

Social Navigation for Loosely-Coupled Information Seeking in Tightly-Knit Groups using WebWear

Scott Bateman^{†‡}, Carl Gutwin[†], Gordon McCalla[†]

[†]Department of Computer Science
University of Saskatchewan
Saskatoon, Saskatchewan

sbateman@upei.ca, gutwin@cs.usask.ca, mccalla@cs.usask.ca

[‡]Department of Computer Science & IT
University of Prince Edward Island
Charlottetown, Prince Edward Island

ABSTRACT

Many web-based information-seeking tasks are set in a social context where other people’s knowledge and advice improves success in finding information. However, when tightly-knit contacts (friends, family, colleagues) are not available, information seeking becomes more difficult. Inspired by previous work in social navigation, we developed WebWear, a system that collects and displays traces of activity for tightly-knit groups. WebWear allows people to use contextual knowledge of contacts’ interests and activities to interpret the meaning of the traces, improving their usefulness. In a comparative study, we found that WebWear helped people complete information-seeking tasks more accurately, without requiring additional effort. A one-week field trial found that WebWear was both usable and useful, and that privacy concerns were reduced in the small-group context. WebWear shows that small-scale social navigation systems are feasible, and that they can improve the effectiveness of information seeking on the World-Wide Web.

Keywords

Social navigation; collaborative information seeking

ACM Classification Keywords

H.5.3 [Group and Organization Interfaces]: CSCW

INTRODUCTION

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

In these situations, we communicate with tightly-knit groups of contacts (friends, family, or colleagues) to seek advice and obtain the knowledge we need in order to most effectively complete our tasks. This knowledge is not

always essential (that is, the work is loosely-coupled and can often continue even without the other person’s input [12]), but, when available, it can provide substantial benefits in efficiency and effectiveness. Tightly-knit groups, then, are groups of contacts who know one another and share common information seeking interests. The contacts in these groups are willing to share information to help one another accomplish their tasks with more success.

However, since people have few tight-knit contacts, they are not always available to provide information when needed – people may be physically away, or too busy to interrupt. In these situations, we lose the benefit of others’ experience and guidance. The unavailability of contacts leads to collaborative work being less efficient or effective. In some cases this causes annoyance and inconvenience (e.g., “I can’t find the article John was telling me about, I will have to ask him about it tomorrow”); in other cases it can lead to more severe problems and task failures (e.g., “I can’t find the online form Sally told me to complete before tonight’s deadline, and her flight just left”).

Although CSCW researchers have studied several aspects of web-based collaboration, existing solutions are still insufficient to solve the problems described here. First, collaborative information seeking systems (e.g., [10,18,24]) are traditionally oriented towards more tightly-coupled and focused collaboration, and the explicit-collaboration approach often does not match the loose coupling of web tasks [1]. Second, social navigation systems (e.g., [6,25,26]) traditionally take a view that is much broader than that of a tightly-knit group, leading to information that is difficult to interpret because it cannot be contextualized through personal knowledge of the provider.

We have developed a third approach – called WebWear – that fits between these existing paths: WebWear focuses on support for loosely-coupled information-seeking tasks in which a tightly-knit group of contacts can provide valuable knowledge and advice. WebWear uses a basic idea of social navigation – collecting traces of activity from other users – but only gathers information from people within the tightly-knit group, allowing the viewer to better interpret the traces based on knowledge of the people in the group, their preferences and habits, and their current activities.

We have developed a working WebWear system that gathers, stores, and visualizes traces of shared web activity (Figure 1). WebWear works in a standard browser and uses

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.

CSCW '13, February 23–27, 2013, San Antonio, Texas, USA.

Copyright 2013 ACM 978-1-4503-1331-5/13/02...\$15.00.

passive methods for gathering and displaying information – which can better support loosely-coupled information-seeking tasks. To test the effectiveness of this approach, we conducted a user study and an initial field trial. In the user study, participants carried out realistic information-seeking tasks in a hypothetical small-workgroup setting. We compared WebWear, a broad social-navigation display, and a standard browser. We found that WebWear led to significant improvements in task success, user confidence, and preference, without requiring additional effort. These results show the potential benefit that can be gained when social-navigation systems are built on a smaller scale: when information comes from known contacts, navigation traces become more interpretable and more useful.

We also deployed WebWear in a small field trial with sixteen users, to test real-world usefulness and to explore issues of overlap and privacy. The study showed that there were numerous times that people encountered others' traces, and that there were several situations where this information was valuable to the user. Although participants had privacy concerns, these did not prevent them from actively using WebWear, and users reported diminished concerns due to the workgroup context.

Our work provides three main contributions. First, we present a new kind of social navigation system that is designed specifically for tightly-knit groups. Second, we show that the contextualization possible in this approach improves web-based information-seeking tasks. Third, we demonstrate a working and freely-available WebWear system, and provide initial evidence that the approach can be deployed successfully in the real world.

RELATED WORK

Information Seeking Tasks

Information seeking is “a process in which humans purposefully engage, in order to change their state of knowledge ([16], p.5)”. Pirolli and Card identified two loops that take place during information seeking [23]: a *foraging* loop where useful information is identified; and, a *sensemaking* loop where the identified information is consolidated with current knowledge. The concept of *information scent* describes the “(imperfect) perception of the value, cost, or access path of information sources obtained from proximal cues, such as WWW links [3].” Thus, the information scent for alternative paths on the Web can be improved by providing informative proximal cues, which in turn improves information access [26].

Loosely-Coupled Collaboration

Kraut and colleagues observed that most workplace interactions (81%) occur in an informal, unplanned, and opportunistic fashion [15]. Similar patterns have been observed in collocated information-seeking; e.g., in one study, 62% of collaborations occurred without pre-planning [1]. These kinds of collaborations often involve *loosely-coupled work* [15]. Coupling describes the amount of work that people can do individually before they require communication or actions from another person [12]. In

loosely-coupled work, communication with collaborators must be low-cost, because the collaboration is often secondary to the underlying task; if the cost of communication is too high, people will simply not engage in collaboration [1,2,15,22].

