

Towards a Model of Understanding Social Search

Brynn M. Evans
University of California, San Diego
La Jolla, CA
bmevans@cogsci.ucsd.edu

Ed H. Chi
Palo Alto Research Center
Palo Alto, CA
echi@parc.com

ABSTRACT

Search engine researchers typically depict search as the solitary activity of an individual searcher. In contrast, results from our critical-incident survey of 150 users on Amazon’s Mechanical Turk service suggest that social interactions play an important role throughout the search process. Our main contribution is that we have integrated models from previous work in sensemaking and information seeking behavior to present a canonical social model of user activities before, during, and after search, suggesting where in the search process both explicitly and implicitly shared information may be valuable to individual searchers.

ACM Classification Keywords

H.5.3 Information Systems: Group and Organization Interfaces—*theory and models, collaborative computing, web-based interaction.*; H.3.3 Information Storage and Retrieval: Information Search and Retrieval—*Search process.*

Author Keywords

Social search, social navigation, information seeking, sensemaking, web browsing.

INTRODUCTION

Web search has changed dramatically how we interact with the knowledge of the world. Its success in impacting our everyday lives in the last two decades is perhaps unparalleled. Surprisingly, however, researchers have mostly thought about navigating and browsing for information as a single user activity, centered on eliciting users’ information needs and improving the relevance of search results.

This view is somewhat in conflict with prior research by library scientists looking at users’ information seeking habits [11, 21, 31, 34, 35]. This work was done by scientists before the wide availability of web search engines; but this research demonstrated that other individuals may be valuable information resources during information search. More recently, researchers have observed direct user cooperation

during web-based information seeking. Morris [25] conducted a survey of 209 enterprise users, revealing that nearly half engaged in explicit collaboration on joint search tasks on the web. Certainly, active collaboration by multiple parties does occur under some circumstance (e.g., enterprise settings); at other times, and perhaps for a greater majority of searches, users may interact with others remotely, asynchronously [29], and even involuntarily and implicitly.

The general term “social search” has been applied widely in the field of Web 2.0 to describe searches that: utilize social and expertise networks; are done in shared social workspaces; or involve social data-mining or collective intelligence processes to improve the search process. Our definition of “social search” is intended to be broad, to include a range of possible social interactions that may facilitate information seeking and sensemaking tasks:

“Social search” is an umbrella term used to describe search acts that make use of social interactions with others. These interactions may be explicit or implicit, co-located or remote, synchronous or asynchronous.

Our focus, therefore, is to bring some clarity to how *social search* occurs in the real world. We analyzed the self-reported search experiences of 150 users from Amazon’s Mechanical Turk, and mapped their complex social activities onto a single, canonical model of the extended search process. Our main contribution is to present this integrated model of social search, specifically noting where and why social interactions occurred before, during, and after a search event. We support the model with observations from our data: users’ search motivations; their pre-search preparation process (seeking guidance, advice, and clarifications on the information need); how they conducted searches according to those information needs (*transactional, navigational, informational*); and how they shared end results after the search.

In the rest of this paper, we briefly review past research, describe our survey and data collection procedure, and present the canonical model, both as a diagram and with quantitative support and anecdotal case studies of actual behavior. Finally, we conclude with design implications, limitations, and some general remarks.

RELATED WORK

Until quite recently, researchers have mostly thought about navigating and browsing for information as a single user ac-

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’08, November 8–12, 2008, San Diego, California, USA.

Copyright 2008 ACM 978-1-60558-007-4/08/11...\$5.00.

tivity [5, 7], even among those who developed behavioral models of information seeking [2, 6, 8, 23]. Ellis' early work in understanding the behavioral patterns of users as they search for information led to a basic model of information seeking characterized by six general categories: starting, chaining, browsing, differentiating, monitoring, and extracting [8]. Marchionini expanded this model to describe how the user acknowledges and defines the information need, formulates the query, executes the search, and examines and internalizes the results [23]. Choo, Detlor & Turnbull [6] and Bates [2] have subsequently presented more integrated models of search on the web. However, the "modes of information seeking" that they discuss focus more on categories of search behaviors and motivations, than on how people actually accomplish search. Like much of the information seeking literature, this work overlooks the role of other individuals in search, instead focusing on the search act from a single user's perspective.

It is interesting, therefore, that library scientists have recognized for some time that other individuals may be valuable information resources during the search process [35]. Even prior to the search, "the inquirer decides whether to discuss his problem with a colleague or go to whatever literature or information center may be available" [33]. Later, users refine their topic selection in preparation for the search by consulting friends and colleagues who serve as sounding boards for ideas [21] and who provide pointers to key references in the literature [11, 31]. More recently, Twidale, Nichols & Paice have highlighted the prevalence and benefits of collaborative searching [34]. By studying the behavior of library users at computer terminals and throughout the library, they observed that co-located, synchronous collaboration occurs both during the search *process* and after obtaining the *end product* (e.g., query results).

