

Around the Water Cooler: Shared Discussion Topics and Contact Closeness in Social Search

Saranga Komanduri

Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA
sarangak@cs.cmu.edu

David Huffaker

Google, Inc.
1600 Amphitheatre Parkway, Mountain View, CA
huffaker@google.com

Lujun Fang

University of Michigan, Ann Arbor
2260 Hayward Avenue, Ann Arbor, MI
ljfang@umich.edu

Jessica Staddon

Google, Inc.
1600 Amphitheatre Parkway, Mountain View, CA
staddon@google.com

Abstract

Search engines are now augmenting search results with *social annotations*, i.e., endorsements from users' social network contacts. However, there is currently a dearth of published research on the effects of these annotations on user choice. This work investigates two research questions associated with annotations: 1) do some contacts affect user choice more than others, and 2) are annotations relevant across various information needs. We conduct a controlled experiment with 355 participants, using hypothetical searches and annotations, and elicit users' choices. We find that domain contacts are preferred to close contacts, and this preference persists across a variety of information needs. Further, these contacts need not be experts and might be identified easily from conversation data.

1 Introduction

In the past few years, search engines have introduced new products that utilize social network data. In October 2009, Google introduced Google Social Search¹ (Heymans and Viswanathan 2009) and in May 2011, Bing introduced Facebook integration (Mehdi 2011). Both of these products add *social annotations* to search results, as shown in Figure 1. Social annotations are visual indicators of endorsement from a user's social network contacts, which are shown along with search results. This work takes some initial steps toward understanding how social annotations can affect user choice.

However, a scientific examination of this issue is very tricky. Annotated search results are novel, and it is difficult to disentangle the effect of novelty from utility. If a user selects an annotated result over an unannotated result, is the annotated result more useful? Or does the user have an initial curiosity about annotated results that will fade over time? To control for the confound of novelty, we conducted a controlled experiment with hypothetical search queries and only make comparisons between annotated results from different contacts. By manipulating the search query and the

contacts used in annotations, we find that *domain* contacts, individuals with whom the user discusses the domain of the search query, have a significantly stronger effect on user choice than those individuals the user reports as *close* contacts. To address concerns with external validity that naturally arise when using hypothetical search results, we asked our participants about the credibility of the annotations used in our study, and accounted for this in our analysis. Randomization was also used in our study design to control for ordering effects with contacts and search results.

In the future, as social search products become more prevalent and the number of social annotations increases, users will be faced with social annotations more often. This could create new problems with using social annotations effectively. For example, search engines might have multiple annotations available for a single result and be forced to select or emphasize a subset. There is also much work in HCI on issues like change blindness (Simons and Rensink 2005) and inattentional blindness (Simons and Chabris 1999) that could cause users to ignore social annotations if they become habituated to them. Human faces can be distracting and increase the user's cognitive load, decreasing their performance on cognitive tasks (de Fockert et al. 2001). We believe that all of these problems can be addressed by showing only relevant annotations to users. This work suggests that domain contacts are more relevant than traditional close contacts and also that annotations are relevant to a larger class of search queries than might be expected.

This work is structured as follows. Section 2 discusses related work and our contributions to the field. Section 3 presents our research questions, study design, and analysis methods. Section 4 presents our main results and the results of subset analyses, while Section 5 discusses limitations. Finally, Section 6 concludes with high-level findings and some implications for how they might be applied.

2 Background and Related Work

Social network data

At the time of this writing, we know of only one other work studying social annotations in web search. Muralidharan et

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Google recently expanded this product (Singhal 2012).

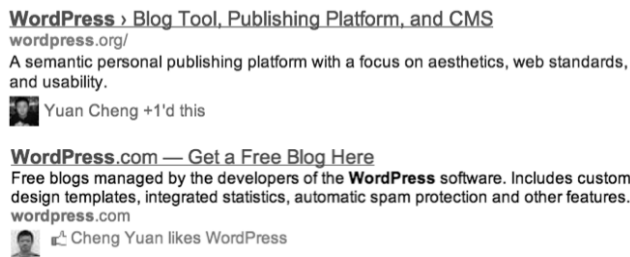


Figure 1: Screenshots of social annotations from Google Social Search (top) and Bing (bottom).

al. used an eye tracker to examine user perception in a small-sample ($N = 23$), lab-study environment (Muralidharan, Gyongyi, and Chi 2012). The authors found that, in the particular case of the Google Social Search product, social annotations are unseen by users because their “visual parsing behavior” ignores the areas of the page in which annotations are placed. Only by altering the page structure, or increasing the annotation size, could they get users to pay attention to them. Our study worked with a much larger sample of participants and focused on the search results they chose. We were also able to select the contacts for our annotations based on conversation domains.