It has been observed in loosely-coupled situations that traces of work left in the physical environment are a preferred channel of communication [2,22]. Traces in shared documents or work artifacts allow collaborators to communicate indirectly about the work they have undertaken [22], cost little to gather (since they are stored as a result of performing other primary tasks), and are easy for others to discover and retrieve [22]. Researchers have suggested that digital environments should, therefore, represent traces of work, since these are an important part of collaboration in physical environments [2].

Based on this previous work, we describe *loosely-coupled information seeking tasks* as tasks in which other people can provide advice or guidance, but where direct communication between an information seeker and a guide may not be practical or possible.

Social Search and Collaborative Information Seeking

Social search encompasses any scenario where people work together to seek information on the Web [4,8,10,18]. Social interaction in web search is a common occurrence [4,8,18]; for example, a recent survey found that 59% of search respondents shared the results of their most recent search [8]. Even when search is solitary, the underlying task can involve communication and collaboration with others [4, 8]: before searching – to help describe and focus what is sought; during search – to receive guidance and advice on avenues to pursue; and after search – to share what has been found. However, most existing search systems do not support any form of collaboration [10,16]. This leads people to adopt workarounds [10]; e.g., exchanging URLs by email [1].

Research in *collaborative information seeking* addresses situations where users have “explicitly-shared information needs” ([10], p.47), and actively work together to fulfill these needs. Work in this area has produced systems that support different forms of collaboration (see [18] for a review), including: collocated, synchronous distributed, and asynchronous distributed collaborations. This research has produced several successful systems (e.g. [19,20,24]); however, the focus on “explicitly-shared information needs”, means that they are oriented towards pre-planned work and more tightly-coupled interaction. As described above, however, most collaboration occurs in an unplanned way, and may be better supported by lower-cost solutions that fit into existing individual tools and practices.

Computational Wear

Hill and colleagues [13] described the idea of *computational wear* – capturing “on computational objects the events that comprise their use” ([13], p.3). They developed tools that visualize traces of reading and editing history for people working on a shared document. Their

goal was to allow people to easily answer questions about who wrote and edited different parts of the document [13].

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

Social Navigation

Social navigation describes how people can be guided through an information space based on the direct advice or interpreted actions of others [7]. The idea of indirect social navigation from interpreted behavior (e.g., computational wear) led to Web systems such as Footprints [25], CoWeb [6], or Scented Widgets [26] (which show traces in a user interface). These systems provide visualizations of others' use of paths and links – they augment the information scent in the environment, providing information about social activity that allows users to make inferences about the choices made by others. Studies showed that these systems can reduce the navigation steps needed [25] and can increase the number of alternatives explored [26].

However, these social navigation systems typically aggregate the behavior of all users [9], limiting what can be inferred from the traces and what questions can be answered. For example, a system showing the navigation behavior of a large group can indicate general visitation patterns and popularity [25], but is less useful for specific questions that are set into a local social context.

Some social navigation systems that offer narrower support are in active use on several social networking websites (e.g., Facebook, Twitter, or Academia.edu), and on search engines (e.g., Bing or Google). These systems provide indirect social navigation from the traces of a seeker's social networking contacts. For example, Facebook displays contacts' reading behavior on particular news sites, and Google displays social annotations from Google+ next to search results. These systems narrow social navigation support, by providing details of *who* and *when*. These systems are limited in scope, however, because they work only for certain types of behavior and on certain websites. For example, Facebook only displays friends' history for certain sites (e.g., Yahoo! News). This means the scope of the activity does not extend to most information-seeking tasks, which can occur over many different websites. Whether current narrow systems that have a limited scope can provide support in typical information-seeking tasks is unclear.

DESIGNING WEBWEAR

The goal of WebWear is to support loosely-coupled Web information seeking that occurs within a social context, by providing narrow social navigation support throughout information seeking tasks. Here we describe four aspects of the design: tasks, information, gathering, and display.

What tasks should be supported?

Previous research has identified several kinds of collaborative information seeking that are relevant to loosely-coupled scenarios:

- *Following*: purposefully repeating someone else's actions [6] to find an object of shared interest.
- *Division of labor*: dividing up work so that there is reduced duplication of effort [18].
- *Seeding*: traces left purposefully for discovery by a collaborator (e.g., [22]).
- *Serendipity*: unexpected encounters with a trace that provides useful information [25].

Direct communication tasks are also a common part of loosely-coupled work; however, in WebWear we focus on implicit information rather than explicit communication.

What information do people need in these tasks?

Much of what people need for loosely-coupled information tasks has been examined in research on group awareness (e.g., [12]). Providing the appropriate level of awareness in collaborative information seeking has been shown to facilitate collaboration [24]; as described above, current social navigation systems may provide information at too high a level for tasks that exist in a social context.

Therefore, supporting information-seeking in tightly-knit groups involves showing traces from the group, at all stages of a person's information seeking activities, including starting points such as search engines. Further, support can be narrowed by providing awareness information that relates to activity from the past [11], including *what* activities were carried out, *who* performed them, *where* on the web they went, and *when* the activity happened.

What can be gathered from web interactions?

Browsers can be instrumented to capture a wide variety of user behavior, from details of scrolling behavior to the pages users bookmark (see [14] for a review). However, the use of most web browser functionality – including how bookmarking functionality is used, how many tabs are used, or if a page has been printed – is highly idiosyncratic.