In addition to such explicitly collaborative interactions, a number of researchers have focused on personalizing web search through more implicit means [13, 14, 16, 20, 32]. Social recommendation systems, for example, use techniques like collaborative filtering to provide recommended items to information seekers [13, 14, 20, 32], based on the opinions or activities of other people. These approaches are arguably also "social search" systems since they make use of latent community behavior, even though they typically do not exploit users' explicit social networks to personalize search.

Considering the related work in the library sciences and social recommender research communities, a reasonable hypothesis is that a significant portion of web search involves social acts, and that social interactions can improve the search process. Indeed, an increasing number of Web 2.0 sites provide various types of social inputs which could be used to augment search. Tagging services allow users to socially annotate web links for personal (del.icio.us [22], Ma.gnolia, diigo), academic (CiteULike, Connotea), and enterprise use (Dogear [24], ConnectBeam). Networking sites allow users to maintain social connections with peers (MySpace, Facebook) and business partners (LinkedIn). Even search engines are taking novel approaches to listing search results—

for example, through votes of link relevance (Wikia) or human-written entries (Mahalo). Yet despite the potential of these sites, only a few researchers have explored how such "social web services" might improve the search process [15, 36].

More importantly, we still do not fully understand users' social information seeking and sensemaking needs. How and where in the search process do users interact with others? How can social networks help? What benefit do social interactions have on the quality or efficiency of the search process? Morris [25] recently looked at collaboration surrounding web search activities and found that office workers often coordinated with others during joint search efforts. In response to these findings, she has worked with colleagues to implement several search tools to explicitly support collaboration among small groups of people for shared web search tasks (SearchTogether [26], CoSearch [1]).

We believe that the benefits of social search will extend beyond joint collaborative (synchronous and co-located) search to more implicit, asynchronous, and remote interactions. Our focus in this paper is to begin to explore a model of social search that may offer suggestions for supporting social interactions in the information seeking process.

PROCEDURE

We surveyed users about their most recent searching behaviors using Amazon's Mechanical Turk, a type of micro-task market, which can engage a large number of users to perform evaluation tasks both at low cost and relatively quickly [19]. Following Kittur, Chi & Suh [19], we took special care to formulate our task to probe for specific information and reduce invalid responses.

The Survey

Our survey was also designed as a critical-incident self-report, in which users describe events of a certain class or quality that occurred relatively recently [4, 10]. We recruited users with a specific statement of our purpose:

We are interested in how you search for digital information on your computer. Please answer the following questions about your most recent search experience.

The survey prompted users for details surrounding the selected incident, presented as yes/no, multiple choice, or free-form responses. We collected information related to the search *context* and *purpose*, additionally asking how (or if) users interacted with other individuals prior to and following the primary search act. For example, users were shown the following series of search-related questions.

1. When was the last time you searched for information?

- today, recently
- today, earlier in the day
- yesterday
- 2 days ago
- more than 2 days ago

2. What kind of information was it? [free-form answer]

Profession	%Users	Job Role	%Users
Education	9.3	Manager	19.3
Financial	8.7	Assistant	18.7
Healthcare	6.7	CEO/Director	8.0
Govt. Agency	6.0	Customer Support	7.3
Retail	6.0	Teacher	6.0
Software	6.0	Programmer	6.0
Research	5.3	Analyst	4.0

Table 1. The most frequently occurring professions and job roles reported by users in our sample.

3. What were you doing just before you searched? [free-form]
4. Did you talk with anyone (face-to-face, email, phone, etc.) before you searched? Why or why not? [free-form answer]
5. If yes, was your conversation related to your current search? [free-form answer]
6. What prompted you to perform the search? [free-form answer]
7. What steps did you take to find this information? [free-form]
8. What did you do just after you searched? [free-form answer]
9. If other people were nearby, were you interacting with them or were they influencing your search process?
 - There were no others in the room.
 - There were others, but I was not interacting with them.
 - There were others, and my interaction with them was related to my search.
10. If other people were nearby, please explain your interaction with them. [free-form answer]
11. After you found the information, did you share it with anyone?
 - yes
 - no
12. Why or why not? [free-form answer]
13. If yes, how did you share the information? [free-form answer]

Finally, users provided background information on their profession, job roles, and job expertise. They commented on how the reported incident was similar to and different from related search experiences, when the search occurred, and how long it took to complete.

Sample Data

We collected 150 complete survey responses from anonymous individuals on Mechanical Turk. Since critical-incident self-reports are most accurate in recounting recent experiences [10], we elicited user responses about their *latest* search acts. Consequently, about 2/3 of search acts occurred on the same day that users filled out our survey (48.7% occurred “recently” and 14.7% occurred “earlier in the day”). 19.3% of searches occurred the day before, and 17.3% occurred more than 2 days ago.

