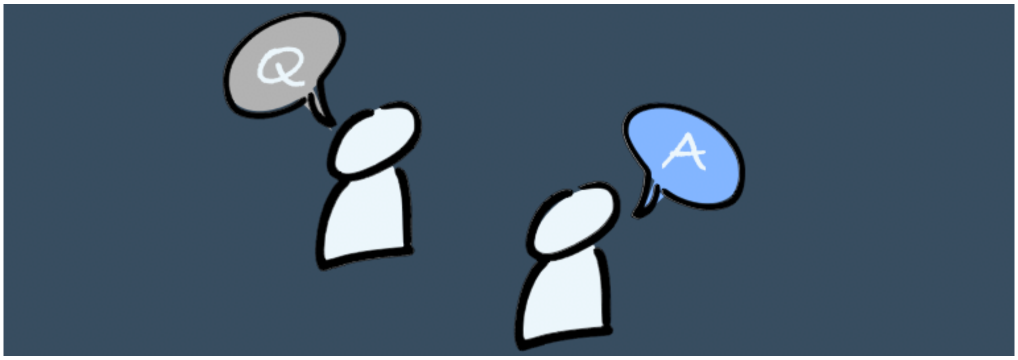

A NetLogo model of question-answering

Brynn Evans • March 19, 2009

1. Motivation

Search engines have grown quite sophisticated in the past decade. Algorithms have been refined to take into account semantics, user intentions, and broader social patterns on the web. Consequently, they serve many of our daily needs and we have a number of positive experiences with them.

Google (Yahoo, AOL, and others) don't always succeed in producing answers to our questions, however. Certain queries are better suited for web-based search (e.g., *how is Apple stock doing?*); others are not (e.g., *why is my software program crashing?*). Yet, the web is often our "go-to" source because it is convenient, familiar, and fairly reliable.

In physical settings such as libraries and offices, friends, colleagues, and neighbors are natural information resources—sometimes as a complement to or in place of impersonal sources (e.g., databases, journals). Social relationships are critical for acquiring new information, new solutions, and new perspectives (Granovetter 1973; Allen 1977; Burt 1992; Brown & Duguid 2000). They are known to provide guidance, advice, and assistance to each other (Bandura 1989; Hatch & Gardner 1993), to provide pointers to databases or other people (Fox et al. 1993), and even to act as memory aids (Karasavvidis 2002)—more obvious cognitive support structures.

We might naturally ask: could we make better use of our friends to help with search? Could social interactions be introduced to the search process to address users' unmet cognitive and informational needs? Recent work has shown that social interactions are both prevalent and useful during search by providing guidance, advice, assistance, and brainstorming opportunities for users (Evans & Chi 2008; Morris 2008; Golovchinsky et al. 2008). Current web facilities do a poor job of incorporating natural social interactions and inputs into the search process, although a number of new approaches to "social search" are emerging.

One approach to social search is a question-answering model, where a user poses a question and receives a reply from one or more friends. Yahoo! Answers and Live Q&A are popular web-services that provide question-answering platforms. Early academic work in this area includes Mark Ackerman's Answer Garden (Ackerman 1998). In the present model, my goal is to depict a canonical question-answering system. The social factors modeled within have been informed by the literature and by current behavioral trends on the web, but there are certainly other important factors to consider in such a complex, socially- and cognitively-distributed ecosystem.

In particular, this model is designed to see how likely it is that a user would receive an answer to her query after posing a question to her social network. "Social network" is used both broadly in the colloquial sense and specifically to refer to individual people within the network. For example, the model below distinguishes between social networks that a user accesses through the web and the people who are available in the "real world." The likelihood of getting an answer will depend on who the user has access to face-to-face, how knowledgeable they are about the topic, the size and diversity of online social networks, how difficult the topic is to answer, and how knowledgeable the user is about the topic. In practice, there are many other factors that would also matter, such as the user's relationship with people in the network and the last time they interacted, which are not modeled here.

