

Aspects of Augmented Social Cognition: Social Information Foraging and Social Search

Ed H. Chi, Peter Pirolli, and Shyong K. Lam

Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304 USA
echi@parc.com, pirolli@parc.com

Abstract. In this paper, we summarized recent work in modeling how users socially forage and search for information. One way to bridge between different communities of users is to diversify their information sources. This can be done using not only old mechanisms such as email, instant messages, newsgroups and bulletin boards, but also new ones such as wikis, blogs, social tags, etc. How do users work with diverse hints from other foragers? How do interference effects change their strategies? How can we build tools that help users cooperatively search? We seek theories that might help us answer these questions, or at least point us toward the right directions.

Keywords: social processes, information search, information foraging.

1 Introduction

“There is a growing mountain of research. But there is increased evidence that we are being bogged down today as specialization extends. The investigator is staggered by the findings and conclusions of thousands of other workers - conclusions which he cannot find time to grasp, much less to remember, as they appear. Yet specialization becomes increasingly necessary for progress, and the effort to bridge between disciplines is correspondingly superficial.” -- V. Bush [6]

V. Bush envisioned a device he called the Memex that would allow scholars to forage through personal stores of multimedia documents, and to save traces of paths through content that could then be shared with other scholars as a way of communicating new findings [6]. The Memex was envisioned as a tool that would increase the capacity of individuals to attend to greater spans of emerging knowledge, and would increase the cooperative information sharing that Bush viewed as necessary to improvements in scientific discovery, which he expected to result in increased benefits to society. Bush’s vision was not only to improve the foraging ability of the individual user, but to also improve communication and collaboration. There appears to be a number of intuitions about why collaboration might improve productive work.

- (1) Like over-the-horizon radar, an individual analyst may receive information otherwise unseen because of the information flowing to him or her from a social network of collaborators.

- (2) Collectively, by arranging the spotlights of attention of individual users to insure maximum, exhaustive coverage of the available information, one can diminish the chances of failing to bring to light some crucial data that might otherwise be missed.
- (3) Coordinated teams of experts may be assembled in order to exploit years of specialized skill and relevant knowledge about background and precedents, in order to better recognize and interpret data and information.
- (4) Diversity of viewpoints can be brought to bear to provide mutually corrective forces to overcome the cognitive heuristics and biases that often create blindness to impending threats.

Each of these beneficial effects is crucial to overcoming the most frequent impediments to situation awareness and knowledge sharing [9]. Each of these effects has been the focus of some fields of research that we draw upon for theoretical guidance.

Despite a large amount of Computer Supported Cooperative Work research and literature in this area, significant progress still can be made. For example, what has been amazing is the amount of new social mechanisms invented on the Internet for social foraging. To some extent, the Web, blogs, email, internet groups, wikis, and other mundane technologies are all aimed at supporting cooperative information sharing and their success implies their effectiveness.

However, there is also a considerable body of evidence from social psychology and decision support systems that indicates that collaborative work often results in worse outcomes than working alone. Further, it is common knowledge that collaboration involves overhead costs that, for the individual, act as disincentives for collaborative activities. Despite the potentially negative features of collaboration, many new tools for social information foraging have emerged as disrupting technologies on the Internet including peer-to-peer information sharing, social search, and blogs. The aim of this paper is to develop theory, evaluation techniques, and disruptive technology to increase the benefits and reduce the costs of augmented social information foraging.

The major question confronting the field, therefore, is to not only understand how these mechanisms are different or the same as previous generation of tools, but also build models of how social foraging appears to be benefiting individuals, and how these models might in turn tell us how to build new collaboration tools.

In this article, we outline several explorations on models of social information foraging, and how some of these ideas might impact social search engines. We will touch upon general results concerning the costs and benefits of cooperative foraging, the effects of group diversity, and patterns of social structuring that are correlated with innovative discovery.