Search acts were relatively brief, nearly half requiring fewer than five minutes to complete. Respondents came from a range of professional sectors, job roles (Table 1), and levels of job experience, as rated on a 5-point Likert scale (Table 2).

A majority of our users only searched for information on the Internet (111/150, 74.0%). The remainder used tools that

Search Duration	%Users	Job Expertise	%Users
< 5 minutes	44.7	5	33.3
5–10 minutes	23.3	4	35.3
10–20 minutes	10.7	3	20.7
20–30 minutes	13.3	2	7.3
> 30 minutes	8.0	1	3.3

Table 2. Information reported about duration of the search act and level of job experience.

were internal to their professional organizations (e.g., Outlook, software programs, company Intranet). Only two users reported using both Internet and internal tools over the course of their search act.

Finally, we categorized each reported incident according to Broder’s taxonomy of information needs in web search [3]:

transactional: performing a transaction and extracting information after a source or website is located.

navigational: following a series of steps to identify a known fact or website.

informational: searching for information assumed to be present, but otherwise unknown.

RESULTS

Our main contribution is that we have integrated our findings with models of sensemaking and information seeking from the literature, and we present a canonical model of user activities throughout the search process (Figure 2, below). We will discuss our model in three phases: before search, during search, and after search, providing both quantitative data and anecdotal case studies of actual, reported user behavior. Within the model, we highlight the places where information exchange occurred through social interactions, as indicated by the illustration in Figure 1.




Figure 1. Users exchanged information with others through social interactions in a number of places throughout the search process.

Before Search

Context Framing

Information-seeking behavior is rooted in a “need” to find information [3, 35] or a motivation that drives the search process. This may be thought of as the *context framing* stage of search, where user motives and information needs are defined. Requests for information may come from an external source (e.g., specific request from a boss, customer, or client) or may be self-initiated (e.g., finding information related to personal or work endeavors).

External requests. (47/150 users, 31.3%). About 1/3 of the searches were motivated by external sources (e.g., specific request from a boss, customer, or client). As an example, a Dell customer support representative was searching




Figure 2. Canonical social model of user activities before, during, and after a search act, including citations from related work in information seeking and sensemaking behavior.

for information related to a promotion advertised on another website—a search prompted by a specific customer request. The service representative was tasked with finding this information and consequently reporting back relevant information to the customer on the phone.

Self-initiated search. (103/150 users, 68.7%). Our data reflect that over two-thirds of user searches were self-motivated prompts to find information related to personal or work endeavors. As an example, while creating a spreadsheet of songs to teach to students, a piano teacher was searching for JPG files of sheet music that highlighted certain musical concepts. Since this teacher always works alone, she is constantly “seeking new resources to keep piano students interested, and to teach them concepts through enjoyable songs and activities.” Therefore, this search behavior is a career-related, self-motivated process for continually finding new and better material.

Requirement Refinement

After information needs and motives are established, search requirements are refined. Previously described as a *generation loop* [30], this phase involves gathering requirements and formulating relevant schemas such that an effective search may result. As an example, an operations technician at Lexus-Nexus needed to collect certain information in order to prepare a new client file. Prior to the search, a meeting was called with several other colleagues to discuss what type of information should be included in the file. After establishing the guidelines, the technician searched on an internal program to extract the information agreed upon in the meeting. This account demonstrates how one individual effectively formulated requirements for the subsequent search, and, more importantly, how social inputs improved the process.

In fact, this phase in the model is marked by social interactions 42.7% of the time (64/150 users) both to “influence the information need” [34] and refine the task guidelines. However, social behaviors differed depending on users’ search motivations. Externally-motivated searchers interacted socially quite frequently during the pre-search phase, but almost always out of obligation. Self-motivated searchers engaged with others for a greater variety of reasons, including seeking advice, feedback, and personal guidance.

Of the 47 externally-motivated searchers, 33 performed information exchanges with others (70.2%) to identify requirements or guidelines for the search. Typically, the information exchange was with the source of the request, such as a client or boss. A social worker for the State of Washington was helping a homeless, disabled client find temporary housing, which required learning about the “client’s disability accommodations.” The need for more information resulted in an obligatory conversation between the social worker and the client, after which the social worker had the resources necessary to complete the request.

Social interactions occurred much less frequently for self-motivated searches (31/103 users, 30.1%), although the rea-

sons for engaging with others were more varied. In addition to establishing the guidelines for the search task, information exchange also occurred to seek the advice of others (8/31 users, 25.8%): When an associate at Circuit City wanted to improve his commission-based sales, he solicited a colleague’s personal opinion before searching on Yahoo! for additional suggestions. Furthermore, pre-search social interactions were used to brainstorm (3/31 users, 9.7%) and collect search tips (e.g., keywords, URLs) from friends and colleagues (2/31 users, 6.5%).