A natural way for search engines to include social network data is as a factor in determining search result ranking, thus partially addressing the “potential for personalization” (Teevan, Dumais, and Horvitz 2010). Indeed, Bing will “surface results, which may typically have been on page three or four, higher in its results based on stuff your friends have liked” (Mehdi 2011). However, such personalization could be performed invisibly, without displaying annotations to the user. Ranking search results based on social network data will obviously affect user choice and has been studied previously (Bao et al. 2007; Yanbe et al. 2007; Heymann, Koutrika, and Garcia-Molina 2008; Zanardi and Capra 2008; Carmel et al. 2009), but the visual display of annotations will have its own effect that we seek to isolate in this work.

Other previous work in social annotations focused on social bookmarking systems similar to del.icio.us (Hong et al. 2008; Nelson et al. 2009), While these works highlight the utility of social signals in information seeking, they do not address the scenario we examine here: social annotations inline with web search results.

Taxonomy of web search

Broder suggested a taxonomy of web search that was later discussed by Evans and Chi in the context of social search (Broder 2002; Evans and Chi 2008). Broder classified search queries into three categories based on *information need*, or “the need behind the query.” These categories are: navigational, informational, and transactional. With a navigational query, the user’s need is to reach a particular site, e.g. searching with the term “netflix” to reach the site netflix.com. In a transactional query, the user wishes to conduct a transaction with some website, e.g. searching with

the term “movie tickets” to find a site for buying tickets. Finally, and not surprisingly, the purpose of an informational query is to find information.

Evans and Chi analyzed the information-seeking habits of web searchers across various information needs. They found that when performing navigational or transactional queries, searchers did not interact with others. Therefore, they hypothesized that “it is unlikely that socially-augmented search would improve or facilitate transactional or navigational information retrieval.” One of the goals of this work is to evaluate this hypothesis.

Social q&a and recommender systems

There have been many studies of social question and answer (q&a) systems (Mamykina et al. 2011; Ackerman and McDonald 1996; Horowitz and Kamvar 2010; Golbeck and Fleischmann 2010). These studies consistently find that the active participation of experts is necessary for a successful q&a system, where experts are those who can provide correct answers to many questions within a domain. Offline studies of information seeking in organizations also corroborates this reliance on experts (Borgatti and Cross 2003), as do studies of recommender systems (Golbeck 2006). However, the experts in these systems are often strangers to those whose questions they answer. Our work focuses on the user’s close social network – people that the user has discussions with. These contacts need not be experts. Cornwell and Cornwell found that when a person is inexpert, as often occurs in disadvantaged populations, they are less likely to have close ties to an expert in their social network (Cornwell and Cornwell 2008).

3 Methodology

Research Questions

This work is focused on two main research questions:

1. Are domain contacts more relevant than close contacts in annotations?
2. If so, does this effect hold across various information needs and domains?

These questions were heavily inspired by the work of Evans and Chi in their hypothesis that socially-augmented search would only be beneficial to some web queries (Evans and Chi 2008). Due to time constraints, we were only able to investigate six different web queries. The queries along with their information needs and domains are given in Table 1.

Information needs were taken mostly from Broder’s taxonomy (Broder 2002) with one addition. After considering how the taxonomy would apply to social search, we expanded the informational category into two categories: *informational-objective* and *informational-subjective*. Informational queries can be placed on a continuum from questions with specific answers, e.g. “how tall is the Empire State Building,” to questions that seek an opinion, e.g. “what is the most fun city in the world.” We consider a search for facts to be pursuing an informational-objective need, and a search for opinions an informational-subjective

Query	Information Need	Domain
“current news”	Informational-subjective	Current events
“harry potter movies in order”	Informational-objective	Entertainment
“how to use twitter”	Informational-subjective	Technology
“movie times”	Transactional	Entertainment
“movie trailers”	Informational-subjective	Entertainment
“netflix”	Navigational	Entertainment

Table 1: Queries used in our study with their information need and domain.

need. We hypothesized that, like navigational and transactional queries, informational-objective queries would not benefit from social annotation.

We considered a large set of possible queries before settling on the six that are examined in this study. We selected the entertainment domain for investigating the information needs taxonomy because it seemed easiest to construct realistic queries for. We selected the informational-subjective need to investigate across domains because we expected it to receive the strongest effect from annotations.

Recruitment

We recruited participants from a broad pool of testers, the majority of whom have college degrees and are within 24-45 years of age. Slightly more than half of the pool is male. All studies were completed online with no direct interaction between the users and the authors of this paper. Like Amazon’s Mechanical Turk,² participants are compensated for their work. However, unlike Mechanical Turk, the participant pool is reviewed periodically to ensure that any respondents who are not giving study tasks the necessary attention are removed from the pool.

On average, participants were heavy social network users. 73% of participants reported visiting a social network at least once per day. Participants in this pool are also frequently asked to evaluate online content.

Study Design

The study consisted of three parts: two surveys and the annotation study.