2 Costs and Benefits of Cooperative Foraging

2.1 Problems of Specialization, Bias Barriers, and Tools

Specialization is a natural consequence of too much public knowledge for the individual mind to comprehend. The worry is that knowledge specialization leads to

situations in which all the information required to make an important discovery is in the available record somewhere, but it is distributed across specialization boundaries with no single set of eyes in a position to see it all and make sense of it. In the information retrieval and library sciences, this is known as the *undiscovered public knowledge problem*. These problems may include hidden refutations, hidden cumulative strength of individually weak studies, or other hidden links in the logic of discovery.

Indeed, in the intelligence analysis community, since 9/11, there has been increasing focus on the need to reduce individual bias in analysis. Some barriers to information sharing has been systematically studied and publicly debated, including computer security issues, political turf concerns, and multiple intelligence agencies with differing sources of information.

While much of the problem remain organizational, analysts also lack the most basic tools for information sharing, such as the ability to identify the experts in a particular domain, search in a database that has been tagged with the most promising leads, or automatically download and attend to the most promising information sources. The tools are truly quite impoverished.

2.2 Studies of Existing Solutions

Existing research showed that the benefits of social information foraging can be systematically studied and understood. Cooperation may yield more benefits than simply making search more parallel and making it less prone to failure. Membership in a group provides actual or potential resources that can be utilized or mobilized to achieve individual goals. This is known as *social capital* [4,17].

Pirolli and Card [16] describe a business intelligence agency whose analysts were tasked to write monthly newsletters about core areas such as computer science or materials science. In addition to culling material for their own newsletters, analysts would also notice articles pertinent to the specialties of other analysts, and would have such articles copied and routed to the appropriate specialist. An analyst would typically receive about 6–12 relevant articles per month from other analysts, at very little cost. The general belief of the analysts was that such cooperation enhanced the individuals' search capabilities, and reduced the risk of missing something relevant to a specialty area that had emerged in a non-specialty publication. The social network amongst the analysts is precisely their social capital for solving complex problems of their own.

Interestingly, recent recommender systems use an implicit form of these social networks to provide social capital to their users. Recommender systems exploit social information to make recommendations (documents, movies, music) to individuals. These include collaborative filtering systems [10] in which people typically indicate their preference for items in some way (e.g., by rating things such as books), and they receive recommendations based on the preferences of others with similar tastes.

2.3 Effects of Group Diversity

Existing research studied the effects of group diversity on cooperative information foraging, as well as the theory that people who provide brokerage of ideas across social clusters are often in position to make valuable novel discoveries.

Organization and management studies [8] suggest that effective work groups are ones that share information and know-how with external members, and that effectiveness is improved by *structural diversity* of the group. *Structural diversity* is variability in features of the group that expose members to different sources of task information, know-how, and feedback. Such features include geographic locations, functional assignments, number of managers to whom members report, and number of business units associated with the group.

The findings of Cummings [8] are consistent with the theory of social structural holes (*structural holes theory*) proposed by Burt [5]. Typically, such social networks of information flow will contain densely connected clusters. The sparse linkages between such clusters constitute *structural holes*. People who bridge such structural holes have an advantage of exposure to greater diversity of information and know-how, and *brokerage* across structural holes becomes a form of *social capital* that translates into the discovery of greater amounts of useful, productive knowledge.

2.4 Epidemiological Model

One of the exciting prospects for the study of social information foraging is improved ability to find and measure social networks using on-line resources. For instance, it appears that e-mail flow and Web links among personal home pages provide data that can be used to accurately construct social networks [12] and to study information flow.

Recent researchers have tried to model the information flow in a social network of information foragers. One set of models for such flow is based on epidemiological models [12,21]. The model of Wu et al. [21] is based on the application of random graph theory [3,14] to the epidemiological models of the spread of diseases. The key idea is that information spreads from person to person in ways analogous to the spread of an infection from person to person. Wu et al. [21] assume a social network whose nodes have an outdegree distribution that follows a scale-free power law. This means that there are a few members of the social network who have a very large number of social connections, but the typical number of connections is small (and less than the arithmetic mean).

Information tends to be selectively passed on to people that the host believes will find the information useful. Individuals tend to form social associations based on the similarity of their characteristics (known as *homophily*), and Wu et al. assume that the similarity of two people diminishes as a function of their distance in a social network (a common observation in sociometrics). One is lead to infer that transmission probabilities for information should decay as a function of network distance from the host source.