Establishing guidelines occasionally required a preliminary search, where users would engage with others. A salesperson for a company that imported tile and natural stone received a customer request for 5,000 square feet of Pennsylvania Blue Stone for a pool deck. This “search” act involved finding a supplier who had the stone in stock, and really began when the salesperson called several local suppliers to see if they carried the material. Only after learning that his regular suppliers were out of stock, he completed the request by performing an *informational* search on Google for additional suppliers. This example demonstrates both that search is a fluid process (making it somewhat inaccurate to draw a line between each search phase), and also that social interactions prior to computer-based search may provide necessary background information or may serve as the first stage of an extended search act.

In summary, these accounts of pre-search activities suggest that many users take explicit actions to prepare for and improve their search process. Establishing search parameters and guidelines as well as appropriate keywords and reference sites are particularly important. Although these actions may occur individually or in collaboration with others, nearly half of users exploited social interactions to facilitate this process.

During Search

Although search can be a cyclical process, the search stage in our model represents the active instantiation of representations or “encodons,” as part of a “data coverage” loop [30]. In other words, this is the stage where users engage in traditional information seeking [35] and foraging activities [27, 28]. We detail three types of searches based on Broder’s information needs (*transactional*, *navigational*, and *informational* [3]), drawing special attention to the social interactions.

Transactional Search

With a *transactional* search, users locate a source where they can subsequently perform a transaction or other “web-mediated activity” [3]. In our sample, this typically involved navigating to a website through a series of routine steps and requesting specific information such as driving directions, weather at a destination, movie listings, or data from a customer’s account. As an example, an ambulance chief for Acton Emergency Medical Services was required to include in a patient’s file the distance from the patient’s home to the hospital. To perform this routine and *transactional* search,

the chief navigated to MapQuest.com, entered the start and end locations, and retrieved the mileage information.

Although *transactional* searches (19/150 users, 12.7%) were less common in our sample than *navigational* or *informational* ones, over a third involved pre-search social interactions (8 of the 19 users, or 42.1%). In most cases, these interactions were themselves transactional—a necessary engagement to obtain details from the source of the request before proceeding with the actual search. For example, a placement advisor for a dental staffing agency spoke with a new applicant to get credentials and position availability. Following this interaction, the advisor went to a “specific program/site that lets you verify people’s licenses,” entered the applicant’s information, and verified that the license was current. Although *transactional* searches occasionally involved social exchanges in the pre-search phase, social interactions never occurred during the search act itself.

Navigational Search

During a *navigational* search, users perform a series of actions to identify content from a particular, often familiar, location. The content is often known in advance, or will be easily recognized once it is (re)discovered. For example, a hospital nurse found a drug listed in a patient’s medical chart, but was unfamiliar with how it was used. Before blindly administering it, the nurse decided to look up the drug on the NIH website, a familiar source often used to recover this type of information. As a result, the search act was *navigational*: the nurse logged onto Google, looked up the NIH’s web address, and then searched for the drug on the NIH website. The nurse reports: “I knew exactly where [the information] would be—just couldn’t recall what the answer was.”

Navigational searches occurred in 42/150 users (28.0%), and of these, nearly half involved pre-search interactions with others (20/42, or 47.6%). Of course, a number of users participated in discussions because they were asked to find specific information for a client or supervisor. In some cases, however, the interaction was necessary to establish the information need. The Membership chairman of the Embroiderers’ Guild of America recalled:

I spoke with the prospective member who needed the information [about membership dues and meeting times]. I did not contact anyone else regarding the searches I was about to do, since the information was needed immediately, and I knew what I was looking for. So I Googled to find the home site for the Embroiderers’ Guild. After finding the link, I went to the pages that had the information I needed.

As with *transactional* searches, information exchange occasionally occurred prior to the *navigational* steps but not during the search itself. Social interactions were primarily used for establishing guidelines and obtaining more information about the topic to be searched. However, six self-motivated *navigational* searchers (6/42 users, 14.3%) used social interactions as the first stage of the extended search, seeing if

others had answers or advice prior to searching. Generally speaking, it appears that users do not take advantage of social or expertise networks to facilitate *navigational* (or *transactional*) information retrieval, save for the occasional self-motivated searcher.

Informational Search

In comparison to *navigational* and *transactional* searches, social search may greatly improve tasks involving *informational* search, which is typically an exploratory process, combining foraging and sensemaking [27, 28], of searching for information that may or may not be familiar to the user. In our sample, over half of the search experiences were *informational* in nature, involving clear foraging and sensemaking processes (89/150 users, 59.3%).

Foraging. The basic “information assimilation” process described by Evans and Card [9] illustrates this early foraging phase where users search for information within a specific patch, followed by skimming, reading, and extracting information from source files. Throughout this process, users may update and shift their search representations [30] as they discover new items.

