
Exploring the Cognitive Consequences of Social Search

Brynn M. Evans

Department of Cognitive Science
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093
bmevans@cogsci.ucsd.edu

Sanjay Kairam

Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94116
sanjay.kairam@gmail.com

Peter Pirolli

Palo Alto Research Center, Inc.
3333 Coyote Hill Road
Palo Alto, CA 94116
pirolli@parc.com

Copyright is held by the author/owner(s).
CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract

To what extent can social interactions augment people's natural search experiences? What factors influence the decision to turn to a friend for help? Our paper presents the preliminary results of a social sensemaking task that begin to address such questions by examining the cognitive consequences of social search.

Keywords

Social search, social information seeking, cognitive task analysis, verbal protocols.

ACM Classification Keywords

H.5.3. Group and organization interfaces: Computer-supported cooperative work; Theory and models.

Introduction

Where do you turn when you have a question or pressing inquiry? It depends, of course, on the problem at hand, but increasingly we have come to rely on personal computers and pervasive Internet technologies to support our daily information needs. The problem with this approach is that Google and Yahoo! searches rarely capture the context or complexity of our situated, practical problems.

Churchill recently argued that we should rethink the

system design behind web search because many of our natural inquiries extend over a period of time and often involve the help of real-world (social) entities [2]. Although library scientists have used similar conceptual frameworks to describe search behaviors [11, 13], only recently have information and computer scientists begun to consider the naturally distributed, social settings of user behaviors on the web. For example, there have been recent efforts to design social systems that facilitate collaborative actions (tagging [9] and search [8]) and provide recommendations to others [10]. Others have explored the role of naturally occurring social interactions during web search [4, 7]. While these approaches, among others, are part of an early effort at *social search*, our current work is concerned specifically with the way individuals make use of peers and other available social resources during search tasks.

Both Morris [7] and Evans & Chi [4] discussed the prevalence and importance of social interactions during web search. Inputs from colleagues and friends provide advice, guidance, and brainstorming opportunities during the search process; and social outlets are important for sharing information and seeking feedback following the primary search task. Additionally, peer support provided the greatest benefit to users performing *informational* searches—or when exploring unknown problem spaces where information needs were poorly defined [4]. This type of exploratory searching can be hard to support from a system's perspective, due to the occasional gulf between users' concepts and keywords and the jargon of the problem domain (known as "the vocabulary problem" [5]). This problem may be mitigated, however, in human-human communication, suggesting the potentially powerful role of social inputs in *informational* sensemaking tasks.

Project Goals

Given this, we wanted to look more closely at the cognitive consequences of social interactions during search in novel problem domains. More specifically, how do *social resources*—from friends and neighbors, to social networks, question-answer sites, and blogs—enhance the search experience over typical web search alone? What factors influence the decision to ask a friend for help?

Methods

Using Cognitive Task Analysis techniques, we collected the verbal protocols of 8 users as they performed two sensemaking tasks related to U.S. energy policy—in one condition being restricted to only social resources, in another to only web (information) resources.

Task Questions

Our specific task questions were chosen for several reasons. First, we wanted to observe the sensemaking process in a technical (presumably novel) domain, so we required questions that were not easily "Googleable". Second, we wanted to preserve a sense of real-life relevancy. At the time of the study, energy policies were hotly debated since gas prices had been soaring. The following two questions about U.S. energy policy seemed suitable given these goals:

- [1] If we lowered the speed limit nationally to 55 mph, how many fewer barrels of oil would the U.S. consume every year?
- [2] What role does pyrolytic oil (or pyrolysis) play in the debate over carbon emissions?

Participants

We screened 35 people from the local community with a detailed questionnaire, selecting eight to participate based on high scores on measures of social activity and expertise in searching the web (3 males).

Procedure

All participants engaged in two, untimed blocks, each consisting of a talk-aloud search task, per standard verbal protocols [3], followed by a recap and discussion of the search procedure. Participants used their own personal laptops, which we video recorded with screencast software.

Task questions were counterbalanced between subjects; conditions were not. The Social Condition always preceded the Web Condition since we expected occasional delays in communication with social resources. If this occurred, we could minimally observe the natural time-course of events—one of the expected limitations of social search.

The Social Condition restricted participants to make use of only *social resources* (e.g., friends, coworkers, social networking sites, question-answer sites online, blogs), but there were no restrictions in *how* they could reach out to others (face-to-face, phone, or web-based interactions were all fine). Social resources did not include online information databases, Wikipedia, or any traditional search engines (Google, Yahoo, MSN).

The Web Condition restricted participants to make use of only *information resources* (e.g., search engines, Wikipedia). Sources used in the first task were not permitted here.