The sum of all the predicted infected individuals across all distances from the host provides a measure of the average total size of an “outbreak” (the number of individuals “infected” by the information). Interestingly, Wu et al. showed that the assumption of decay in transmissibility with network distance typically leads to outbreaks of finite size. Empirical studies of samples of email, and numerical simulations using parameters based on observation and a social network graph based on email patterns suggest that outbreaks in an organization of about 7000 people is typically limited to under 50 people.

2.5 Social Information Foraging Model

Pirolli [15] derived a basic model of social information foraging based on the Information Foraging Theory. The basic model derivation is too complex to present in detail here. Instead, we will provide a brief summary here. A key insight of the theory is that there is an intricate relationship between the interference costs of collaboration and the benefits of the diversity of view of point.

The model seeks to understand the effect of the diversity of information foragers and their hints to each other, and how this diversity directly affects the size of cooperating groups.




Fig. 1. Key predictions of an Augmented Social Cognition theory. (a, top left): collaboration in a diverse group improves chances of discovery; (b, top right): cooperation improve chances of discovering the more important facts; (c, lower left): cost and benefit of collaboration determine effective group size; (d, lower right): diversity decreases confirmation bias.

Figure 1 summarizes key predictions of the theory. This model assumes that there are some processing steps that are required to find useful patches of information and this can be modeled as a Poisson process. Cooperating information hints from other foragers can be modeled as the number of search steps that are eliminated due to a distinct hint.

The model suggests that so long as the diversity of agents increases with group size, then the size of a group increases the overall power of cooperative discovery. As individual foragers increase the diversity of their cooperating contacts they will improve in performance.

Figure 1a shows how collaboration with a diverse group improves the rate of return to the individual analyst. Being embedded in a cooperative social network of analysts

provides the individual analysts with the ability to explore more of the space of information more rapidly than could be done alone -like an over-the-horizon radar [5].

Figure 1b shows how the theory predicts how cooperation among a set of analysts improves the probability of making important (or difficult) discoveries. This means that unseen patterns, connections, inferences that are latent in the raw data may come to light as the number of diverse but cooperative analysts increases [19].

The model also showed that, assuming there are interference costs that grow with the number of members in a group, these interference costs may reduce the social capital of group foraging down to the level of individual foraging, at which point one expects an equilibrium group size.

Figure 1c illustrates that a combination of expected benefits and interference costs to an individual determines the effective size of a group. People typically join a group only if the benefits (to the individual) outweigh the costs of cooperation [7].

Figure 1c also illustrates how a reduction in the costs of collaboration is predicted to increase the effective size of the collaborating group. Lowering the costs of collaboration leads naturally to larger groups of collaborating analysts.

Figure 1d illustrates the effects of diverse heuristics and biases in a cooperative group on the mitigation of confirmation bias [11,20]. Diversified groups show greater mitigation of confirmation biases than homogenous groups [2].

3 Implications from the Models

The effects of group diversity and the social foraging models tell us that the structure, composition, and the size of the social network are extremely important.

In terms of **composition**, Cummings research show that the diversity of hints is important in bringing fresh information into a group.

In terms of **structure**, the findings from Burt suggest that if we can locate structural holes in the social network, we will be able to bridge across these holes by enabling information to directly flow across more efficiently. One way this can be operationalized is by identifying two workgroups that has tight information sharing, and finding the highest connected members from each group to inform each other.

In terms of **size**, the results from Pirolli and Wu et al. [21] suggest that there is an optimal size of the information outbreak for a given situation. Too large of an outbreak will generate interference, while too small of an outbreak causes information to be disseminated too inefficiently.

The crushing pressure of information overload leads to specialization by people in fields of knowledge discovery. Overspecialization may lead to increases in undiscovered public knowledge [19]—a failure for any one mind to grasp and connect all the dots. Information sharing is usually recognized as a strategy for extending the grasp of the solitary mind across specializations, to reduce the risk of failing to make discoveries implicit in the existing collections of data. To the extent that individual members of a core specialty can devote some effort to exploring related peripheral specialties, and sharing possible leads with others, then one might expect the group to perform more effectively.