For example, an environmental engineer began searching online for a digital schematic of a storm-water pump while simultaneously browsing through printed materials to get “a better idea of what the tool is called.” This search was iteratively refined as the engineer encountered new information, first on metacrawler.com and then on Google, that allowed him to update his representation of the search space, or what might be called a “search schema.” He finally discovered a keyword combination that provided the desired results.

During this foraging and reformulation process, users may seek input from others for feedback and further refinement of their search [34]. A circulation clerk at a public library was asked to find the Cheetah Girls 2 (movie) soundtrack for her boss’s daughter. She began by using an internal search tool where she entered “Cheetah Girls 2,” although this failed to return the movie soundtrack. After the boss joined in the search process, the clerk reported:

We had to deduce a number of combinations in an attempt to obtain the item in question. We tried a number of ways to write Cheetah Girls, including hyphens and spelling out the number two. We even had to look up the actual product on Amazon.com for additional information.

The circulation clerk later commented that the boss was a useful resource during this process, providing “variations on the words in an attempt to solve the problem.” This example illustrates the well-known “vocabulary problem” of human-system communication [12].

Sensemaking. In fact, the foraging and sensemaking loops are tightly coupled. After an initial pass at foraging, users may identify preliminary “evidence files” [28] from which they can reflect and, if necessary, use to further modify their

search schema and query. This process of sensemaking may occur individually or through interactions with others. While an English Professor at Bradley University was preparing a lecture on a poem by Robert Frost, he “wanted to look up information on Frost himself as well as any information...specifically about the poem.” To do so, he recalled:

I went to msn.com and typed in “Robert Frost” first. I found several websites that were helpful, then copied and pasted information into the Word document I had already set up. [Then] I typed “Frost + Out Out.” [Finally] I went back to the Word document and began sorting my information. I placed like information together; summarized some information; deleted irrelevant information; and paraphrased other information.

Not only did the professor’s foraging involve a series of search queries, he engaged in a classic sensemaking process. He reflected on the material he collected as he summarized and paraphrased the information.

Although the professor performed “information assimilation” by himself, the sensemaking process could include social interactions. As an example, a programmer from Intuit was searching for a file transfer protocol (FTP) application programming interface (API) built in core JAVA. He began by brainstorming with the project’s technical architect. Then he performed a search online, followed by “another round of discussion with the technical architect” on whether the API he found (the “evidence file”) would be sufficient for their purposes. In this way, social input and exchange augmented the programmer’s sensemaking process.

Notably, many *informational* searchers engaged with others both before and during the search (36 of the 89 users, or 40.4%), even though our survey did not ask for social incidents or explicitly collaborative search acts. Social exchanges were useful for a variety of reasons: to establish search guidelines (to obtain “the exact specifics of the product”); to seek coworkers’ opinions (e.g., a behavioral consultant first spoke to “other consultants for advice and ideas for a novel way of teaching colors” to autistic children); to brainstorm (“through instant messenger...to confer with my group”); or to improve search schemas (“to know what kind of material would be useful” and “for wider search options and different opinions”). Moreover, a majority of users who engaged with others in this stage were self-motivated (22 of the 36 users, or 61.1%).

In summary, these accounts of search activities, including the role of social interactions during foraging and sensemaking, suggest that social search could be immensely valuable for *informational* searches, and for self-motivated users. It appears there is a demand and a desire for social inputs where the information need is undeveloped or poorly specified.

After Search

Following the active search phase, an “end product” is often obtained [34]. This target information—new IRS mileage

reimbursement rates, next week’s weather forecast, or the latest news about autism—may then be “acted” upon through *organization* and/or *distribution* [9].

Organizing Information

The organization of material includes saving or augmenting new content, bookmarking a webpage, or creating a new document or presentation. Pirolli and Card referred to this process as *schematizing*, where raw evidence is organized and “represented in some schematic way” [28].

Surprisingly, nearly half of our users (71/150, or 47.3%) organized their end products in some fashion. One real estate agent printed and reviewed the results of a search (public records of a property owner) before “giving them to an attorney for legal inspection.” The president of a design company bookmarked online articles about web mashups to read later in the week. And the manager of a data center created a PowerPoint presentation (for training subordinates) after finding technical information on cooling towers through a Google search. In each case, users created artifacts based on their search products that served to organize, save, or synthesize important information.

Distributing Information

Such organizational acts additionally served to distribute the search products to others (e.g., documents and presentations were created with the intention of sharing with colleagues). Pirolli and Card observed that the end products of a search may be delivered to an audience as a “presentation or publication of a case” [28]. Evans and Card also remarked that users serve as information filters for others through their organizational, and consequent distributional, acts of bookmarking, tagging, or annotating items [9].