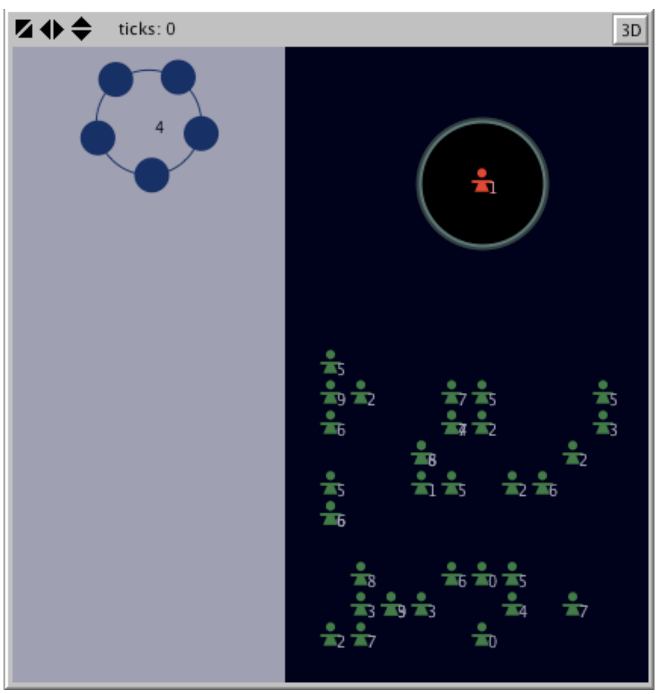

Figure 1. Screenshot of the setup of one trial in the model. The user is in red; friends are green; social networks are blue. The user poses a hard question [1] on this trial.

2. The model

The NetLogo world has three types of turtles: the user (red); real-life friends (green); online social networks (blue connected-dots) [Figure 1]. The gray background (left-most patches) represents cyberspace while dark blue background (right-most patches) represents physical space. On each trial, the user will pose a question and collect responses from real-world friends and online social networks in the hopes of getting an answer. The question type, knowledge of friends, knowledge of the user, and other factors are considered when computing success on each trial.

2.1 The user.

More specifically, each user has an attribute of knowledge (on a scale from 0–9) [not shown in Figure 1], and an easy (0) or hard (1) question to solve [shown in pink]. The easy—hard distinction is a simplification of two classes of common web queries. Easy questions are more simply fact-retrieval searches, where the goal is to find an optimal path to the requisite information or documents (White et al. 2007). For example:

- *How is Apple stock doing?*
- *What are the contraindications for Clonazepam?*
- *What is the local DMV's phone number?*

Hard questions tend to involve ill-structured problems with open-ended goals. Users may perform multiple queries to explore the problem space of the question, iterating on and reformulating the problem along the way. These processes are often more concerned with learning about a topic than about answering a specific question (Marchionini 2006). For example:

- *How do I export video in an uncompressed format from my Canon HD camcorder?*
- *What are the latest regulations in enlisting army personnel with GED credentials?*
- *How much would it cost to get our company t-shirts embroidered by an outside vendor?*

2.2 Friends.

By default, a random number of friends from 0–49 will be generated on each trial. The upper-limit of friends can be tweaked with the *maxFriends* slider. Each friend will have an attribute of knowledge (scale from 0–9), also randomly assigned on every trial. This score is supposed to represent their knowledge about the topic posed in the question. They will reply to a question if their knowledge is high enough [see Calculating replies].

2.3 Social networks.

Online social networks provide access to individuals as well. By default, the number of social networks is between zero and two (randomly generated on each trial), but the upper-limit can be modified by the *maxSocialNetworks* slider. Each social network has an attribute of diversity (scale from 0–9), which is supposed to represent diversity of knowledge within the network. When networks are well represented, they have high diversity measures. This typically means that the people within the network have many disparate interests, backgrounds, and beliefs. High diversity doesn't necessarily equate to large size, but it often may. Therefore in the model, networks with high diversity are assumed to have more individuals (e.g., a social network with a diversity of five will be converted to 250 individuals; diversity of nine will be converted to 450).

When the user has access to two or more online social networks, the model assumes that the user is overall more socially connected. In this case, social network diversity scores will be multiplied by a larger factor (150 instead of 50) so that two networks with diversity of two and eight, respectively, will be converted to one score of 1500 individuals. All computations going forward use the number of individuals computed to be part of the social network(s) collectively. This becomes the "social network size" regardless of the number or diversity of social networks in the initial setup.

2.4 Calculating replies.

Each question may receive replies from real-life friends or people presumed to be in online social networks. Friends will give a reply if their knowledge of the topic is high. A knowledge score of 8 will give a correct response 70% of the time; a knowledge score of 9 will give a correct response 85% of the time. Only friends with knowledge scores above 7 provide replies, but since we assume some inaccuracies in replies, the model reduces the likelihood of a (correct) reply for scores of 8 and 9. This way every response collected can be considered correct.