Preliminary Results

Our procedures resulted in a rich data set of quantifiable events and qualitative observations, from both video screencasts and detailed questionnaires. The preliminary results reported here are intentionally restricted to video analysis of the user behaviors, tactics, and outcomes in the Social Condition only. Our first goal is to understand how using social resources in web search affects problem solving in new domains.

Social Tactics

As expected, some participants were more successful in their social search tasks than others. Our measure for success was based on two *four-point* rating scales: one for exploration of the problem space, one for accuracy of the final answer.¹ We collected success scores from two independent raters, and later compared the sum of both rating scales. This placed our subjects into three success conditions: High, Medium, and Low.

Table 1. Success groupings for the social search task.

High	Medium	Low
s01, s03	s02, s05, s08	s04, s06, s07

A natural question is what accounts for these differences in success? We identified three social tactics that may account for the user behaviors and outcomes observed in the social search condition:

- *Directed Asking*: targeting specific friends or colleagues to ask for help (e.g., over email, IM).

¹ Accuracy was based on government and academic reports on these topics [6, 12].

- *Public Asking*: posting a question in a public venue, typically an organized social network (e.g., Twitter, Facebook, Ask MetaFilter, Aardvark).
- *Searching*: searching over repositories of social data (e.g., Yahoo! Answers, MetaFilter, blogs).

It may be noteworthy that none of these tactics involved *pushing* information to the user. This is surely an artifact of both our experimental design and currently available web technologies. Future work could compare how user strategies differ when passively receiving social data (e.g., recommendations) versus actively seeking it.

Strategies and Outcomes of Social Tactics

Each of these social tactics has potential merits. Targeting individuals (through *directed asking*) can be useful if they are knowledgeable and available to respond. *Public asking* distributes a question over a wide (diverse) audience, theoretically increasing the likelihood of reaching an individual with the appropriate knowledge and availability. *Searching*, of course, could provide ample information (with large databases), but is limited to the content already present in the database.

The most successful participants used all of these social tactics in combination, interacting with multiple different entities (people and databases) over the course of their search. For example, s01 began by publicly asking friends on Twitter, and then targeting individuals on instant messenger (IM) for help with her task. She later switched to searching over Yahoo! Answers, while simultaneously discussing the problem with a friend on IM. S03's process was reversed: He first searched over MetaFilter, then asked his question

publicly on MetaFilter and Aardvark (both question-answer sites), and finally targeted specific friends on IM.

Surprisingly, any one social tactic seemed to produce lower success when it was the entirety of the search strategy. For example, *directed asking* was the only information channel for s02 and s08. Both participants contacted close friends (via email and phone) who they identified as "knowledge mavens," or domain experts. Their conversations were short, succinct, and produced reasonable information; but the answers were incomplete, and participants did not continue their problem solving when the conversations ended. After observing the success of *directed asking* in combination with other social tactics, why does it fail on its own? Is one possible shortcoming of "expert" (social) opinions that people place *too much* trust on their data?

Even more surprisingly, we found that the solo-strategies of *public asking* and *searching* were the least successful. This is somewhat misleading since searching can be an effective technique—and it was—for certain inquiries. S04 relied exclusively on searching, however, and never managed to reach the heart of the problem.

Similarly with *public asking*, we expected to see an increase in responses when people distributed questions over large—presumably diverse—social networks. But as a solo strategy with s06, this was not the case. He very deliberately posted a question to multiple online social networks through a service called Ping.fm and quickly received a reply; only the response was inaccurate. In some ways this was typical of other subjects' experiences: the largest social networks (>1000 individuals) did produce quick replies, with content that was short, incomplete, and often off-topic.

On the other hand, some participants were successful at incorporating *public asking* into their overall strategy. In fact, a series of quick, short replies may help illuminate relevant aspects of the problem space or expedite the sensemaking process, especially when combined with other social tactics. For example, s01's early Twitter post produced several responses that may have changed the framing and formulation of her question, as evidenced by her subsequent search queries on Yahoo! Answers.

Of course, a number of factors could affect search strategies at any given moment—the specific information needs, intrinsic motivations, external pressures, and cultural expectations. Individual differences and personal preferences undoubtedly also account for the varying strategies and outcomes we observed. Future analyses will look more closely at these, and other, factors that likely influenced many of the social and tactical decisions our participants made.

Social Decisions

Our second immediate goal was to understand why someone would ask a friend for help. We are calling these *social decisions* after the set of *social* factors that go into the *decision* to use certain social resources during search.

