

UC Berkeley

UC Berkeley Previously Published Works

Title

What's Missing from Collaborative Search?

Permalink

<https://escholarship.org/uc/item/9gf057hd>

Journal

Computer, 47(3)

ISSN

0018-9162

Author

Hearst, Marti A

Publication Date

2014

DOI

10.1109/mc.2014.77

Peer reviewed

What's Missing from Collaborative Search?

Marti A. Hearst, University of California, Berkeley

People are Satisficing with Current Tools

It is fair to say that, according to the literature, tools dedicated to collaborative information seeking have not enjoyed widespread success to date (Morris et al. 2008, 2013). Surveys, interviews, and observations show that although now even more than ever, people are searching in pairs and in larger groups, they are doing so without the benefit of specialized search tools (Twidale and Nichols 1996). Instead, they are making do by coordinating via out-of-channel use of communication tools, especially email, texting, phone calls, and social media communication (Morris et al. 2008, 2013, Capra et al. 2010).

While Morris et al.'s recent study (Morris et al. 2013) showed that users were generally satisfied with their most recent collaborative search, the results also showed that there is much room for improvement. Participants expressed a need to increase group awareness of mutual activities, which is a functionality that has been addressed in collaborative search research systems (Shaw et al. 2013, Morris et al. 2007, Yue et al. 2012) as well as commercial systems. Other participants requested (i) the ability to compare results with others in real time, and (ii) a way to reduce redundant work. Furthermore, there may be specific categories of tasks that were not uncovered in the Morris et al. study, in part because participants reported only their most recent use of collaborative search. The "open answer" portion of the survey suggested that users do desire support for more such as complex tasks or tasks that extend over longer periods of time.

Research tools that allow for shared content and shared presence in both real time and asynchronous collaboration include SearchTogether (Morris et al. 2007), Coagmento (Shah 2013), CollabSearch (Yue et al. 2012), and Design Space (Capra et al. 2012). These tools allow participants to find, save, and share documents, and see the activities of others in the collaboration group.

A study by Shah 2013 compared an interface in which pairs of participants could not see any status information, could see their own personal action history, or could see

both their own and the action histories of their search partner. This work found that while the different conditions did not change the final search outcome in terms of quantitative evaluation metrics, awareness of the partner's search history did reduce the number of coordination messages that needed to be exchanged after the search was completed.

Similarly, the work of Jetter et al. 2011 introduced a highly novel interface for finding hotel reservations in a group, setting constraints via a combination of visualization and haptic displays. A key feature of this interface was that it allowed individuals to set constraints 'privately' in one corner of a tabletop display and then combine the constraints in the group publicly in the central view. Conflicts among the constraints were then adjusted collaboratively until a set of hotels that met some subset of the constraints could be agreed upon. This approach was compared to a standard web-based faceted interface. The overall outcomes were not distinguishable by the standard metrics, but the communication patterns were seen as having less 'noise' in the when using the interface that was tailored to collaboration, because it was easier to keep track of system state and individual's preferences.

Morris et al. talk about the friction inherent in having to switch from tools currently being used to a special purpose new tool. People just do not seem willing to move from a standard search tool to another tool for collaborative search to date. The reason for this may be that there must be enough additional value as yet in the tools offered, and/or they may not yet be easy enough to use, to justify using a specialized tool. If that is the case, the question is, what, if anything, can a collaborative search tool offer that would make it worthwhile to switch to a different tool, or add in a new plugin?

There is reason to hope that with the right approach, a tool could be developed to support collaborative search successfully. There are some recent prominent examples of new user interface solutions arising and becoming very widely adopted for long-standing problems in coordination and collaboration. For example, for many years there was no popular web-based way to schedule small-group meetings among people who did not share a calendaring tool; instead, people made use of email chains to schedule meetings. Now online web polls like Doodle are a standard solution (Reinecke et al. 2013).

As another example, for years there were no simple, high-usability tools available for an organization, such as a government agency, to use to put out a request for ideas and have thousands of lay people participate with comments and votes. The rise of online comment and moderation tools such as UserVoice and Google Moderator, not

to mention the rise of third generation question–answer forum tools like StackExchange and Quora, show that it is possible to make progress in this space.

