Interfaces for Collaborative Exploratory Web Search: Motivations and Directions for Multi-User Designs

Meredith Ringel Morris
Microsoft Research
One Microsoft Way, Redmond, WA, USA
merrie@microsoft.com

ABSTRACT
In this position paper, we examine exploratory Web search as a collaborative activity and propose that such collaborations are commonplace. We present survey results that support this claim, and argue that current search interfaces provide limited support for common collaboration needs. We identify four features of the exploratory search experience (*coverage*, *confidence*, *exposure*, and *productivity*) that could be enhanced by providing explicit support for collaboration during the search and subsequent sensemaking processes.

Author Keywords
Exploratory search, computer-supported cooperative work.

ACM Classification Keywords
H5.3. Group and Organization Interfaces: Computer-supported cooperative work.

INTRODUCTION

Currently, searching the Web is considered a solitary task. The most commonly used tools for Web search, sites such as Google, Yahoo!, and Windows Live Search, are designed for single-user scenarios. However, joint and sequential activities by people collaborating on information-gathering tasks are commonplace. To date, there has been little work on providing means of enhancing collaborative exploratory search activities. Tools supporting collaboration during Web search promise to be valuable to users.

We provide motivation for focusing attention on multi-user information-gathering scenarios, review related work on this topic, and report on our own surveys regarding collaborative search experiences. Then, we discuss key features of the exploratory search experience that could be enhanced by providing explicit support for collaboration.

Motivation

Though not the target use scenario of current Web search engines, collaboration with friends, family, or colleagues during a search task is surprisingly commonplace. School-age children, for example, frequently work together at a single computer for the completion of group projects (both as a consequence of the limited supply of computing resources and to experience the pedagogical benefits of team work). Large *et al.*'s fieldwork with elementary-school students [4] found that, "Much class work, including project assignments, is undertaken by school students in groups rather than individually. Collaborative information-seeking, then, is not unusual."

Ethnographic work in higher education has also identified collaborative information-seeking practices. Twidale *et al.* [9] observed students using databases in a university library, and identified two types of collaborative information-seeking behaviors, neither of which was explicitly supported by the library's search application. The first, "joint search," involved groups of two to four students clustering around a single computer, frequently pointing to the screen while discussing ideas and planning actions. The second, "coordinated search," involved the group sitting with each member at a separate, but adjacent, workstation. The students would discuss the task, compare results, compete to find information quickly, and lean over to look at their teammates' screens.

Our own investigations, discussed in the following sections, provide evidence that collaborative searching is a strategy that is also employed in non-educational contexts, such as information-oriented workplaces, as well as in more casual settings. For example, colleagues in a research lab who are working together to write a scholarly article would benefit from a collaborative interface to facilitate a joint literature search. A group of friends planning a vacation to Hawaii might benefit from an interface that allows them to combine their efforts in identifying cheap airfares, appropriate hotels, and interesting tourist activities.

Related Work

Most research in the area of collaborative search focuses on passive forms of collaboration (*i.e.*, using data generated by large numbers of users' interactions with a system to impact current system behavior). Examples of such work include using query logs and clickthrough data to generate query substitutions or recommendations [1, 2]. Implicit data such as clickthrough actions of a large number of users can also be used as a collaborative filtering mechanism to re-rank search results [7].

There have been a few research efforts toward providing more active collaborative experiences among groups of users who know each other. However, these systems tend to be designed for very specialized domains and/or devices. TeamSearch [6] is a system that enables co-located groups of up to four people to simultaneously search collections of digital photographs, using a visual query language designed

for a multi-user interactive tabletop. Maekawa *et al.* [5] describe their system for groups of co-present people who each have a small, Web-enabled mobile device – to improve the efficiency of searching for information within a Web page (since scrolling through long Web pages on small screens is time-consuming), they allow a page to be split into several parts, each of which is displayed on a different user's device to facilitate parallelization of visual search. Krishnappa [3] developed a system that supports synchronous, remote collaboration between two people searching a medical database. Users of his system perform standard, single-user searches, but have a built-in textual chat facility as well as the ability to press a "share" button that sends some metadata about what they have found to the other user.