Replies from social networks is based on the cumulative social network size (described above). A social network will give a different number of replies based on its size (larger networks tend to provide more responses). It's unclear if there is one straight-forward metric to predict this behavior, so I estimated number of replies by extrapolating from data in my current study (Evans et al. 2009). For example, a network of 150 people will give 1 response. Other network sizes and replies provided is shown in the table to the right.

Network size	# Responses
150	1
425	2
648	3
790	4
2300	5
4890	12
7389	19
8888	23
10000	27

The total number of replies from friends and social networks is combined and plotted in the *#Replies from friends* plot in the model [**Figure 2**].

2.5 Determining whether the question was answered.

For a question to be answered, the following conditions will need to apply. If the question is easy [0], the question will be answered if the user receives 2 total replies from any combination of friends or social network responses. Although this class of questions tends to be fact-finding queries, the model assumes the user would want two corroborating data points before deciding that the information need was satisfied.

Hard questions [1] will be answered differently depending on the amount of knowledge the user has about his posed topic. In the organizational learning literature, it has been observed that knowledge about a problem domain has an effect on how well new information is assimilated for inquiries requiring problem reformulation and brainstorming (Cross & Sproull 2004)—in our model, this applies to hard questions. Thus, when the user's knowledge is high (a score of 8 or 9), the question will be answered if 3 replies are received. If the user's knowledge is low (below 7), the question requires 5 total replies before getting answered.

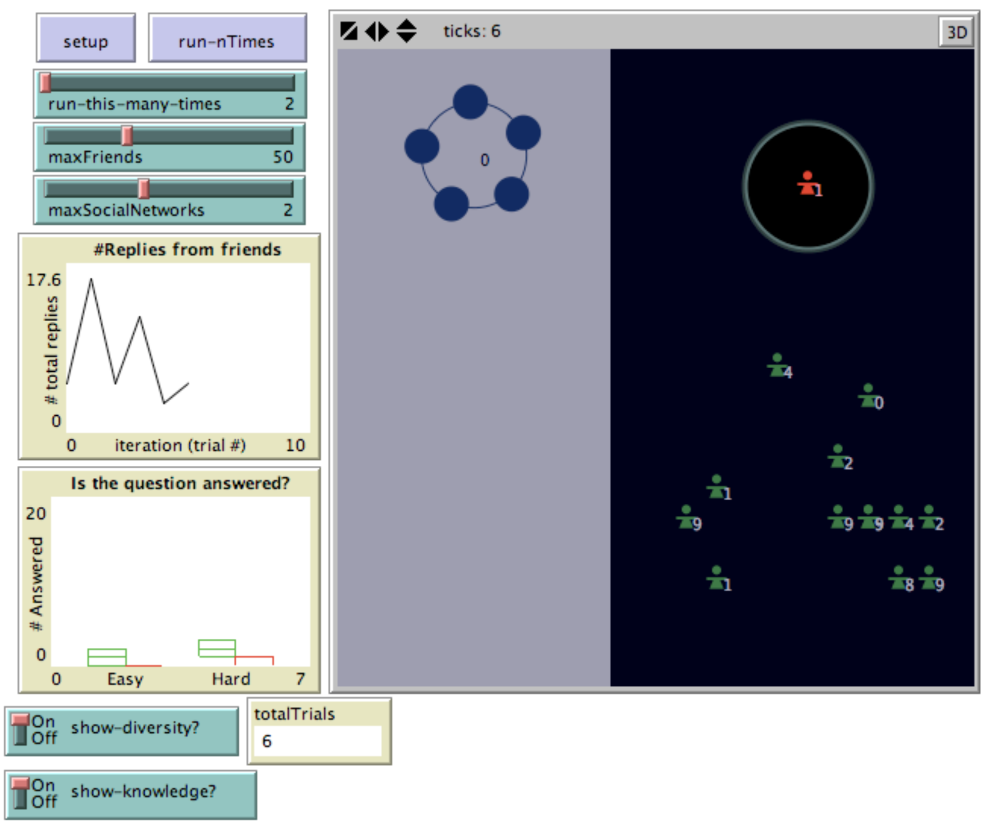

Figure 2. The model after 6 trials where 5, 16, 5, 12, 3, and 5 replies were received, respectively (*#Replies from friends* plot). Both easy questions were answered; of 4 hard questions, 3 were answered (*Is the question answered?* plot).

Answer success is shown in the *Is the question answered?* plot [**Figure 2**]. Each question type has two outcomes: answered [in green] or not answered [in red].