What Are Key Missing Features?

The question is, what should the next generation of collaborative search tools be focused on? Below, we consider the special properties of three use cases in which collaboration is likely to be needed in a search task, and what therefore needs to be supported in a collaborative search tool above and beyond what existing tools have already successfully implemented.

Scenario 1: Selecting a Few from Many Similar Choices

Sometimes when people collaborate, it is because they need to make a decision together to meet the various group members' desires. The classic example is travel planning: Johnny wants a hotel with a good bar, Khoa wants to be near the beach, Pat doesn't want to pay very much, Sonali can't come till the third day, etc. Or some people engage in 'competitive shopping' in which they split up the work looking at many similar items offered at different venues to try to negotiate best prices.

For this case, a collaborative search tool should offer structure to let individuals define what their personal constraints or preferences are in some manner, as they search. People should not be required to specify this information up front, as that requires too much cognitive overhead and because the process of searching is likely to be what reveals reasonable values for those constraints in the first place. Perhaps as the searchers find important constraints, double-clicking on them or swiping them into a special location turns them into a tag that is visible to all in the group as something to pay attention to (price, near-the-beach) for others to pay attention to for further searches. A smart search engine will eventually learn what the common ones are and automatically recognize these across users and populate the constraints when search results are retrieved, perhaps using XML microformats or some other common representation.

Research by Fisher et al. 2012 has shown that people can build on simple representations that have already been begun by others, and if a search tool can recognize when such a representation is being built and suggest it as an organizing tool, or simply match the structure being built with those already seen to make more intelligent suggestions.

Scenario 2: Covering a Topic Thoroughly

The PhD student researching their dissertation work or a paralegal looking through a document trove trying to find all instances of some concept is the classic information retrieval challenge problem (Blair and Maron 1985). Collaborative search should be especially useful for this problem, since it should be amenable to divide-and-conquer, and since it can be implemented in special-purpose search tools such as academic reference search (Capra 2010) and legal document search.

But there are several fundamental functionalities that are missing from current collaborative search tools. For one, the tool should be aware both of what has already been stashed away in the bibliography and what has been viewed by anyone in the group of searchers. It should rerank based on this information, hiding what has already been assessed (but allowing users to override this setting). Prior research systems, including Querium (Golovchinsky et al. 2012) and Results Space (Capra et al. 2012) have made collaborators' explicit ratings visible, but assessments of these tools are done in laboratories in which users explicitly set up their tasks at fixed starting times, over fixed document collections. What to do with leftover ratings, implicit information (Fox et al. 2005), and new group members joining and leaving the collaboration, are still open questions.

It may be useful to allow a trio of searchers to work together, with one doing triage using a general query and a general ranking algorithm, another looking at promising documents in more detail, and classifying those that are relevant using a set of agreed-upon categories or tags, a third using a different ranking algorithm to further search within one of the refined categories, perhaps along the lines suggested by Pickens et al. 2008.

Such a tool should have a clear depiction of the landscape covered and the landscape yet to be looked at, organized by pre-determined categories as well as user-defined keywords and sortable by citations or usage, where available. Searchers should be able to achieve a feeling of accomplishment, of walking over the landscape and having their bearing on the terrain, as they work, and be able to strategize about who is to do what next over that map.

Scenario 3: Discovering Unknown Information

By far the most challenging search task is that of trying to help people make a new discovery, such as solving the mystery of why honeybees are dying in North America. For a problem like this, search over known materials is only one part of a multi-faceted effort, wide-ranging collaborative effort. But it is worth considering

the role of collaborative search for this type of problem. Any collaborative search tool should always be comparing what has already been added to a community collection to what is currently being viewed to see if the information is redundant. Text mining algorithms have long been proposed as an aid to discovery of new information (Hearst 1999) but data mining and knowledge discovering algorithms have not emphasized collaboration for the most part. Collaborative challenges, such as the DARPA red balloon challenge, in which teams were given 10 hours to collaboratively discover the physical position of 10 red weather balloons released before dawn across the continental United States (Tang 2011) or the Polymath Project (Gowers et al. 2009, Cranshaw and Kittur 2011), in which the world was invited to collaboratively solve a difficult mathematics proof, and the goal was achieved by 27 people in a matter of months using comments on blogging and wiki platforms, suggest directions forward.