We therefore focused our attention on the basic unit of interaction on the Web – the webpage visit. Regardless of how a user makes use of their browser, everyone must visit pages in order to find and gather information. Therefore, we limited our data to what can be collected from time-stamped URLs with a user identifier. This simple data can provide rich awareness of the information seeking process, including: what pages were visited, who visited the pages, when a visit occurred, how long a visit lasted, and the content of search queries (from the query string of a URL). User-specific URLs are, of course, a potential privacy concern; we discuss this issue further later in the paper.

How can the information be displayed?

Design Constraints

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

coupled collaboration. Thus, WebWear must work within established interaction patterns based on webpages, web browsers, and common tools such as search engines.

Design Factors

Location: External versus Internal to the page. Previous work on collaborative information seeking has provided support widgets that display or summarize participants’ activities (e.g., [19,20,24]), which allows a large amount of detailed information to be displayed (e.g., using a sidebar). In contrast, work on information scent often augments webpages to provide proximal cues (e.g., [6,21]), allowing users to see additional information scent within the page. However, the internal approach means that only limited information can be displayed without obscuring webpage content. We chose an internal approach using graphical representations to allow better integration with existing tools such as search engines; in addition, when no trace information is available, users can fall back to using webpages as they normally would.

Style of Augmentation: Artifact Encoding vs. Glyph. Graphical objects (glyphs) can be placed near the artifacts of interest (i.e., pages and links) [6], or information can be encoded into the representation of the artifacts themselves (e.g., changing the font size of links [4]). We decided to use glyphs to encode the information: although glyphs can occlude information on the page, they can be consistently displayed without altering a page’s style and functionality.

Information to Display. Information seeking on the Web often involves a combination of page visits, search engine queries, and decisions about which search results to visit. We wanted to display only information relevant to making choices in this process [26]; therefore, we include information about the page currently visited, the links that can be chosen on the page, and, when search is available, information that supports query formulation. We show two basic elements that are initially visible in the glyphs: whether or not a contact has visited a particular page, and an indicator of interest in that page (a review indicators of interest from interaction data can be found in [14]).

WebWear

The WebWear system is an add-on for the Chrome browser that collects and adds trace visualizations to web pages (see Figure 1, source code: <https://github.com/scottbateman/WebWear>). WebWear provides glyphs that encode the number of contacts who visited a page or a link, and the recency of the visit. If no glyph is present, it means that none of the users’ contacts have visited the page.

Page-Level Glyphs. The basic glyph used in WebWear is a small colored bar. Bars representing visits to the current page are colored purple. The fill of the bar represents the number of unique contacts who visited a page (see Figure 2). The bars fade away over time to indicate the recency of the last visit. The bar is darkest if the visit occurred in the last few days, and fades to its lightest if the visit was over a month old. Glyphs are 8x10-pixel marks for page-level visitation and are placed at the right of each link (Figure 1

and Figure 2). A larger version of the glyph (300x15 pixels) is shown at the top of the page to indicate visitation to the page currently being viewed (Figure 1).

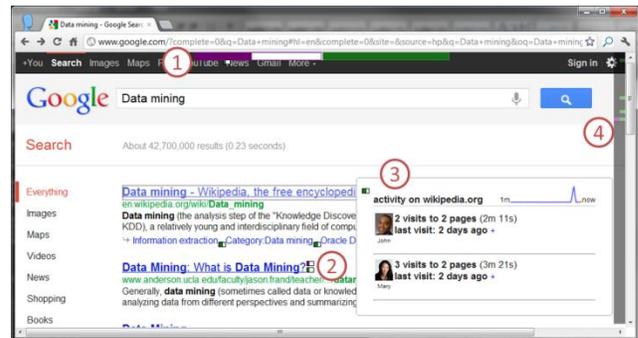


Figure 1. WebWear displayed on a search results page. 1) Glyphs representing visitation to the current page (purple) and website (green). 2) Activity glyphs on a link. 3) Detail popup after hovering on a glyph. 4) Scrollbar overview showing glyph locations on the page.



Figure 2. Left: glyphs indicating the number of unique visitors to a link in WebWear. Right: glyphs as displayed in the BroadNav condition of the study.

Website-Level Glyphs. To increase the likelihood of encountering useful information, WebWear provides a second indicator that aggregates the activity on all pages visited on a given website. Website-level aggregation was chosen because visits to a particular website are likely to have some tasks or information needs in common. We used the same visual encoding as the page-level glyphs, but colored these bars green (Figure 1).

Details-on-Demand. All glyphs are interactive, and can provide details about the represented activity. Mousing over a glyph produces a popup that shows a recency-ordered list of visitors, and an overview of their visitation for the page or website. The popup displays the total number visits on the website, the number of pages visited, the total amount of time spent, and when the last visit occurred. A list of the occurrence and duration of each individual visit can also be displayed. Popups relating to website visits show a recency-based list of pages visited on that site. Further, a sparkline visualization in the popup allows users to get a sense of the activity of all contacts’ visits to a page or website in the last month (see Figure 3).

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

Contact Filtering. Users can filter the traces by selecting and de-selecting contacts in the WebWear control menu located on the browser toolbar.



Figure 3. Details popup for activity on wikipedia.org, showing details of all contacts' visits to the website.

Group-based Query Completion. For any page that uses a standards-based query field (i.e., with an HTML name attribute value of “q”), WebWear also provides a recency-based query completion list based on the query history of contacts (Figure 4). The group-query completion overrides existing query completion. Clicking on the search field displays the most recent queries executed by all contacts, providing a simple way to see the activity of all users. Letters typed in the field updates the filter. The query completion works effectively on Google, Bing, and Yahoo, and also in site-specific search fields for sites such as stackoverflow.com. Group search activity is shared across all search fields, regardless of the site where it came from.



Figure 4. Group-based query completion displayed on Bing; a recency-based based on the queries of contacts.