In our sample, two thirds of respondents distributed end products either to others or to themselves for retrieval at a future date (101/150 users, 67.3%). Of these, 88 shared information with others (87.1%), typically face-to-face or verbally over the phone: For example, a floral designer relayed information about local spring blooming flowers to a bride-to-be. 15 users “shared” the content with themselves by printing out documents or bookmarking websites (14.9%), important actions for re-accessing and re-finding information in the future [18]. The IT Director for an automotive sales group reported:

I went to google.com, typed in Customer Relationship Management Software Solutions, then searched within results for Automotive Dealerships. I clicked several links and proceeded to read up on several programs that are available. Afterwards, I saved several links that I will be going back to and researching further. Once I am done with my research, I will put together a comparison of my top 3 choices and present it to the owner.

Additionally, two users reported sharing information “simultaneously for themselves and others” [9]. In one case, the manager of a retail store searched for an email containing a PDF memo describing the correct method of shipping goods

to customers. He “printed the PDF for future reference” and subsequently explained the procedure to his employees, requiring each to read and sign a copy.

The majority of information was distributed within a small radius of collaboration, to proximate others (86 of 101 distributing users, or 85.1%). In other words, most shared information with close friends or colleagues—people who had requested the information previously, whom the searchers thought would have an interest in it, or from whom searchers wanted to get feedback. Only two users found information that they shared to public (largely unknown) audiences: the professor from Bradley University giving the Robert Frost lecture and a graphic designer who maintains the website for a small bookstore. The graphic designer searched on Google and the New York Times online to find the current best selling mystery novels, and subsequently edited the bookstore’s public website to reflect the up-to-date information.

Motivation for Distribution

It is interesting to consider the reasons why searchers chose to distribute their end products. Items saved for future selves were all intended to be revisited at a later date. Information shared with public audiences (from only two examples in our sample) was essentially required of the users as an extension of their natural work.

Reasons for distributing information to close friends and colleagues, however, were quite varied. In fact, there are interesting differences between externally- and self-motivated searchers:

Search acts motivated by external requests for information resulted in sharing 83.0% of the time (39 of 47 users who received external prompts). Of these, nearly all sharing was out of obligation. Users were required to report on product availability to customer requests, provide details on upcoming conferences and workshops to the CEO, or send a map to a friend because he had asked for help. Only five of the 39 users shared information because they thought others would find it interesting.

In contrast, self-initiated searches resulted in post-search sharing slightly less frequently (49 of 103 self-motivated users, or 47.6%), but for a greater variety of reasons. Some still distributed information out of obligation, but many shared because they thought others would find the information interesting. They also shared information to get feedback from friends and colleagues: to make sure the information was accurate and valid, or to see if the search should be refined and repeated. For example, an audio transcriptionist, unsure of the spelling of one person’s name, performed a sound-alike search on Google before sharing the translation over instant messenger for colleagues to verify.

Activities reported in this phase suggest that post-search actions are important for organizing, reflecting upon, and distributing search products. Social interactions are still important even after the primary search act: to share information with specific individuals, and to obtain feedback and vali-

dation on search results. Anecdotes from our data indicate the need for post-search organizational tools for distributing information to interested parties, especially to close friends and colleagues.

DISCUSSION

Our results support recent findings [25, 34] that social interactions play a key role throughout the search process. Over one-third of our sample engaged socially *both* before and after their search; and two-thirds interacted with others at some point during the course of searching. Importantly, the majority of social interactions we observed across our 150 unique search events were accomplished through real-world interactions, emphasizing the social importance of others, even where digital tools do not directly support user needs.

Our user sample and findings are complementary to related work on collaborative information seeking [25]. While Morris’ study [25] examined active cooperation, we observed fewer instances of co-located, synchronous search. Occasionally users talked face-to-face with coworkers and supervisors, but often information was exchanged through email, phone calls, or instant messaging (remote interactions) and nearly always before or after the search act (database query) itself.

From the results of this study, we can begin to understand the contextualized role of social search in daily life: where in the search process users interact, why they socialize, and subsequently how these social interactions might improve the search process. We can then use the new understanding to create new social search tools.

Design Suggestions

In the following section, we review the social elements within the model and suggest a few design principles that might support social sharing behaviors.

Before search. As shown, information exchanges occurred through social interactions in each phase of our model. Even while formulating representations and search schemas prior to searching, nearly half (42.7%) of users talked with friends, colleagues, and clients as primary sources of information. These interactions allowed searchers to receive clarifications and guidelines on the task, seek advice, suggestions, keywords, and URLs, and at times, explore their existing social networks for the necessary information first.