A few commercial products also offer support for collaboration during search tasks. For example, the Cha Cha search engine[1] pairs searchers up with a live person who assists them in formulating their query and suggesting interesting websites; however, this is not intended as an interface for allowing people to collaborate on an exploratory search task, but rather an attempt to support novice searchers by providing the ability to interpret natural language queries via human cycles. The Live Messenger[2] IM client provides a "shared search" feature whereby conducting a Web search through the client allows the list of returned URLs to be displayed to both the searcher and his/her IM partner. Google Notebook[3] allows a user to store clippings from several Web sites in one document; the tool provides a facility for allowing multiple users to add content to a single notebook document. Both Live Messenger's "shared search" feature and Google's notebook software are important steps toward facilitating collaborative exploratory search; however, both technologies focus on collaboration during the sensemaking [9] portions of an exploratory search task (*i.e.*, the viewing, selecting, and organizing of search results). Tools that allow users to collaborate in the formulation and refinement of queries during an exploratory search in addition to supporting collaborative sensemaking could offer many benefits, which we discuss in the following sections.

INITIAL EXPERIMENT: THE "AUNT EDNA" TASK

We conducted an online survey of ten knowledge workers in our research lab, in which they were asked to perform a specific exploratory search task. The task they were given was:

> *Your elderly Aunt Edna was recently diagnosed with high blood pressure. She's asked you to send her a link to a good source of information on treatments for her condition. Using your favorite search engine, find a single website that you would recommend to your Aunt Edna.*

Participants were asked to keep track of and submit the following information: the URL of the site they chose to recommend, the query terms they used to discover the chosen site, the search engine used to discover the chosen site, and their level of confidence (on a 5-point scale) that the site they chose was the best one for the task.

Note that this was not a collaborative task. We began our investigation with a single-user task in order to better understand how the single-user experience might be improved through the addition of collaborative support.

Results and Discussion

A total of seven distinct Web pages were "recommended" to Aunt Edna by the ten participants. Three sites were recommended by two participants each and the other four were each recommended by one. Eight different query formulations were employed to discover these sites (with two query formulations each being used by two users, and the other six by one each). Two different search engines were used (one by 40% of users, the other by 60%).

This variation in strategies (search engines and keywords used) and results (final URLs recommended) is quite interesting, especially considering that the participants in the exercise were all skilled searchers (all work in technical fields and use search technologies on a daily basis). In particular, this variation demonstrates potential benefits of collaboration for such tasks – the "Aunt Edna" task suggests that the breadth of tools used, query terms generated, and sites considered may increase as multiple users tackle the same search task.

Users' level of confidence in the quality of their chosen website also suggests room for improvement of the exploratory search experience. On a 5-point scale (with 1 = "Very Unconfident" and 5 = "Very Confident"), the mean rating of confidence that the recommended site was the best available was 3.0 (s.d. = 1.05). Notably, no participants selected the "Very Confident" option. Again, this points out a potential benefit of facilitating collaborative exploratory search – there is significant room for improvement in increasing users' confidence that they have succeeded at a search task; allowing them to combine and verify their efforts with others may be one means of accomplishing this goal.

Follow-Up Task

The next day, we sent a follow-up task to the same ten participants. This task, which was completed by six of the ten participants, asked them to visit each of the seven websites collected during the first task, and to answer two questions about each. The first asked them to rate (on a 5-point scale) the quality of the suggested page compared to the quality of the page they had originally recommended. The second asked them whether they had encountered this suggested page during their initial search task.

[1] http://www.chacha.com

[2] http://get.live.com/messenger/

[3] http://www.google.com/googlenotebook/overview.html

For each of the seven sites, at least three of the six respondents reported that they had not encountered that site during the initial search. For four of the sites, at least five of the six respondents reported not encountering it previously. In the previous section, we postulated that the large number of sites suggested in response to the initial task indicated that involving more users in a search task could increase the number of websites considered during an exploratory search. This follow-up data adds additional confidence to that hypothesis by demonstrating that users hadn't even encountered many of these sites during their search (rather than encountering them but choosing not to recommend them).