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

EMPIRICAL USER STUDY OF WEBWEAR

To evaluate our primary hypothesis – that smaller-scale social navigation information can improve information seeking in tightly-knit groups – we carried out a user study

that asked participants to perform web-based information tasks in a variety of different conditions. The study was designed to answer two main questions:

- Whether the small-group interaction histories in WebWear improve task performance compared to either no history, or broad social-navigation information; and,
- Whether the WebWear visualizations (as described above) are easy to access and interpret for users, when used in a real web browser and a realistic task context.

Study Methods

Interface Conditions

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

- *Normal browser* was an unchanged version of the Chrome web browser (version 19.0.1084.52).
- *Broad social navigation (BroadNav)* is a visualization of link popularity that was similar to the simple bar visualizations used in WebWear, and appeared on all pages (Figure 2). Popularity information was simulated with Google PageRank data (obtained by querying the Google Toolbar webservice for each link). PageRank approximates the likelihood of randomly visiting a page. While it is only weakly correlated with real world page visitation [17], as an approximation it provides a reasonable stand-in for true link visitation counts.
- *WebWear* was implemented as described above (see Figure 1). Information traces were pre-determined for the tasks used in the study; traces were pre-recorded and made visible for three members of the participant's task workgroup (described below).

Tasks and Scenarios

For the experiment, participants were asked to imagine working with a hypothetical group, which included: Mary, a close friend to the participant; John, a friendly colleague and a computer programmer; and, Lynn, the group's manager and mentor. Participants carried out two kinds of tasks: *follow* and *divide* tasks. These two task types (as introduced above) were selected for our evaluation because both have been a focus of previous work, are extremely common in real-world workgroup collaborations, and can occur in planned or unplanned scenarios. First, *follow* tasks involved following up on a lead or comment provided by a member of the person's scenario workgroup, and finding a specific article or site that had been mentioned or discussed by the other person. In follow tasks, we were interested in whether participants could successfully find specific items that had been partially identified by another person. For example, one of the follow task descriptions was:

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

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

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

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

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

Participants and Apparatus

Twenty-four participants (14 women, 10 men, mean age=26.1 years) were recruited from a local university. Eighteen were students from a range of disciplines, others were researchers (5) or software engineers (1); none were familiar with the research project. Participants all reported searching the Web at least once a day and spending at least 2 hours/day using a computer, during a normal workday.

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

Design and Procedure

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

Participants first completed a demographics survey, and were given a handout that introduced them to the hypothetical scenario and collaborators. They were then given a practice task using the normal browser. They next completed one of the three interface condition blocks

(balanced with a Latin square), which included both a *follow* and a *divide* task (alternating order between participants). Before each of the *BroadNav* and *WebWear* interface blocks, participants were given a short tutorial (approximately 5 minutes), where they were introduced to the new interface and asked to perform a random query to show how they system displayed information. They were then given an additional practice task. The experimenter interacted with the participant during the practice tasks and pointed out the displayed information in the context of the task. Participants then completed the two experimental tasks. After each task they completed a short questionnaire concerning their experiences, and after the *WebWear* and *BroadNav* interfaces they completed a second questionnaire that asked their opinions of the interface. At the end of the session, participants were asked about their preferences.

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

Data Analysis

Performance data was analyzed using RM-ANOVA, and post-hoc tests used Bonneferoni corrections for multiple comparisons. Because the accuracy score was a binary value, we initially transformed accuracy scores using the Aligned Rank Transform [27]. Subjective responses were analyzed using Friedman's test, and post hoc comparisons were done using the Wilcoxon signed-rank test. However, because Friedman's is a one-way test only, we split our task-related questionnaire data by task type, and looked for differences between interfaces under the different task conditions. We did not look at the differences in task type data, as we were mainly interested in participants' subjective responses to interface, not task type.

Results

We present our results organized by our main hypotheses about success and effort in completing the tasks. We first present performance data and subjective responses, and then report on experiences and preferences.

Were participants more successful with WebWear?

Participants completed their work more successfully when using *WebWear* (see Figure 5). There was a main effect of Interface on accuracy score ($F_{2,46}=13.76, p<.001$) and task type ($F_{1,23}=6.29, p<.05$), but there was no interaction effect between interface and task type ($F_{2,46}=2.11, p>.05$). Pairwise comparisons show that participants were significantly more accurate in completing the task with *WebWear* than with either *BroadNav* ($p<.005$) or the normal browser ($p<.001$). There was no difference in accuracy between *BroadNav* and normal ($p>.05$). Participants were significantly more accurate in completing the divide tasks than the follow tasks ($p<.05$).

Participants also felt more certain about having successfully completed the task when using WebWear. In follow tasks, participants felt more certain about having found the correct webpage (Figure 5). In divide tasks, participants felt more certain about having not duplicated any answers (Figure 5). There was a significant effect of Interface on certainty about divide tasks ($\chi^2(2)=39.78, p<.001$). Pairwise comparisons show that participants were more certain with WebWear than with BroadNav ($z=-2.94, p<.005$) and normal ($z=-3.63, p<.001$). Participants were also significantly more certain with BroadNav than normal ($z=-2.27, p<.05$). There was a significant effect of Interface on certainty for follow tasks ($\chi^2(2)=18.25, p<.001$); pairwise comparisons show that participants were significantly more certain with WebWear than BroadNav ($z=-4.07, p<.001$) or normal ($z=-3.32, p<.001$), and more confident with BroadNav than normal ($z=-2.78, p\leq.005$).

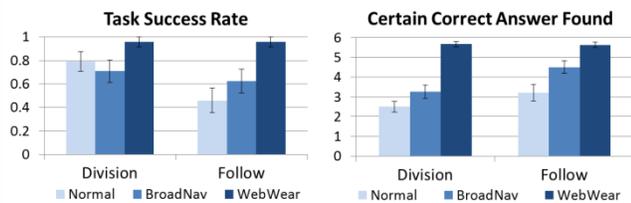


Figure 5. (Left) Task success rate. (Right) Subjective certainty about task success.