3. The analysis

3.1 The canonical question-answering model

The first goal of the project was to depict a canonical question-answering system. Thus, the model was designed to be run for multiple random iterations [see the *run-this-many-times* slider] so that an analysis of the data across many trials would inform our understanding of social information seeking. The discussion below reports the results of running the model 1000 times.

3.1.1 Questions are easy to answer

I was surprised to learn that easy questions received answers 96% of the time. I was even more surprised that hard questions were answered 86% of the time. This is easily visualized in the highlighted plot in **Figure 3**.

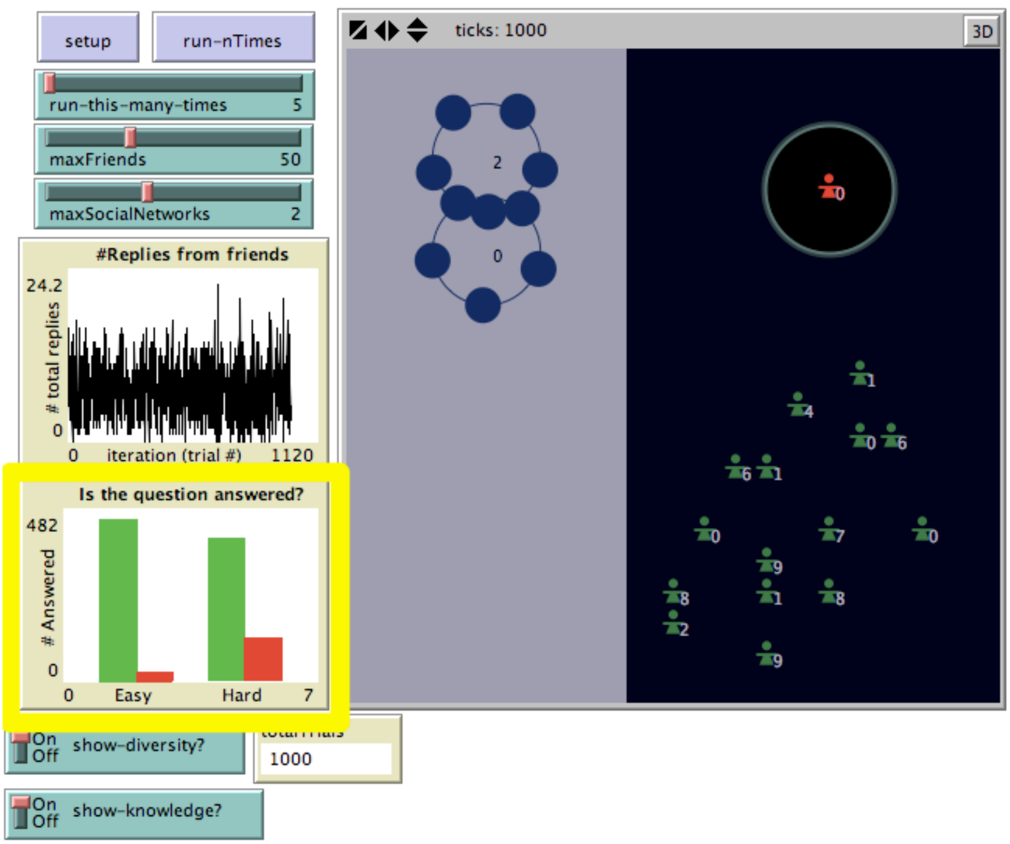

Figure 3. After 1000 trials of randomized network configurations, the highlighted plot shows the percentage of questions that were successfully answered.

While it surprised me that so many questions received correct answers, this actually reflects the notion that several other researchers have echoed recently: that information seeking can be social (EVANS & CHI 2008; MORRIS 2008; CHI 2009). On the other hand, the model's specific parameters determined the outcome of the 1000-trial run. Since the average number of replies was 7.3, even hard questions with low user knowledge (the most stringent use case) would have received answers. We would need to (re)consider or update the parameter settings after being informed by new data or new literature to be fully confident in the conclusion of this simulation.

3.1.2 Friends provided most replies

Another interesting observation was that most answers came from friends. The average number of friend replies was 5.2, whereas the average number of social network replies was 2.2. Again, this is likely an artifact of the model since social network size had to grow fairly large before it provided many responses. These 2.2 replies would come from a network size of just over 425. In practice, online social networks are rarely much larger than 425. A recent report in the Economist stated that the average Facebook network size is 120, but that some users had networks of up to 500 friends (ECONOMIST 2009). This would suggest that the model's social network sizes may have been estimated with accuracy; but since response rates from networks have not been tested empirically in any study, we should be cautious about over-interpreting the present results.