Conclusion

Because people are searching together on a regular basis, there is a need for support for this activity in search user interfaces. The fact that no such interface has caught fire does not suggest there is no need for such a design, but rather that the best parts of the design space have not been investigated fully yet. What they are and how to present them remain an open question; this article has made a few suggestions about directions for exploration.

References

Blair, D.C., and Maron, M.

An evaluation of retrieval effectiveness for a full-text document-retrieval system.

Communications of the ACM 28}, 3 (1985), 289--299.

Capra, R., Chen, A.T., Hawthorne, K., Arguello, J., Shaw, L., and Marchionini,

G. Design and evaluation of a system to support collaborative search.

Proceedings of the American Society for Information Science and Technology 49}, 1 (2012), 1--10.

Capra, R., Marchionini, G., Velasco-Martin, J., and Muller, K.

Tools-at-hand and learning in multi-session, collaborative search.

In Proceedings of the SIGCHI Conference on Human Factors in

Computing Systems}, ACM (2010), 951--960.

Cranshaw, J., and Kittur, A.

The polymath project: lessons from a successful online collaboration in mathematics.

In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems}, ACM (2011), 1865--1874.

Fisher, K., Counts, S., and Kittur, A.

Distributed sensemaking: improving sensemaking by leveraging the efforts of previous users.

In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems}, ACM (2012), 247--256.

Fox, S., Karnawat, K., Mydland, M., Dumais, S., and White, T.

Evaluating implicit measures to improve web search.

ACM Transactions on Information Systems (TOIS) 23}, 2 (2005), 147--168.

Golovchinsky, G., Diriye, A., and Dunnigan, T.

The future is in the past: designing for exploratory search.

In Proceedings of the 4th Information Interaction in Context Symposium}, ACM (2012), 52--61.

Gowers, T., and Nielsen, M. Massively collaborative mathematics.

Nature 461}, 7266 (2009), 879--881.

Hearst, M.A. Untangling text data mining.

In Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics}, Association for Computational Linguistics (1999), 3--10.

Jetter, H.-C., Gerken, J., Zöllner, M., Reiterer, H., and Milic-Frayling,

IEEE Computer, March 2014, pp 58-61

N. Materializing the query with facet-streams: a hybrid surface for collaborative search on tabletops.

In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems}, ACM (2011), 3013--3022.

Morris, M.R.

A survey of collaborative web search practices.

In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems}, ACM (2008), 1657--1660.

Morris, M.R.

Collaborative search revisited.

In Proceedings of the 2013 conference on Computer supported cooperative work}, ACM (2013), 1181--1192.

Morris, M.R., and Horvitz, E.

Searchtogether: an interface for collaborative web search.

In Proceedings of the 20th annual ACM symposium on User interface software and technology}, ACM (2007), 3--12.

Pickens, J., Golovchinsky, G., Shah, C., Qvarfordt, P., and Back, M.

Algorithmic mediation for collaborative exploratory search.

In Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval}, ACM (2008), 315--322.

Reinecke, K., Nguyen, M.K., Bernstein, A., N{ "a}f, M., and Gajos, K.Z.

Doodle around the world: Online scheduling behavior reflects cultural differences in time perception and group decision-making.

In Proceedings of the 2013 conference on Computer supported cooperative work}, ACM (2013), 45--54.

Shah, C.

Effects of awareness on coordination in collaborative information

seeking.

Journal of the American Society for Information Science and
Technology 64}, 6 (2013).

Tang, J.C., Cebrian, M., Giacobe, N.A., Kim, H.-W., Kim, T., and Wickert,
D.B., Reflecting on the darpa red balloon challenge.
Communications of the ACM 54}, 4 (2011), 78--85.

Twidale, M., and Nichols, D.

Collaborative browsing and visualisation of the search process.
Aslib Proceedings 48}, 7-8 (1996), 177--182.

Yue, Z., Han, S., and He, D.

An investigation of search processes in collaborative exploratory web
Search. Proceedings of the American Society for Information Science and
Technology 49}, 1 (2012), 1--4.