Did participants require less effort when using WebWear?
To consider efficiency, we looked at the time taken to complete the task, subjective perceptions of efficiency, and subjective need for communication with a collaborator.

Completion times are shown in Figure 6. There was a main effect of Interface on completion time ($F_{2,46}=20.89, p<.001$). Pairwise comparisons show that BroadNav was slower than both WebWear ($p<.001$) and normal ($p<.001$), but that there was no difference observed between WebWear and normal ($p>.05$). There was no effect of task type on completion time ($F_{1,23}=.32, p>.05$), with divide tasks ($mean=2.9$ min., $SEM=.19$) taking roughly the same amount of time as follow tasks ($mean=2.7$ min., $SEM=.35$). There was no interaction effect observed ($F_{2,46}=.56, p>.05$).

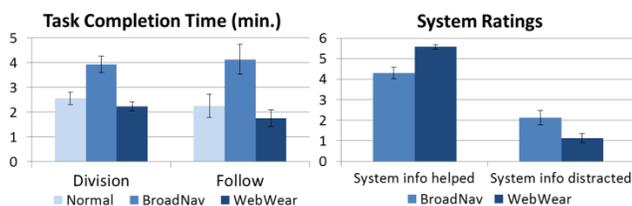


Figure 6. (Left) Average task completion time. (Right) Participant ratings of the system being helpful and distracting.

Participants felt they were most efficient when using WebWear (see Figure 7). There was a significant effect of Interface for both divide tasks ($\chi^2(2)= 10.84, p<.005$) and for follow tasks ($\chi^2(2)= 13.92, p<.001$). For divide tasks, participants felt they were more efficient with WebWear than both BroadNav ($z=-2.65, p<.01$) and normal ($z=-2.81,$

$p\leq.005$). There were no differences between Normal and BroadNav ($z=-.84, p>.05$). For follow tasks, participants again felt they were more efficient with WebWear than both BroadNav ($z=-3.39, p\leq.001$) and normal ($z=-3.21, p\leq.001$), but there was no observed difference between Normal and BroadNav ($z=-.24, p>.05$).

Participants felt less need for communication when using WebWear (see Figure 7). There was an effect of Interface for both divide tasks ($\chi^2(2)= 25.42, p<.001$), and for follow tasks ($\chi^2(2)=11.17, p<.005$). For division tasks, the perceived need to communicate was lower for WebWear than both BroadNav ($z=-3.31, p\leq.001$) and normal ($z=-4.01, p<.001$). Participants also felt they needed to communicate less with BroadNav than normal ($z=-2.07, p<.05$). For follow tasks, again the perceived need to communicate was lower for WebWear than both BroadNav ($z=-2.86, p<.005$) and normal ($z=-3.84, p<.001$), but there was no difference between BroadNav and normal ($z=-.68, p>.05$).

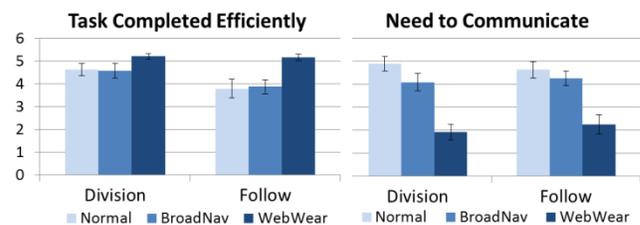


Figure 7. (Left) Participant agreement with having completed the task efficiently. (Right) Ratings of needing to communicate to be successful in the task.

Perceived Usefulness of WebWear and BroadNav
Participants felt that the information available in both WebWear and BroadNav was useful in completing the tasks (Figure 6), but they rated the information in WebWear as being significantly more helpful ($z=-3.447, p<.001$). Participants also felt that the information displays in both BroadNav and WebWear did not get in the way or distract participants from their work (see Figure 6); however, they felt that WebWear was significantly less distracting than BroadNav ($z=-2.55, p<.05$). Finally, 22 of the 24 participants felt that WebWear was the best tool for the type of tasks seen in the study, and 22 agreed that they would use WebWear for information seeking activities ($mean\ agreement=5.13, SEM=0.26, max=6, min=0$).

FIELD TRIAL

To better understand issues of WebWear’s use in the real world, we also held an initial one-week field trial within our research lab. We were interested in three main issues:

- would people’s information-seeking activity overlap, providing opportunities for WebWear to be used;
- would WebWear lead to useful experiences;
- what are peoples’ impressions of, and concerns with, sharing activities during their real-world searches.

Sixteen members of our lab (1 professor, 1 visiting researcher, 10 graduate students, 4 undergraduates; none of them authors) installed the WebWear extension and used it on their lab computers for one week. Typical work

activities for this group cover a wide range of HCI and computer science topics. Trial participants were given a 10-minute tutorial about how WebWear worked; none had previously used or seen it. Participants completed an optional survey four times during the week, which asked about their experiences. At the end of the week, we held an hour-long group interview. The deployed WebWear system was the same as described above, but additionally provided a facility for participants to view and permanently delete any of their own traces from the system.

Participants were relatively active during the week of the trial, sharing 6140 page visits and 1022 searches. Participants reported that encountering others' traces was frequent, and 88% of the surveys report exploring other users' activities. Three participants even reported changing their search behavior so that they could encounter traces more often – rather than searching by typing a query in to the address bar in Chrome as they normally do, they would first visit the Google homepage, so that they could see WebWear's group-query completion (Figure 4).