3.2 The effect of network configuration on results

A second goal of the project was to understand how various network configurations provided different kinds of support in question-answering. Therefore, I looked at the average social network size and the average number of friends for questions that were both answered and unanswered [Figure 4].

I was surprised to see very few differences between conditions. The average social network size for hard questions that got answered was just over 1000; for easy questions that got answered it was about 750. The relative differences between network sizes for answered and unanswered questions did not strike me as significant. However, since networks' reply rates produce different results at 700 versus 1000, it is possible that the absolute differences between these conditions is significant.

Average number of friends reflects a similar finding. Both hard and easy questions had an average of ~27 friends in trials where the question was answered. When the question was not answered, trials with hard questions had about twice as many friends as trials with easy questions. Even though friends were most responsible for providing answers in the model, it seems that more friends are needed to answer hard questions. This makes sense given the model's parameters, actually, since most hard

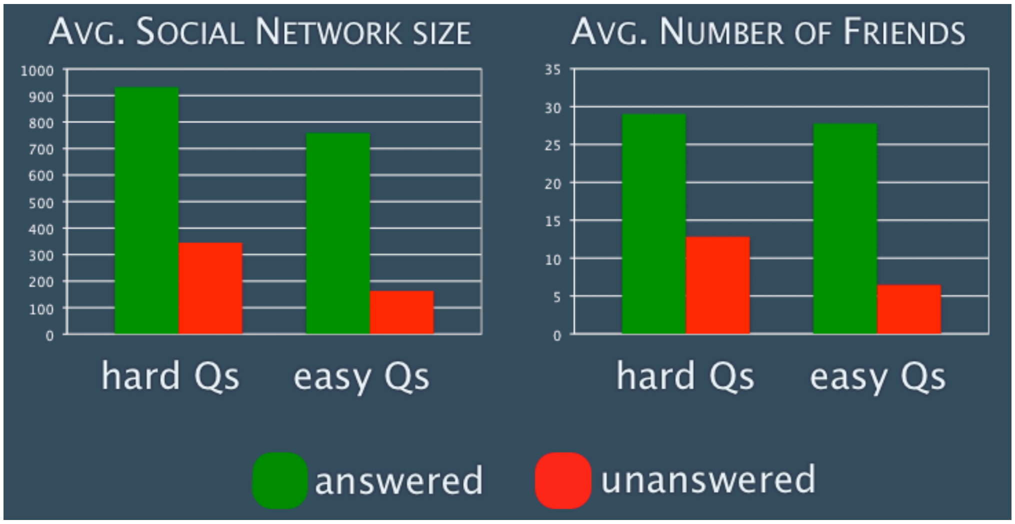

Figure 4. Average social network and friend group sizes for hard and easy questions that got answered [in green] or remained unanswered [in red].

questions are asked by users with low knowledge on the topic. In these cases, a larger friend network is necessary to successfully answer questions. In practice, this also makes sense. Exploratory questions that require brainstorming or reformulation would be harder to answer with the help of only one or two other individuals, unless each person was willing to provide substantial support to the searcher. Thus, we would expect a higher threshold of help for question-answering. This is reflected in the plot in Figure 4: even with 12 friends, hard questions fail to get addressed.

4. Conclusions

4.1 Extensions to the model

There are many social and cognitive factors that influence whether a searcher turns to a colleague for help or combines personal and impersonal sources (e.g., friends + Google). Most of these are not modeled here and would make natural extensions to the model. A few examples include:

1. *Relationship with friends*: The interpersonal relationship between people surely affects information seeking behaviors. Users are more likely to seek information from close friends and colleagues than from distant ties (Cross et al. 2001).

2. *Time since last interaction*: Searchers are also more likely to turn to friends with whom they've recently interacted (Evans et al. 2009). In addition to the relational qualities [above] which we could model, the history of interactions and time since last interaction may be important factors, too.

3. *Availability*: Interestingly, availability and willingness to engage may be more important for information seeking than expertise (CROSS & SPROULL 2004). These variables will be affected by relational and structural qualities of a network: close friends may be more willing to engage or have a history of similar interactions; and colleagues who are physically close are often more accessible (and hence, available).