All but one site had at least one person (not including the person(s) who initially recommended that site) who felt the site was at least of equal quality to the site they had recommended. Three of the sites had at least one person who felt that the site was even better than the site he/she had recommended during the previous task. Again, these results point to the potential for collaboration to assist users in discovering high-quality information on the Web that they may not encounter on their own, and perhaps thereby increasing users' confidence that they have successfully met an information need through exploratory search.

SURVEY ON WEB SEARCH HABITS

While the results of the "Aunt Edna" task suggest potential benefits in allowing users to combine their search efforts, the task itself did not involve nor ask directly about collaboration. To learn more about current and desired practices regarding collaboration during search, we conducted a survey of employees at Microsoft. The survey was sent to 740 people, and was completed by 204 (giving a response rate of 27.6%). Entry into a prize drawing served as incentive for completion of the questionnaire.

Respondents represented a variety of knowledge workers, including researchers, software developers, managers, and administrative assistants. The population surveyed had a relatively high level of expertise in Web search – 99.5% report using a Web search engine at least once per day, and no respondents self-identified as "novice" users of search technologies.

53.4% of respondents answered "Yes" to the question "Have you ever cooperated with other people to search the web?". This is a surprisingly high number, considering that current technology is not explicitly designed to support cooperative searching. Of the people who said they had cooperated on a search task, 26.6% reported engaging in cooperative search activities at least once a week, and an additional 48.6% reported such activities occurring at least once a month.

Additionally, when we described specific collaborative behaviors, even some respondents who initially said they had never cooperated to search the Web reported engaging in a variety of "workarounds" that in fact enabled cooperative searching. For example, 87.7% of all respondents said they had watched over someone's shoulder as he/she searched the Web and suggested alternate query formulations to that person. 30.4% of respondents reported using instant messaging applications to coordinate real-time information seeking with a remote partner. 18.1% reported having divided up the responsibilities for sub-portions of a search task among several people, and then shared the results afterward.

18.1% reported needing/wanting to cooperate with others to search the Web, and being unable to find an effective workaround. Respondents provided some specific examples of search tasks where they had wanted the ability to cooperate with others. For example:

- "Ever try buying an airplane ticket at the cheapest possible price with someone else? Yikes."
- "Helping less computer-savvy users search the web (*e.g.*, my parents)."
- "Vacation planning: looking at hotels, dates, etc. … Goal: agreeing on vacation details."
- "We were trying to do a lit search but we both have different strategies for how we traverse the space. It was difficult to do together (because we wanted to follow different paths) … however, when we did it separately we weren't sure how much redundant information we were gathering."

DISCUSSION

Our survey on Web search habits revealed that a large number of people already engage in collaborative Web searching activities, even though these activities are not well-supported by current systems. The fact that people employ a variety of workarounds in order to collaborate suggests that collaboration support should be a first-class design requirement in Web search systems. It is likely that the frequency and pervasiveness of cooperative searching would increase if appropriate tools were available to end-users.

As indicated by the results to our "Aunt Edna" task, users employ different strategies to locate information (*e.g.*, different search engines, different query formulations, etc.), and subsequently encounter different, but still highly-relevant, websites during the search process. Creating user interfaces that support collaboration during exploratory search has the potential to improve search experiences and outcomes in several ways:

(1) better *coverage* of the space of relevant, high-quality sites.

(2) higher user *confidence* in the completeness and/or correctness of the search.

(3) *exposure* to varying search strategies and query syntax

(4) increased *productivity* due to a decrease in redundant information-seeking

Coverage

Facilitating collaboration during exploratory search could increase the number of relevant results discovered via several means. First, the simple increase in the number of pairs of eyes exploring the data could decrease the likelihood that a useful result is overlooked, and the increase in total person-hours devoted to the task is likely to increase with the number of collaborators as well. Additionally, the increased variety of search strategies (search engines, keywords, etc.) used is likely to increase the breadth of information encountered, giving another opportunity for additional relevant results to be discovered.