Design suggestion. Social search may be useful here, especially for self-motivated searchers looking for information in domains with which they are not familiar (*informational* search). These users want to gather as much information as possible before the actual investigation to maximize the cost/benefit tradeoff associated with searching. Software tools could support this need by exploiting online social or expertise networks. This might include instant messaging access to one’s personal connections alongside the MSN, Google, or Yahoo! search box. Or it might exploit a website’s existing community to reveal domain-specific experts who would be willing to advise searchers. In other words, a

method for explicitly or implicitly making available knowledge from single individuals or aggregated social networks may help users prepare for their search tasks.

During search. During the main search act, our users engaged with others during foraging and sensemaking to refine their query and get feedback on their preliminary results. Here again, social search might help individuals during the active investigation of new material.

Design suggestion. Since search websites theoretically maintain a history of users' query terms, hit results, and search trails, a framework exists for exploiting social and expert resources to augment search. Put another way, the experiences of other users could be aggregated and presented to guide new searchers through a search episode. For example, the user's initial query reveals some clues about the search topic. Presenting expert tag clouds of semantically related concepts may help users refine their search while they are actively foraging and browsing for information. Alternatively, the website could display related and successful keyword combinations or search trails from previous users that may provide high-level feedback on the general search topic. Related work in the social recommendation literature has explored personalizing search by displaying *related queries* from a community's collective history [13] or from like-minded searchers [32].

After search. Post-search sharing was fairly prevalent. Over half the users in our sample engaged with others after their end product was located (88/150, or 58.7%). Sometimes information was shared out of obligation or because others might find it interesting. At other times, searchers wanted additional feedback or advice, especially for self-motivated, *informational* inquiries. Moreover, several users saved information for their future selves.

Design suggestion. In all cases, users may benefit from having features built into the web interface that facilitate post-search communication and sharing, such as email, instant messaging, bookmarking, or tagging. The Spartag.us system [17] currently supports low-cost tagging of webpages, and subsequent storage of tagged content in personal notebooks. Such data can then be aggregated and shared publicly as collective tag clouds of related information topics (from a personal social network or the extended community), serving as ambient, implicit, and socially-generated feedback.

Limitations of the Study

Despite the surprisingly thorough self-reported search experiences we collected, our study faces several limitations based on our sampling technique. Mechanical Turk is a relatively new tool for academic research, although other researchers have reported on the utility of it [19]. Regardless, our sample was anonymous, raising several questions about how user demographics might affect search behavior. For example, we do not take into account socio-economic status, technical and computer literacy, or individual personality traits that could influence search experiences or users' ability and inclination to access social resources.

There are additional limitations with the critical incident technique. Each user only recounted one search act. While we prompted users for their latest search experience, the episodes "fresh in their minds" may have been salient or noteworthy, not necessarily the most recent. One way of overcoming these concerns is through a large sampling where individual (and conspicuous) differences become attenuated. Nevertheless, self-reported critical incidents surely vary from daily, typical, or less significant search events.

Finally, our survey explored user interactions with the social and technical environments in which they are currently embedded. In fact, this is both a limitation and a benefit. On the one hand, many of our users did not have tools that effectively supported social search; therefore, we were only able to measure existing behaviors, with users' current tools and practices. On the other hand, this revealed users' natural inclination to interact with others, suggesting where in the search process both explicitly and implicitly shared information may be valuable to individual searchers.

CONCLUSION

In this paper, we presented a critical incident analysis of 150 individual search episodes. Our main contribution is that we have integrated our findings with models from previous work in sensemaking and information seeking behavior to present a canonical model of social search. We also offered some design suggestions for supporting social behaviors throughout the search process.

Our results and analysis demonstrated that users have a strong social inclination throughout the search process, interacting with others for reasons ranging from obligation to curiosity. Self-motivated searchers and users conducting *informational* searches provided the most compelling cases for social support during search.

Although the general term "social search" may not have a single precise definition, we hope that the canonical model of search we presented offers suggestions for a notion of social search that may support web-based information seeking with both explicit and implicit social processes. Most importantly, we hope that our work encourages fellow researchers to explore and expand "social search," so that future searchers can work more productively.

REFERENCES

1. S. Amershi and M. R. Morris. CoSearch: A system for colocated collaborative web search. In *Proc. CHI'08*, ACM Press, pages 1647–1656, 2008.
2. M. Bates. Toward an integrated model of information seeking and searching. In *The 4th Int'l Conf. on Information Needs, Seeking and Use in Different Contexts (Keynote address)*, September 11–13 2002.
3. A. Broder. A taxonomy of web search. In *ACM SIGIR Forum*, pages 3–10, 2002.
4. J. Castillo, H. R. Hartson, and D. Hix. Remote usability evaluation: Can users report their own critical incidents. In *Proc. CHI'98*, ACM Press, pages 253–254, 1998.