4.2 Summary and conclusion

There may be many additional metrics that would strengthen the fidelity of a question-answering model. The difficulty with modeling is that factors may interact with each other in unknown ways, making multiple variables difficult to model accurately. Human relations and social network structures (especially across the physical-digital divide) are particularly complicated. Although the present model is a simplification of question-answering behavior, it did reveal two issues that are relevant to researchers and designers of social search and question-answering systems.

Real-life friends should not be taken for granted. They are known to provide emotional support to individuals; they also provide critical informational support. The challenge on the web is to replicate the positive benefits of real-life or face-to-face friends such that users who are digitally (but not socially) isolated may continue to make use of this important resource.

Online social networks may not provide great question-answering support [yet]. My current research echos this sentiment (EVANS ET AL. 2009)—social networks above 1000 individuals are well-suited to question-answering because of a presumed underlying diversity among a few hundred people. Smaller networks provide few, if any, responses. This does not mean that online social communities will never support information seeking habits. People may be hesitant to reply to anonymous requests for information (personalizing questions may be more effective). Social networking users may simply be unaccustomed to high levels of online participation. Services like Facebook and Twitter make participation with friends fun and enjoyable; in time, these services may better support question-answering behaviors. Alternatively, new services could algorithmically exploit social metadata from online networks as support for information seekers. This may avoid the need for direct user contributions, while still providing high-level feedback on concepts and related concepts in a given subject area.

5. References

Ackerman, M.S. (1998). Augmenting organizational memory: A field study of Answer Garden. *ACM Transactions on Information Systems 16*(3), pp. 203-224.

Allen, T.J. (1977). *Managing the Flow of Technology*. Cambridge, MA: MIT Press.

Bandura, A. (1989). Social Cognitive Theory. In R. Vasta (Ed.), *Annals of child development, 6*. Six theories of child development (pp. 1-60). Greenwich, CT: JAI Press.

Brown, J.S. and Duguid, P. (2000). *The Social Life of Information*. Harvard Business School Press.

Burt, R. (1992). *Structural Holes*. Cambridge, MA: Harvard University Press.

Chi, E.H. (2009). Information Seeking Can Be Social. *Computer 42* (3), pp. 42-46.

Cross, R., Rice, R.E., and Parker, A. (2001). Information seeking in social context: Structural influences and receipt of information benefits. *IEEE Transactions on Systems, Man and Cybernetics—Part C 31*(4), pp. 438-448.

Cross, R. and Sproull, L. (2004). More than an answer: Information relationships for actionable knowledge. *Organization Science 15*(4), pp. 446-462.

The Economist. (Feb 26, 2009). Primates on Facebook: Even online, the neocortex is the limit. Retrieved from: http://www.economist.com/science/displaystory.cfm?story_id=13176775.

Evans, B.M. and Chi, E.H. (2008). Towards a model of understanding social search. In *Proc. CSCW'08, ACM Press,* pp. 485-494.

Evans, B.M., Kairam, S., and Pirolli, P. (2009). Exploring the cognitive consequences of social search. *Work in Progress, CHI'09, ACM Press*.

Fox, E.A., Hix, D., Nowell, L.T., Brueni, D.J., Wake, W.C., Heath, L.S., and Rao, D. (1993). Users, user interfaces, and objects—envision, a digital library. *JASIS 44*(8), pp. 480–491.

Golovchinsky, G., Pickens, J., and Back, M. (2008). A Taxonomy of Collaboration in Online Information Seeking. In *Proc of the Workshop on Collaborative Information Retrieval*, June 20th, Pittsburgh, PA.

Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology 78*, pp. 1360-1380.

Hatch, T. and Gardner, H. (1993). Finding cognition in the classroom: An expanded view of human intelligence. In G. Salomon (Ed.), *Distributed cognitions: Psychological and educational considerations* (pp. 164-187). New York: Cambridge University Press.

Karasavvidis, I. (2002). Distributed Cognition and Educational Practice. *Journal of Iterative Learning Research 13*(1/2), pp. 11-29.

Marchionini, G. (2006). Exploratory search: From finding to understanding. *Communications of the ACM 49*(4), pp. 41-46.

Morris, M.R. (2008). A survey of collaborative web search practices. In *Proc. CHI'08, ACM Press*, pp. 1657–1660.

White, R.W., Drucker, S.M., Marchionini, M., Hearst, M., schraefel, m.c. (2007). Exploratory search and HCI: Designing and evaluating interfaces to support exploratory search interaction. *Extended Abstracts CHI'07*, *ACM Press*, pp. 2877-2880.