elements (topics and queries) that are meaningful to the user, which allows a user to browse a dense hierarchy more efficiently than a date- or site-based hierarchy.
- (b) Enabling rating-based filtering of the history.
- (c) Highlighting a past topic if a query contained by that topic is performed again later, thereby alleviating the need to manually explore the hierarchy in many cases.

In addition to studying the utility and usability of SearchBar, our study allowed us to explore the mechanisms used by both experimental and control participants to conduct search tasks and to manage search tasks across interruptions. One particularly surprising result was the

extensive use of browser tabs among our participants: although we did not recruit participants for familiarity with tabs and did not recruit from a highly technical user population, all sixteen participants opened multiple tabs during the study. This is in contrast with results reported by Weinreich et al [27], who noted that seven of thirteen Firefox users in a 2006 log study “hardly opened any tabs”. We hypothesize that this dramatic increase in tab use represents increased adoption of tabs even in the short span of time since that work was published.

We also specifically observed a previously-unreported use of browser tabs as a mechanism for managing interruptions in search tasks: during the first session, all sixteen participants opened a new tab in response to an interrupting search task. This indicates that the creation of a new tab is potentially a valuable statistical indicator of a boundary between tasks, and may be applied to automatic delineation of Web search topics or to intelligent timing of interruptions and alerts. This may help address one of SearchBar’s primary limitations, namely its reliance on the user to delineate task boundaries.

Additionally, our evaluation suggests that typical Web users make use of multiple search engines: 88% of participants used multiple general-purpose and/or vertical search engines. This motivates the need for a client-side tool for logging query histories; server-side history tools hosted by a search engine will log only searches conducted at that site. Coupled with the privacy implications of using search engines as primary re-finding instruments, which requires storing a persistent personal browsing history on a remote server, client-side support for re-finding may be desirable.

One possible limitation of our evaluation’s ecological validity is that no additional content was inserted into SearchBar between sessions, which may not be typical of a one-week interruption in a search task. However, the tools typically used for restoring search context – history, favorites, link coloring, autocomplete notes – would all also have been subject to increased “clutter” over a week of real-world use, and we hypothesize that due to SearchBar’s use of a semantically-meaningful hierarchy based on topics and queries, users would be more easily able to sort through clutter using SearchBar than using existing mechanisms.

RELATED WORK

Browser Features

Several mechanisms for representing browsing history and re-finding information are ubiquitous in current Web browsers. Browsers provide a comprehensive and proactively-collected navigation history that can be sorted alphabetically or chronologically (although this history is rarely utilized [1]) and a mechanism for actively collecting relevant links into a hierarchy (referred to as “favorites” or “bookmarks” depending on the browser). Browsers also support coloring of HTML links conditioned on whether a link has been previously visited within a specified time window, in order to assist users in re-tracing (or avoiding

re-tracing) navigation paths. Browsers also offer automatic completion of partial URLs or terms entered into text entry fields for both navigation and search, based on recently-visited pages or recently-executed queries, which can also assist in query recall or re-finding of URLs.

Commercial Systems

Several commercial services assist users in collecting and re-finding Web-based information. Google Web History provides server-side logging of Google searches a user conducts and results subsequently clicked; a user can later view all queries and results clicked in a chronological list. This service also allows users to search pages they have visited directly from Google searches; installation of the Google Toolbar optionally extends this history list and search capacity to all Web pages a user has visited. This service has the advantage of existing on a remote server and thus traveling with a user to multiple computers, but is not integrated directly with Web browsing (as SearchBar is), does not capture queries from other search sites, and requires the user to consent to fairly comprehensive server-side logging and tracking. Del.icio.us is an online bookmarking service, which allows bookmarks to travel with a user and allows sharing of bookmarks among users. Google Notebook similarly allows server-side storage of links, and also allows storage of notes and images. Onfolio is an extension to the Windows Live Toolbar that provides similar note-taking and “clipping” functionality as a client-side application.

Research Systems

Several systems enable users to manually create collections of online content. Session Highlights [10] provides a workspace to which users can drag URLs, similar to a conventional bookmarking system but using thumbnails to represent pages and employing a strict chronological ordering. WebView [4] provides a thumbnail-based web history that can be ordered chronologically or hierarchically by site. Flagging of useful pages both implicitly and explicitly is integrated directly into the history system by showing “dog-ear” marks on page thumbnails, representing either an explicitly-marked page or a frequently-visited page. The Hunter Gatherer system [22] allows users to highlight components of a Web page and, using a shortcut key, place those components into a persistent notebook. Dontcheva *et al.* [7] present a system for summarizing personal web browsing sessions, which presents a richer summary than that provided by SearchBar, but at the cost of more manual intervention and annotation. Jones *et al.* [12] explore the needs of users managing complex projects that span multiple applications, including Web search, and introduce the Project Planner system, which allows users to organize a variety of task-related content.

The SearchPad [2] system was designed as a lightweight notebook that assisted users in keeping track of a search’s progress, and is similar to SearchBar in that its organization is centered around queries. SearchPad supported passive storage of queries, and it modified search engine result

pages to allow active storage of search results for later review (by adding a “mark” button to all search results). SearchPad could optionally be docked into a Web browser to allow integration with search results, but did not persist when a user navigated away from a search results page.

The S³ [17] and SearchTogether [16] systems provide dedicated interfaces for conducting Web searches and record subsequent navigations into a persistent history that can be shared among users. Previous work (e.g. [8],[13],[15]) has explored alternative visual representations for Web history, but those systems have not used search queries as the primary organizational system for Web history. Like SearchBar, these tools all aim to assist users with informational search tasks; however, SearchBar differs from prior art in that it is a search-centric history mechanism that aims to assist users in regaining the context of multi-session investigations by providing both proactively gathered information (queries and subsequent pages visited) as well as optional explicit user annotations (topic designations, thumbs-up ratings, and notes).