There are still other open questions, such as issues surrounding trust and privacy. An information sharing tool should enable further enhancement of trust between users

or at least not diminish it. Much more difficult to design for, the privacy of the user must be protected. These are open future research issues.

4 Social Search

We have been interested in applying these findings and implications to the design of a social search engine. In considering the implications of composition and structure, we first wanted to understand the variety of information that different people look at and if there is enough structure and composition in these groups to cover many topics but with enough commonalities to infer possible structural holes to bridge.

We wanted to do this analysis to understand how well people can collaborate in their search process. We do this by first comparing the history and bookmarks of some real users. We performed some analysis of user browsing histories and their bookmarks, trying to understand to what extent their information sources are diverse but yet overlap. Given the vastness of the web, it turns out that there is a lot more overlap than we might otherwise have expected.

USERS	3x3eyes	biotech	creap	dinopixel	ed	elindbloom	fabio	ftg	gbuser6	guruzilla	jiim	jugglebird	lowellk	rainer	terrorism
3x3eyes	-	0.00%	2.13%	5.80%	0.02%	2.02%	0.06%	0.46%	0.01%	0.52%	1.07%	1.77%	1.71%	0.61%	0.00%
biotech	0.00%	-	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
creap	2.13%	0.00%	-	4.66%	0.01%	2.10%	0.00%	0.08%	0.00%	0.04%	0.54%	1.60%	0.93%	0.40%	0.00%
dinopixel	6.80%	0.00%	4.66%	-	0.01%	6.13%	0.02%	0.59%	0.00%	0.13%	1.18%	2.41%	2.63%	0.34%	0.00%
ed	0.02%	0.00%	0.01%	0.01%	-	0.01%	0.72%	0.04%	0.03%	0.00%	0.03%	0.03%	0.01%	0.00%	0.00%
elindbloom	2.02%	0.00%	2.10%	6.13%	0.01%	-	0.00%	0.20%	0.00%	0.00%	0.60%	1.90%	0.78%	0.16%	0.00%
fabio	0.06%	0.00%	0.00%	0.02%	0.72%	0.00%	-	0.00%	0.08%	0.02%	0.02%	0.00%	0.02%	0.00%	0.00%
ftg	0.46%	0.00%	0.08%	0.59%	0.04%	0.20%	0.00%	-	0.00%	0.28%	0.35%	0.55%	0.18%	0.27%	0.00%
gbuser6	0.01%	0.00%	0.00%	0.00%	0.03%	0.00%	0.08%	0.00%	-	0.00%	0.03%	0.00%	0.06%	0.00%	0.00%
guruzilla	0.52%	0.00%	0.04%	0.13%	0.00%	0.00%	0.02%	0.28%	0.00%	-	0.14%	0.19%	0.03%	0.09%	0.00%
jiim	1.07%	0.00%	0.54%	1.18%	0.03%	0.60%	0.02%	0.35%	0.03%	0.14%	-	0.60%	0.45%	0.21%	0.00%
jugglebird	1.77%	0.00%	1.60%	2.41%	0.03%	1.90%	0.00%	0.55%	0.00%	0.19%	0.60%	-	0.41%	0.37%	0.00%
lowellk	1.71%	0.00%	0.93%	2.63%	0.01%	0.78%	0.02%	0.18%	0.06%	0.03%	0.45%	0.41%	-	0.23%	0.00%
rainer	0.61%	0.00%	0.40%	0.34%	0.00%	0.16%	0.00%	0.27%	0.00%	0.09%	0.21%	0.37%	0.23%	-	0.00%
terrorism	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	-

Fig. 2. Browsing history overlap study show (1) overlap between random del.icio.us users is surprisingly high; (2) Experts (such as gbuser6) may be foraging in entirely different areas of the Web than curated collections (terrorism and biotech)