We found that our participants' social decisions matched those laid out in the formal model of social information seeking by Borgatti & Cross [1]: perception of another person's knowledge and authority ("I trust what she says" [s06]); perception of their accessibility ("I know he's online and I know he's at home" [s01]); and perception of the social costs, or obligations that would be incurred ("He's going to hate me for this...bugging

him about stuff that's not relevant to our usual interaction" [s03]).

We noted several additional factors, as well: the history of the relationship, time since the last interaction, and the technology used to communicate (in the present and historically). For example, s08 tells us why she chose David: "We have an engagement over email where we can do this kind of thing." S01 instant messages Alex not because he's an expert, but because she had "been IM'ing with [him] recently." S07 hopes to use her Facebook buddy list to find someone to ask: "I can see who's online right away...who's even on here?"

Understanding how these social decisions relate to users' overall objectives in social search tasks is absolutely critical for developing supporting tools. While our study cannot address the reasons for turning to social resources in the course of natural search tasks, it does begin to illuminate some of the cognitive consequences of social interactions in search.

Conclusion and Future Directions

Our preliminary findings already build upon related work in the collaborative information seeking and social recommender communities. We identified three social tactics for information gathering (directed asking, public asking, and searching), and see early evidence that using these tactics in combination may lead to a more productive social search. Surprisingly, asking questions to large social networks had some drawbacks, as did relying exclusively on just a single tactic. We plan to look more closely at the cognitive effects of these social tactics for individual subjects relative to their own processes and reported habits, and relative to their counter search task in the Web Condition.

We also intend to identify the patterns and rules that appear to guide users as they navigate novel problem domains—and how social inputs can alter or augment the natural course of sensemaking, beyond providing merely informational support. A specific conjecture is that two-way conversations, in the manner of *directed asking* in semi-synchronous channels (e.g., phones, IM), will cause subjects to think critically about the problem space earlier or more completely than with asynchronous communication or in solitary search. It could be added that we expect strong ties to contribute here more so than weak ties, an effect that we hinted at

earlier while discussing the quick, short nature of replies from large social networks.

We're beginning to see that *social search* is about more than just providing access to new information—social resources may change how we formulate a search problem. Already, social interactions have been shown to improve cognitive functions in lab settings [14]; we are at the point of understanding how social inputs provide cognitive benefits to web searchers, too. With new insights into this process, we hope to design social tools that provide cognitive (and contextualized) support to users in their natural working environments.

References

- [1] Borgatti, S.P. and Cross, R. A relational view of information seeking and learning in social networks. *Management Science* 49, 4 (2003), 432-445.
- [2] Churchill, E. Of candied herbs and happy babies: Seeking and searching on your own terms. *Interactions* 15, 6 (2008), 46-49.
- [3] Ericsson, K.A. and Simon, H.A. Verbal reports as data. *Psychological Review* 87, 3 (1980), 215-251.
- [4] Evans, B.M. and Chi, E.H. Towards a model of understanding social search. *Proc. CSCW'08*, ACM Press (2008), 485-494.
- [5] Furnas, G.W., Landauer, T.K., Gomez, L.M., and Dumais, S.T. The Vocabulary-Problem in Human-System Communication. *Comm. of the ACM* 30, 11 (1987), 964-971.
- [6] Lehmann, J. Bio-energy in the black. *Frontiers in Ecology and the Environment* 5, 7 (2007), 381-387.
- [7] Morris, M.R. A survey of collaborative web search practices. *Proc. CHI'08*, ACM Press (2008), 1657-1660.
- [8] Morris, M.R. and Horvitz, E. SearchTogether: An interface for collaborative web search. *Proc. UIST'07*, ACM Press (2007), 3-12.
- [9] Naaman, M. and Nair, R. ZoneTag's collaborative tag suggestions: What is this person doing in my phone? *IEEE MultiMedia* 15, 3 (2008), 34-40.
- [10] Smyth, B. A community-based approach to personalizing web search. *Computer* 40, 8 (2007), 42-50.
- [11] Twidale, M.B., Nichols, D.M., and Paice, C.D. Browsing is a collaborative process. *Information Processing and Management* 33, 6 (1997), 761-783.
- [12] United States Government Accountability Office. *Energy Efficiency: Potential Fuel Savings Generated by a National Speed Limit Would be Influenced by Many Other Factors*. Washington: Government Printing Office, 2008. (GAO-09-153R).
- [13] Wilson, T. On user studies and information needs. *Journal of Documentation* 37, 1 (1981), 3-15.
- [14] Ybarra, O., Burnstein, E., Winkielman, P., Keller, M., Manis, M., Chan, E., and Rodrigues, J. Mental exercising through simple socializing: Social interaction promotes general cognitive functioning. *Personality and Social Psychology Bulletin* 34, 2 (2008), 248-259.