5. L. D. Catledge and J. E. Pitkow. Characterizing browsing strategies in the World-Wide Web. *Computer Networks and ISDN Systems*, 27(6):1065–1073, 1995.
6. C. W. Choo, B. Detlor, and D. Turnbull. Information seeking on the web: An integrated model of browsing and searching. In *Proc. ASIS'99, ACM Press*, pages 3–16, 1999.
7. A. Cockburn and S. Jones. Which way now? Analysing and easing inadequacies in WWW navigation. *International Journal of Human Computer Studies*, 45:105–129, 1996.
8. D. Ellis. A behavioural model for information retrieval system design. *Journal of Information Science*, 15(4/5):237–247, 1989.
9. B. M. Evans and S. K. Card. Augmented information assimilation: Social and algorithmic web aids for the information long tail. In *Proc. CHI'08, ACM Press*, pages 989–998, 2008.
10. J. Flanagan. The critical incident technique. *Psychological Bulletin*, 51:327–358, 1954.
11. E. A. Fox, D. Hix, L. T. Nowell, D. J. Brueni, W. C. Wake, L. S. Heath, and D. Rao. Users, user interfaces, and objects—envision, a digital library. *JASIS*, 44(8):480–491, 1993.
12. G. W. Furnas, T. K. Landauer, L. M. Gomez, and S. T. Dumais. The vocabulary problem in human-system communication. *Commun. ACM*, 30(11):964–971, 1987.
13. N. S. Glance. Community search assistant. In *Proc. Intelligent User Interfaces*, pages 91–96, 2001.
14. D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Commun. ACM*, 35(12):61–70, 1992.
15. P. Heymann, G. Koutrika, and H. Garcia-Molina. Can social bookmarking improve web search? In *Proc. WSDM'08, ACM Press*, pages 195–206, 2008.
16. P. I. Hofgesang. Web personalisation through incremental individual profiling and support-based user segmentation. In *Proc. IEEE'07, ACM Press*, pages 111–120, 2007.
17. L. Hong, E. H. Chi, R. Budiou, P. Pirolli, and L. Nelson. Spartag.us: a low cost tagging system for foraging of web content. In *Proc AVI '08: ACM Press*, pages 65–72, 2008.
18. W. Jones and J. Teevan. *Personal Information Management*. University of Washington Press, Seattle, 2007.
19. A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with Mechanical Turk. In *Proc. CHI'08, ACM Press*, pages 453–456, 2008.
20. J. A. Konstan, B. N. Miller, D. Maltz, J. L. Herlocker, L. R. Gordon, and J. Riedl. GroupLens: Applying collaborative filtering to Usenet news. *Commun. ACM*, 40(3):77–87, 1997.
21. C. C. Kuhlthau. Inside the search process: Information seeking from the user's perspective. *JASIS*, 42(5):361–371, 1991.
22. K. J. Lee. What goes around comes around: An analysis of del.icio.us as social space. In *Proc. CSCW'06, ACM Press*, pages 191–194, 2006.
23. G. Marchionini. *Information Seeking in Electronic Environments*. Cambridge University Press, Cambridge, UK, 1995.
24. D. R. Millen, J. Feinberg, and B. Kerr. Dogear: Social bookmarking in the enterprise. In *Proc. CHI'06, ACM Press*, pages 111–120, 2006.
25. M. Morris. A survey of collaborative web search practices. In *Proc. CHI'08, ACM Press*, pages 1657–1660, 2008.
26. M. R. Morris and E. Horvitz. SearchTogether: An interface for collaborative web search. In *Proc. UIST'07, ACM Press*, pages 3–12, 2007.
27. P. Pirolli and S. K. Card. Information foraging. *Psychological Review*, 106:643–675, 1999.
28. P. Pirolli and S. K. Card. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proc. Int'l Conf. on Intelligence Analysis'05*, pages 2–4, 2005.
29. T. Rodden. A survey of CSCW systems. *Interacting with Computers*, 3(3):319–354, 1991.
30. D. M. Russell, M. J. Stefik, P. Pirolli, and S. K. Card. The cost structure of sensemaking. In *Proc. CHI'93, ACM Press*, pages 269–276, 1993.
31. J. Shepherd. *The Information Needs and Information Seeking Behaviour of Polytechnic History Lecturers*. British Library, London, 1983.
32. B. Smyth. A community-based approach to personalizing web search. *Computer*, 40(8):42–50, 2007.
33. R. S. Taylor. Question-negotiation and information seeking in libraries. *College & Research Libraries*, 29(3):178–194, 1968.
34. M. B. Twidale, D. M. Nichols, and C. D. Paice. Browsing is a collaborative process. *Information Processing and Management*, 33(6):761–783, 1997.
35. T. Wilson. On user studies and information needs. *Journal of Documentation*, 37(1):3–15, 1981.
36. Y. Yanbe, A. Jatowt, S. Nakamura, and K. Tanaka. Can social bookmarking enhance search in the web? In *Proc. JCDL'07, ACM Press*, pages 107–116, 2007.