The figure above depicts our analysis results. The user profiles “ed”, “fabio”, and “gbuser6” (highlighted in yellow and orange) are history traces we obtained by directly examining the browser histories. They all contain 3-6+ months of data. “Biotech” and “terrorism” (highlighted in green) are bookmarks from the Open Directory Project curated listings. Finally, we also examined ten bookmark histories from a social bookmarking site called del.icio.us (users with no highlights). The profiles ranged from 831 URLs to over 9211 (Mean=2243). There are some striking results:

- Ed and Fabio's profiles only overlapped with each other strongly, but not with other profiles. This is interesting, especially since they were co-workers.
- Gbuser6's profile is a real intelligence analyst profile that mainly consists of terrorism and biological weapon URLs. This profile did not overlap with the curated Biotech and Terrorism profiles, and in fact showed very little overlap with any other profiles, despite having 9211 in his profile. This indicates that he used very different information sources. Biotech and Terrorism also had no overlap with other profiles.
- The del.icio.us users overlap strongly with each other, forming their own little community. This is true despite the fact that we have no evidence that these users know each other at all. In particular, dinopixel overlapped with 3x3eyes by 6.8%, which means roughly every 1 in 15 URLs were viewed by both of them! One possible explanation is that they are all uber-geek technologists who tend to use similar information sources on the net such as Slashdot.

What these results seem to suggest are:

- Users with similar interests tend to skim the same surfaces on the web, if they might be using similar news sources in their everyday browsing activities. Co-workers and communities are most likely to skim the same surfaces, but they are not always getting to the same pages. Given the vastness and diversity of the web, people actually get pretty close to each other, and could easily inform each other of sources that they should not have missed.
- People in related fields but not in the same community do not skim the same surfaces on the web, and are prime candidates for more information sharing. These results seem to confirm that structural holes do exist. Gbuser6 is well poised to benefit from both biotech and terrorism profiles, and serve as a bridge between the two communities.
- The similarity between profiles could be used as input to a social search engine. It could be used to find experts in particular areas, share new information sources, and bridge across structural holes. We could use the overlap between users to help find related users, related queries, and result pages.

We are currently constructing foraging agents that build upon personalized search technology developed at PARC. A preliminary version of this system called Proximal Search. Foraging agents can be seeded with starting points and search/browsing histories extracted from the individual users and their knowledge resources. Using social network search one may gain access to the results produced by foraging agents working on behalf of experts, or users with different sensemaking viewpoints.

In this system, context focused foraging is achieved by a focused crawler that has the goal of seeking out and indexing parts of the Web that are highly relevant to a particular user, based on the search history, tags, and other knowledge resources. Such search agents can find deep parts of the Web that are unindexed by common search engines, and provide search results that have higher precision and recall for the specific interests of the user.

One aspect is the ability to exploit social connections to gain access to expertise outside of one's own. Membership in a group provides actual or potential resources that can be utilized or mobilized to achieve individual goals. Cooperation, such as collaborative tagging, may yield more benefits than simply making search more

parallel and making it less prone to failure. Systems such as Dogear can be combined with social search [13].

Expert location using social network search can be used to guide peer-to-peer meta search to retrieve from expert search indexes produced by those expert's agents. Professional knowledge workers, such as scientists, who are engaged in making sense of information and making discoveries, typically search through their networks of social relations in order to find experts on matters on which they themselves are unfamiliar. Searching for expertise in a social network in the electronic world is known as the *expert location problem* [1,22]. A variety of search algorithms have been explored in simulation studies [22], but remain to be implemented.

5 Concluding Remarks

In this short paper, we summarized recent work in the understanding of social information foraging, and models of how users socially foraging with diverse hints.

Information overload results in knowledge specialization, which increases the need to communicate between analysts with diverse viewpoints. Communications lead to interference costs, which can be mitigated by technology. Less interference costs will enable more effective collaboration and cooperation amongst analysts.