Trial participants found WebWear useful in several ways. First, many participants reported that WebWear was useful for maintaining an awareness of the activities of others. For example, "While looking for tutorials on ... XNA I did notice, thanks to WebWear, that [P12] had also been searching for this topic. I didn't even realize that [P12] was using XNA in his project. Now I guess if I run into anything major, I at least know someone who may have also been dealing with the same topic."

Further, even in this brief trial, we counted three instances of unplanned follow or divide collaborations. An example of a follow task: "[P15] had visited Dexsoft, a 3D model website, and had discussed with me the category of models that she would like me to look at. When I visited the Dexsoft website I hardly had to look around the website to find the models she was talking about because I just followed her WebWear 'footprints'."

An example of a division task arose when P8 asked P10 if he knew how to solve an issue with a Bluetooth dongle. P10 did not have an answer and returned to his desk, deciding only later to search for an answer to try to help P8. As P10 started his search he saw P8's queries, and selected different search terms to explore a different path, now knowing that she had already searched without success. These examples illustrate real-world cases of the tasks used in our user study, and also provide evidence that loosely-coupled tasks can be supported through shared traces.

Finally, all field-trial participants agreed that there was potential value in being able to share traces both to maintain awareness and to help coordinate tasks. In the group discussion, many people still reported being wary of privacy despite having mainly positive experiences during the trial. The privacy concerns related mostly to the fear of unwanted activity accidentally being entered into the browser, or other participants misinterpreting the context of search activity. Not all participants experienced these

concerns, though, and some reported forgetting about the data collection because they were not overly worried.

All participants agreed that the work context was suitable for sharing traces, and while some concerns with sharing existed, no one reported having a bad experience in using WebWear, and were open to using it in the future.

DISCUSSION

The empirical user study provides three main results:

1. People were significantly more accurate in finding answers with WebWear than with either a normal browser or with a broad social-navigation system, and were significantly more confident in their answers.
2. The increased accuracy did not come at the cost of additional effort: tasks were faster in WebWear than in BroadNav and no slower than the normal browser. Participants also reported a reduced need to communicate than with the other systems.
3. People subjectively felt that they were more efficient with WebWear, that the information it presented was helpful and did not distract from their tasks, and that they would like to use it in appropriate contexts.

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

4. There was enough overlap in people's real-world browsing that participants were able to make use of WebWear on numerous occasions over one week;
5. Participants reported several cases where they found WebWear information useful, both for specific loosely-coupled information-seeking tasks, and for maintaining awareness of others in the workgroup;
6. All sixteen participants indicated that while privacy from passive tracking remained a concern, using WebWear did not unduly intrude on their privacy, and there were benefits to sharing in a work context.

In the following sections, we provide explanations for these main results, and then discuss issues surrounding the real-world deployment of WebWear.

Explanations for Main Results

The results observed in our user study support our main hypothesis – that trace information is more useful in loosely-coupled collaborative tasks when contextualized to a tightly-knit group. WebWear combines lessons from social navigation and collaborative information seeking, and our studies show that this combination provides good support for loosely-coupled information-seeking tasks.

Why did WebWear work in the user study?

WebWear provided three main tools that led to the successful results in our user study. First, group-query completion allowed users to see if their search terms aligned with those of their contacts. This was useful in both follow and divide tasks. In follow tasks, users could repeat queries to match the results of a previous collaborator; in divide tasks, they could formulate unique queries to get results that had not been seen by their collaborators. Second, the activity glyphs next to links allowed users to quickly identify pages that they should visit or avoid.

Third, the popup details on the glyphs provided appropriate information that allowed users to confirm if an activity was of interest or not. Most often, simply seeing *who* had left the trace was all that was needed (and the user portraits available in the system gave this information at a glance). Further, the details for websites (Figure 3) provided a useful way to quickly find a known target. For example, to complete the “bird book” task (described above) a user could simply visit amazon.com, look at the website popup, and see the list of pages John had visited on Amazon. This form of contextual history provides a natural way for people to quickly find history of interest.

Were the user study tasks biased towards WebWear?

The specific task types examined in this work (loosely-coupled information seeking tasks) have not been previously examined, and show the limitations of many previous social navigation systems. We selected experimental tasks that could be successfully completed by our participants without information from tightly-knit contacts. These tasks did prove more difficult without WebWear; however, this only means that WebWear successfully supports this type of loosely-coupled task. Our literature review suggests that loosely-coupled information seeking may be a more common use case than explicit forms of collaboration, but exactly how often systems like WebWear may be able to lend support during typical information seeking is an important open question.

Why did broad social navigation not work?

Broad social navigation (BroadNav) was no better than a normal browser for most tasks, and was significantly slower than the other interfaces. Reasons for these results may arise from the issue of introduced cost. While we intended our design of WebWear and BroadNav to have a low cost (small glyphs that were easily ignored), it is likely that participants needed more time to consider them. BroadNav offered information that was too general for our tasks (e.g., participants wanted to find John’s bird book, not the book that was most popular); therefore, time spent considering BroadNav information was not beneficial. This effect could diminish with more training – that is, users could learn to only attend to the glyphs for certain tasks. Attending to the information of WebWear also introduced a cost, but this was offset by the benefit provided.

While we could have designed other tasks that would suit BroadNav (e.g., “Find the most globally popular webpage describing attractions in Vancouver”), the value of general social navigation systems has already been established – and our interests were in showing the value of the approach behind WebWear.

WebWear in the Field Trial

In the field trial, it was clear that the most-frequently used feature was group query completion. Participants would visit the Google main page and click on the query box to see what other participants had recently been searching for. Several participants reported that this feature gave them an increased sense of awareness about others’ activities.

Participants also reported frequently checking the activity traces of other participants when they were encountered. The indicators did not always provide task assistance (there were relatively few shared projects going on at the time), but were still followed out of curiosity. For example, several participants reported looking at one another’s traces on YouTube. These participants reported not having seen anything of major interest, suggesting that shared traces on sites like YouTube, which cover many topics, may be too broad for most encounters.