Confidence

The increased coverage of the search space in and of itself may inspire increased confidence among collaborating searchers that they have encountered all of the relevant sites during an exploratory search. Additionally, the ability of collaborators to view, and perhaps rate or comment on, the results contributed by others to a shared search could improve confidence in the quality of the results found.

Exposure

Collaborative search interfaces have the potential to serve as mechanisms not only for performing searches, but for helping users improve their own searching techniques. For example, providing users with awareness of query formulations used by their teammates could expose them to previously unknown syntax, which they could then use in their own future searches.

Productivity

By allowing people to collaborate with friends or colleagues on searches of mutual interest, redundant work could be reduced or eliminated, thus improving net productivity. For example, colleagues attempting to do a shared search now (for instance, by each searching independently on their own computers and then emailing each other lists of relevant links) must either spend time explicitly dividing up the search into subtasks for each user, or potentially duplicate each others' efforts by using the same keywords, finding the same documents, etc. In contrast, a system that provided a shared workspace with awareness of keywords used and sites flagged by each collaborator could reduce the overhead of coordination and duplication of effort.

An additional productivity benefit of interfaces enabling collaboration during Web search occurs in single-user scenarios. One important component of any collaborative system for exploratory search would be a persistent representation of the current state of the search process (*i.e.*, what queries have been issued, what results have been found, etc.), so that each of the collaborators could examine and add to the search session. This persistent representation also adds value to the single-user exploratory search experience since, after all, a user often must "collaborate" with him/herself, either across time (*i.e.*, beginning a search task that is interrupted and resumed hours, days, or weeks later) or across locales (*i.e.*, beginning a search task in the office, and continuing it on a different computer at home). By maintaining a persistent search state, single-user productivity could be improved by avoiding duplication of effort and by potentially increasing the speed with which a person can regain the context of their search task when resuming after a gap in time.

CONCLUSION

We have presented an overview of the current state of collaborative search interfaces, which shows the current lack of systems supporting collaboration for exploratory search tasks. Our survey on Web search habits showed that cooperative search is already a part of current work practice, although users must develop workarounds in order to share the process and results of a search with others. Data from our sample exploratory search task demonstrated potential benefits to providing formal support for group searching. In particular, we have described four aspects of exploratory search (*coverage*, *confidence*, *exposure*, and *productivity*) that stand to benefit from collaborative search UIs.

ACKNOWLEDGMENTS

We thank Eric Horvitz and Susan Dumais for their insightful discussions and comments.

REFERENCES

1. Beeferman, D. and Berger, A. Agglomerative Clustering of a Search Engine Query Log. *Proc KDD 2000*, 407-416.
2. Jones, R., Rey, B., Madani, O., and Greiner, W. Generating Query Substitutions. *Proceedings of WWW 2006*, 387-396.
3. Krishnappa, R. Multi-User Search Engine: Supporting Collaborative Information Seeking and Retrieval. *Master's Thesis, University of Missouri-Rolla*, 2005.
4. Large, A., Beheshti, J., and Rahman, T. Gender Differences in Collaborative Web Searching Behavior: An Elementary School Study. *Information Processing and Management*, Vol. 38, 2002, 427-443.
5. Maekawa, T., Hara, T., and Nishio, S. A Collaborative Web Browsing System for Multiple Mobile Users. *Proceedings of PERCOM 2006*, 22-35.
6. Morris, M.R., Paepcke, A., and Winograd, T. TeamSearch: Comparing Techniques for Co-Present Collaborative Search of Digital Media. *Proceedings of IEEE Tabletop 2006*, 97-104.
7. Pujol, J.M., Sangüesa, R., and Bermúdez, J. Porqpine: A Distributed and Collaborative Search Engine. *Proceedings of WWW 2003 (poster)*.
8. Russell, D., Stefik, M., Pirolli, P., and Card, S. The Cost Structure of Sensemaking. *Proceedings of CHI 1993*, 269-276.
9. Twidale, M., Nichols, D., and Paice, C. Browsing is a Collaborative Process. *Information Processing & Management*, 33(6), 1997, 761-83.