In communities of practice that depend on foraging in overly rich information environments, there appears to be pressure to self-organize into a balance of some division of labor, plus some degree of cooperation. It appears that the division of labor is necessary because of the limits of human attention, but some investment in cooperation can lead to increased returns and less risk of missing something important. The power of cooperation is related to the amount of diversity of the information foragers. Greater diversity leads to greater returns for the group and the individual. This is related to the notion that brokerage (diverse social contacts) provides social capital, and there is evidence that brokers in the flow of information are more likely to be sources of innovative discoveries. Although there are benefits to cooperation, those benefits trade against interference effects.

Recent emerging tools such as blogs, podcasts, and other internet groups serve to bridge across structural holes, but they must at the same time eliminate their potential for interference effects and overwhelm the users.

One way to bridge between different communities of users is to diversify their information sources. We have been examining the possibility of using a collaborative search engine to achieve this effect. Initial studies suggest that a social search engine build with the understanding of the implications of composition, structure, and size of communities should be able to inform users of new potential information sources without creating much additional interference overhead for the users.

References

1. Adamic, L.A., Adar, E.: How to search a social network. *Social Networks* 27(3), 187–203 (2005)
2. Billman, D., Convertino, G., Shrager, J., Massar, J.P., Pirolli, P.: Collaborative intelligence analysis with cache and its effects on information gathering and cognitive bias, Human Computer Interaction Consortium Workshop. Boulder, CO (2006)

3. Bollobas, B.: *Random graphs*. Academic Press, London (1985)
4. Bourdieu, P.: The forms of capital. In: Richardson, J.G. (ed.) *Handbook of theory and research in the sociology of education*, Greenwald Press, New York (1986)
5. Burt, R.S.: Structural holes and good ideas. *American Journal of Sociology* 110(2), 349–399 (2004)
6. Bush, V.: As we may think. *Atlantic Monthly*, 176, 101–108 (1945)
7. Clark, C.W., Mangel, M.: The evolutionary advantages of group foraging. *Theoretical Population Biology* 30(1), 45–75 (1986)
8. Cummings, J.N.: Work groups, structural diversity, and knowledge sharing in a global organization. *Management Science* 50(3), 352–364 (2004)
9. Grabo, C.M.: *Anticipating surprise: Analysis for strategic warning*. Joint Military Intelligence College, Washington, DC (2002)
10. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.: Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems* 22(1), 5–53 (2004)
11. Heuer, R.J.: *Psychology of intelligence analysis*. Center for the Study of Intelligence, Washington, DC (1999)
12. Huberman, B.A., Adamic, L.A.: Information dynamics in a networked world. In: Ben-Naim, E., Frauenfelder, H., Toroczkai, Z. (eds.) *Complex networks*, Springer, Heidelberg (2004)
13. Millen, D.R., Feinberg, J., Kerr, B.: Dogear: Social bookmarking in the enterprise. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '06, Montréal, Québec, Canada, April 22 - 27, 2006, pp. 111–120. ACM Press, New York, NY (2006)
14. Newman, M.E.J.: Spread of epidemic disease on networks. *Physical Review E (Statistical, Nonlinear, and Soft. Matter Physics)* 66(1), 16128 (2002)
15. Pirolli, P.: Chapter 8: Social Information Foraging. *Information Foraging* (2007)
16. Pirolli, P., Card, S.K.: Information foraging. *Psychological Review* 106, 643–675 (1999)
17. Putnam, R.: *Bowling alone: The collapse and revival of american community*. Simon and Schuster, New York (2000)
18. Sandstrom, P.E.: Scholarly communication as a socioecological systems. *Scientometrics* 51(3), 573–605 (2001)
19. Swanson, D.R.: Undiscovered public knowledge. *The Library Quarterly* 56(2), 103–118 (1986)
20. Tversky, A., Kahneman, D.: Judgment under uncertainty: Heuristics and biases. *Science* 185, 1124–1131 (1974)
21. Wu, F., Huberman, B.A., Adamic, L.A., Tyler, J.R.: Information flow in social groups. *Physica A: Statistical and Theoretical Physics* 337(1-2), 327 (2004)
22. Zhang, J., Ackerman, M.S.: Searching for expertise in social networks: A simulation of potential strategies. In: *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, ACM Press, Sanibel Island, Florida, USA (2005)