Did WebWear lead to new information seeking behavior?

In both studies, we found that people adopted behavior that was different from their normal search behavior. First, as reported above, people worked in ways that would allow them to encounter traces as part of the normal search practices. Second, we observed several cases where people would formulate queries in the search box without executing them, to see if they could find someone else’s past query. Finally, we also saw a few people become momentarily confused when activity traces didn’t appear where they were expected on search result pages. In these cases participants seemed to have incorporated expectations of WebWear-style information into their model of information scent on web pages. This problem arose due to inexperience with the system, but highlights the importance of making the source and display of trace information simple and easy to understand.

Generalizing the Results and Deployment Issues

There are four main issues in generalizing our results: overlap, usefulness, tightly-knit groups, and privacy.

Overlap of Traces

The sparseness of traces is an important potential issue for small workgroups – given the size of the Web it could be that traces are encountered too infrequently to be useful. However, two principles provide some indication that there will be reasonable overlap. First, there are a set of common starting points (e.g., search engines) that increase overlap and provide opportunities for shared encounters. WebWear’s group-query completion was designed with this in mind, and the field trial showed that this feature was frequently used, and in the ways that we intended.

Second, a defining characteristic of tightly-knit groups such as workgroups is that they have shared goals, interests, and tasks (even when working in loosely-coupled fashion). These overlaps should translate into people visiting some of the same regions of the web. In our field trial, we saw clear examples: 9 of 16 people had visited stackoverflow.com (a programming help website) within the 1-week trial. This and other examples suggest that there is adequate overlap within many small groups. However, in future work we will test this hypothesis through longer term studies with several different workgroups.

Providing Useful Information

Our work shows that traces can be successfully displayed and used, but there are situations where traces may not provide valuable information. For example, one of the

known challenges of social navigation systems is *herding* (i.e., blindly following others' trails) [9]. Herding is a problem if a trace leads to a poor decision being made.

Further, as previously discussed, there is a cost to interpreting information, so the benefits of making traces easy to encounter and explore must also be considered in terms of the quality of the information. For example, we did not expect a great deal of overlap in our field trial, due to the wide range of activities that people were working on. However, the (green) website glyphs that can appear next to each of the links on a page actually appeared quite often on certain search result pages, which participants did not find overly useful. Systems like WebWear must balance between broadening the possibility of encountering traces and ensuring that the information is still valuable.

There are, however, several design possibilities that can be explored to address this specific issue. For example, a simple approach would be to narrow our website-level traces. Currently, we present website activity for all traces on a top level domain (e.g., google.com); this could be narrowed by including the full domain (e.g., the website would be separated for traces on sites like news.google.com and maps.google.com). However, we have some evidence from our field trial that curiosity leads people to investigate more indicators than necessary (which could partially be an effect of novelty). With more familiarity, and because loosely-coupled collaboration is discretionary, people can decide on a case-by-case basis whether the information provided by WebWear is useful. Further, the interpretation context will help people to determine whether information is valuable or spurious (e.g., "I know Mary is not an expert in biochemistry, so I won't follow her trail in this area").

Privacy

Concerns about privacy have often been the main barrier to the deployment of some awareness systems. Because WebWear passively collects people's interaction history, it could in some cases feel like surveillance. In addition, WebWear attaches people's names to particular actions, meaning that there is no anonymity like there is with broad social navigation. Therefore, systems like WebWear should be used carefully. There are several examples that suggest the technique can work in some cases, however.

First, people have different degrees of willingness to share information about themselves with different kinds of contacts [5]. For example, people are more willing to share with their spouse, friends, and family; therefore, there may be settings and groups of contacts where people are willing to share, even if they would not share with all contacts.

Second, people treat work situations and work browsing very differently from more private settings such as the home. At work (e.g., in an open-plan office) it is easy for people to see what sites are in people's web browsers; therefore, there is a general change in the kinds of browsing that people do at work. This can lead to a greater willingness to have interaction history shared.

Third, people may actually want to share some of their interaction histories. For example, people may wish to highlight a website or demonstrate an expert route through a set of material; examples such as Twitter indicate that some people are quite willing to selectively share their activities with others.

Fourth, it is always possible to switch off the information-gathering aspects of WebWear; this can help to mitigate privacy concerns since people know they will not always be tracked. Further privacy-protection mechanisms could be explored (as discussed below). Last, the ability to edit the recorded information and delete items can help people recover from inadvertent sharing of private activity.

The initial field study provides at least limited evidence that these factors do provide a workable solution to privacy problems. Of the people who initially volunteered, all sixteen completed the one-week trial, indicating that there are people out there who are amenable to this degree of sharing (within a small workgroup). Privacy was a main interest in the group discussion, and some participants did report concerns for their privacy during the field trial. Concerns were not important enough, however, to prevent active sharing by all participants. Further, several participants stated that they do not see their work setting as providing a great deal of privacy anyway, and so they already carry out work on the Web with the expectation that others will be able to see their activities.

Participants were interested in discussing activity sharing during the group interview and felt that logical extensions to WebWear might include whitelists and blacklists to define how personal data should be shared (e.g., I may be willing to share activity on Wikipedia, but keep all my YouTube activity private). However, a few participants were skeptical that such mechanisms could be added without requiring undue effort on the part of a user, and felt that the current on/off mechanism would actually be best. Future work should focus on low-effort and simple sharing mechanisms, because concerns about privacy and decisions about sharing may change on a task-by-task basis, and would likely evolve over time.

Identifying, Maintaining, and Targeting Tightly-Knit Groups

Perhaps the greatest protection for privacy is the targeting of WebWear to tightly-knit groups, groups with mutual trust among members. In our studies we used WebWear with a particular tightly-knit group: a small workgroup. Field trial participants agreed that the use of WebWear in the workplace would be useful and appropriate.