CONCLUSION

The primary contributions of this work are threefold:

1) We presented results from a survey on Web search habits, which found that multi-session investigations are common. Along with the results from this survey, we presented a taxonomy of strategies presently used to maintain and restore context for exploratory searches.

2) We presented SearchBar, a novel tool for organizing Web history that supports re-acquisition of search context after interruptions. SearchBar uses search topics and search queries as its fundamental organizational metaphors to organize the Web pages subsequently browsed. An evaluation of this tool demonstrated that first-time users are able to effectively use this interface to complete a complex search task and make use of its affordances for re-finding information obtained via Web search.

3) We presented additional study results that describe the mechanisms used by Web searchers in general to re-find information and manage searches across interruptions, which will inform future work on SearchBar and related work on supporting Web search and multi-tasking.

Future work will include a broader deployment of SearchBar to explore the use of this tool over longer periods of time and within a larger user population. Future development of SearchBar will focus on automating the creation of search topics and delineating tasks based on user actions such as tab creation and initiation of new queries.

REFERENCES

1. Aula, A., Jhaveri, N., and Käki, M. Information Search and Re-access Strategies of Experienced Web Users. *WWW 2005*.
2. Bharat, K. SearchPad: Explicit Capture of Search Context to Support Web Search. *WWW 2000*, 493-501.
3. Broder, A. A Taxonomy of Web Search. *ACM SIGIR Forum*, 36(2), 2002, 3-10.
4. Cockburn, A., Greenberg, S., McKenzie, B., Smith, M. and Kaasten, S. WebView: A Graphical Aid for Revisiting Web Pages. *OZCHI 1999*.
5. Cockburn, A. and McKenzie, B. What Do Web Users Do? An Empirical Analysis of Web Use. *IJHCS*, 2001, 54(6), 903-922.
6. Czerwinski, M., Horvitz, E., and Wilhite, S. A Diary Study of Task Switching and Interruptions. *CHI 2004*, 175-182.
7. Dontcheva, M., Drucker, S., Wade, G., Salesin, D., Cohen, M. Summarizing Personal Web Browsing Sessions. *UIST 2006*.
8. Hightower, R., Ring, L., Helfman, J., Bederson, B., and Hollan, J. Graphical Multiscale Web Histories: A Study of Padprints. *Hypertext 1998*, 58-65.
9. Iqbal, S., Horvitz, E. Disruption and Recovery of Computing Tasks: Field Study, Analysis, and Directions. *CHI 2007*.
10. Jhaveri, N. and Rähkä, K. 2005. The advantages of a cross-session web workspace. *CHI 2005 Extended Abstracts*.
11. Jones, W., Bruce, H., and Dumais, S. How Do People Get Back to Information on the Web? How Can They Do it Better? *Interact 2003*, 793-796.
12. Jones, W., Bruce, H., Foxley, A., Munat, C. Planning Personal Projects and Organizing Personal Information. *ASIST 2006*.
13. Kaasten, S. and Greenberg, S. Integrating Back, History, and Bookmarks in Web Browsers. *CHI 2001 Ext. Abs.*, 379-380.
14. Kellar, M., Watters, C., and Inkpen, K. An Exploration of Web-Based Monitoring: Implications for Design. *CHI 2007*.
15. Milic-Frayling, N., Sommerer, R., and Rodden, K. WebScout: Support for Revisitation of Web Pages within a Navigation Session. *WI 2003*, 689-693.
16. Morris, M.R. and Horvitz, E. SearchTogether: An Interface for Collaborative Web Search. *UIST 2007*, 3-12.
17. Morris, M.R. and Horvitz, E. S³: Storable, Shareable Search. *Interact 2007*, 120-123.
18. O’Connell, B. and Frohlich, D. Timespace in the Workplace: Dealing with Interruptions. *CHI 1995 Ext. Abs.*, 262-263.
19. Obendorf, H., Weinreich, H., Herder, E., and Mayer, M. Web Page Revisitation Revisited: Implications of a Long-Term Click-Stream Study of Browser Usage. *CHI 2007*, 597-606.
20. Rainie, L. and Shermak, J. *Search Engine Use, November 2005*. Pew Internet and American Life Report.
21. Rose, D. and Levinson, D. Understanding User Goals in Web Search. *WWW 2004*, 13-19.
22. schraefel, m.c., Zhu, Y., Modjeska, D., Wigdor, D., and Zhao, S. Hunger Gatherer: Interaction Support for the Creation and Management of Within-Web-Page Collections. *WWW 2002*.
23. Sellen, A., Murphy, R., and Shaw, K. How Knowledge Workers Use the Web. *CHI 2002*, 227-234.
24. Tauscher, L. and Greenberg, S. Revisitation Patterns in World Wide Web Navigation. *CHI 1997*, 399-406.
25. Teevan, J., Adar, E., Jones, R., and Potts, M. Information Re-Retrieval: Repeat Queries in Yahoo’s Logs. *SIGIR 2007*.
26. Teevan, J. The Re:Search Engine: Simultaneous Support for Finding and Re-Finding. *UIST 2007*, 23-32.
27. Weinreich, H., Obendorf, H., Herder, E., and Mayer, M. Off the Beaten Tracks: Exploring Three Aspects of Web Navigation. *WWW 2006*, 133-142.
28. Wen, K. Post-Valued Recall Web Pages: User Disorientation Hits the Big Time. *IT & Society*, Vol 1, No 3 (2003), 184-194.
29. White, R., Kules, B., Drucker, S., and schraefel, m.c. Supporting Exploratory Search. *CACM*, 49(4), 2006.