Even though we saw real examples of support provided by WebWear, there were likely situations where participants could have used some guidance, but none was available. With a larger number of contacts available, a user would have a higher likelihood of receiving support when it was needed, by being able to access the experiences of more people. At the same time, in the real world people may be part of several distinct tightly-knit groups, and it seems likely that WebWear users would want to share different

activities with different groups. This raises several questions: How might someone create and maintain separate tightly-knit groups? How could different types of information-seeking activities be targeted to different groups? Should a user see a group's traces if the user is not themselves sharing with the group? What should happen to shared history after a member leaves the group?

In this work, the context of a single tightly-knit group allowed us to first focus on assessing the WebWear approach in general. Future work should aim to better understand the nature of information seeking and information sharing in tightly-knit groups – groups of people that have shared interests, have some level of trust, and are willing to share activities to benefit one another. The findings could lead to answers to the questions raised above, and increase the number of situations where people could benefit from each other's previous work.

CONCLUSION

In this paper we presented WebWear, a system that collects and displays traces of shared activity to support information seeking in tightly-knit groups. We showed that WebWear allows people to use contextual knowledge of contacts' interests and activities to interpret the meaning of shared traces. Through two studies we showed that WebWear can improve accuracy and confidence in completing previously untested loosely-coupled collaborative information-seeking tasks, and provided initial evidence that the approach can be successfully deployed in real-world scenarios. WebWear shows that small-scale social navigation systems are feasible, and that they can improve the effectiveness of information seeking on the World-Wide Web.

ACKNOWLEDGEMENTS

This work was supported by NSERC, the GRAND NCE, and the Walter C. Sumner Memorial Foundation.

REFERENCES

- Amershi, S. & Morris, M.R. Co-located Collaborative Web Search: Understanding Status Quo Practices. Proc. *CHI*, 2009, 3937-3640.
- Bellotti, V., & Bly, S. Walking Away from the Desktop Computer: Distributed Collaboration and Mobility in a Product Design Team, Proc. *CSCW*, 1996, 209-18.
- Card, S., Pirolli, P., Van der Wege, M., Morrison, J. B., Reeder, R.W., Schraedley, P.K., & Boshart, J. Information scent as a driver of web behavior graphs. Proc. *CHI*, 2001, 498–505.
- Chi, E. Information Seeking Can Be Social. *IEEE Computer*, 42(3), March, 2009, 42-46.
- Davis, S. & Gutwin, C. Using relationship to control disclosure in awareness servers. Proc. *GI*, 2005, 145-52.
- Dieberger, A. & Lonnqvist, P. Visualizing Interaction History on a Collaborative Web Server. Proc. *Hypertext* 2000, 221-22.
- Dourish, P. & Chalmers, M. Running out of space: Models of information navigation. Proc. *HCI*, 1994.
- Evans, B.M., & Chi, E. An elaborated model of social search. *IPM*, 46 (6), 2010, 656-78.
- Goecks, J. *Understanding the Social Navigation User Experience*. PhD Dissertation. Georgia Tech. 2009.
- Golovchinsky, G., Qvarfordt, P., & Pickens, J. Collaborative information seeking. *IEEE Computer*, 42(3), 2009, 47-51.
- Gutwin, C. & Greenberg, S. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *CSCW*, (11)3, 2002, 411-446.
- Gutwin, C., Greenberg, S., Blum, R., Dyck, J., Tee, K., & McEwan, G. Supporting Informal Collaboration in Shared-Workspace Groupware. *JUCS*, 14(9), 2008, 1411-34.
- Hill, W.C., Hollan, J.D., Wroblewski, D., McCandless, T. Edit wear and read wear. Proc. *CHI*, 1992, 3-9.
- Kelley, D., & Teevan, J. Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum*, 37(2), 2003, 18-28.
- Kraut, R., Fish, R., Root, R., & Chalfonte, B., Informal Communication in Organizations: Form, Function, and Technology, in R. Baecker, *Readings in Groupware and Computer Supported Cooperative Work*, 1993.
- Marchionini, G. *Information Seeking in Electronic Environments*. Cambridge University Press. 1995.
- Meiss, M.R., Menczer, F., Fortunato, S., Flammini, A., & Vespignani, A. Ranking web sites with real user traffic. Proc. *WSDM*, 2008, 1067-75.
- Morris, M.R., & Teevan, J. Collaborative Web Search: Who, What, Where, When, and Why. Morgan and Claypool Publishers, San Rafael, CA, USA, 2010.
- Morris, M.R. & Horvitz, E. SearchTogether: an interface for collaborative web search. Proc. *UIST*, 2007, 3-12.
- Morris, M.R. and Horvitz, E. S3: Storable, Shareable Search. Proc. *Interact*, 2007, 120-123.
- Olston, C. & Chi, E. ScentTrails: Integrating browsing and searching on the Web. *TOCHI*. 10(3), 2003, 177-97.
- Pinelle, D., & Gutwin, C. Designing for loose coupling in mobile groups. Proc. *GROUP*, 2003, 75-84.
- Pirolli, P., & Card, S. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. Proc. *Intelligence Analysis*, 2005.
- Shah, C. & Marchionini, G. Awareness in collaborative information seeking. *JASIS*, 61, 2010, 1970–86.
- Wexelblat, A. & Maes, P. Footprints: history-rich tools for information foraging. Proc. *CHI*, 1999, 270-77.
- Willett, W., Heer, J., & Agrawala, M. Scented Widgets: Improving Navigation Cues with Embedded Visualizations. *IEEE Vis*, 13(6), 2007, 1129-1136.
- Wobbrock, J., Findlater, L., Gergle, D. & Higgins, J. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only ANOVA Procedures. Proc. *CHI*, 2011, 143-146.