Surveys First, we sent participants a survey in which we asked them questions about their social networks and usage of social networking services. After answering the question “How many close contacts do you have (i.e. people that you discuss important matters with),” participants were asked to list the first and last names of at most 5 of their closest contacts. The “discuss important matters” wording was taken from the General Social Survey, an accepted instrument for gathering information about personal networks (Bailey and Marsden 1999; Straits 2000). They were also asked to provide additional domain-specific lists for those they discuss entertainment with and those they discuss current events with. The contact lists were not required and participants could fill out zero to five names for any of the lists. We also

cleaned the data to remove invalid responses, such as “no name.” 286 participants provided at least one response in this phase of the study.

Six weeks later, we recruited a larger segment of the pool to take a second survey. This survey was similar to the first survey but asked participants to provide another list of contacts that they get computer advice from. 280 participants took the survey a second time, along with 69 additional participants.

After an additional two weeks, we administered the annotation study.

Annotation study The annotation study investigated the six search queries shown in Table 1. For each query, participants were shown mock search queries from Google with social annotations on the top two results. Search results were initially captured from a standard Google search made using the Chrome browser in “incognito” mode. For a given query, all participants who received that query saw the same search results – except for the names used in the social annotations and the ordering of the top two results, as explained below.

The mock results were personalized for each participant using the contact lists collected in the two surveys. The top two search results had social annotations, the rest did not. Of these two annotations, one was a domain contact and the other was a close contact. The domain contact was the first name on the contact list provided by the participant for the domain of the search query. For example, given a participant and the search query “current news,” the domain contact was the first name on the list of people the participant reported they discussed current events with (see Table 1 for the mapping of queries to domains.) The close contact was then selected by choosing the first name from the list of a participant’s close contacts which did not appear on the domain contact list. Thus we compare domain contacts to close contacts. If a user did not have two such contacts for a query, they were not given that query. For the 280 participants who took the survey twice, contacts were selected from the first survey, which was administered two months prior. The number of responses collected for each query varied from 269 to 288.

To control for ordering effects, the study design made heavy use of randomization. The ordering of the contacts was balanced so that the domain contact appeared first in half of the search queries, and the close contact appeared first in the other half. Independently, the order of the top two search results was randomly switched for half of the queries,

²<https://www.mturk.com>

to control for potential interactions between annotations and search results at particular ranks. We also randomized the order in which queries were presented to participants.

With each search query, we asked participants “Of these search results, which would you choose for the query [query text]?” The participant was not restricted to the top two results. They were then asked several questions about their closeness to the contacts, the believability of the annotations, and the relevance of the annotations.

Analysis

Our analysis focuses on the search result selected by the participant. For each query type, we compare the rates at which participants selected domain-contact annotated results to their selections of close-contact annotated results. We also construct 95% confidence intervals for these rates to determine statistical significance. This analysis is repeated for certain subsets of our data selected a priori, e.g., all responses where annotations are believable, to ensure that our findings are robust. Finally, to help understand why domain-contact or close-contact annotations were selected, we performed logistic regressions using the other questions from our survey.

4 Results

Overall, we found a strong preference for domain-contact annotations on most search queries. This result was not statistically significant when looking at individual queries, but is significant when aggregating across domains and information needs. We also found a stronger effect when only looking at those queries that participants considered credible.

Preference for domain contacts

For five of the six queries, domain-contact annotations are preferred to close-contact annotations, on average. These results are shown in Figure 2. Many of these results are not statistically significant at the $\alpha = .05$ significance level, with only the “HPMovies” and “Twitter” queries showing statistical significance. Nevertheless, the trend over the six queries suggests a strong preference for domain-contact annotations.

A more fine-grained view of the responses makes the pattern more clear. Figure 4 illustrates the bias created by the top-ranked search result and the ability of social annotations to increase or decrease the rate at which search results are picked. For the “HPMovies” query, a domain-contact annotation is enough to override the top-ranked search result bias. We find this especially surprising since it is an informational-objective query. Similarly, the “netflix” and “movie times” queries, though not showing a significant effect, indicate that a preference for domain-contact annotations might exist for at least some navigational and transactional queries.

Credibility

We use the term *credible* to refer to how believable an annotation is to a participant. For each annotation, we asked our




Figure 2: The fraction of responses that select results annotated by domain contacts (DC-solid) and close contacts (CC-dashed) for each query. The figure includes the mean across all queries (horizontal lines) and 95% confidence intervals (vertical error bars). There was an average of 280 responses for each query.




Figure 3: A reproduction of Figure 2 for only credible responses (75% of collected responses).

participants to agree or disagree on a 5-point Likert scale with the statement “[contact name] would +1 a site like [website of search result].” We use responses to this question in our analysis to account for those responses where our hypothetical search results might have been too surprising. For 57% of responses, participants agreed that both annotations were credible. For 75% of responses, participants agreed or were neutral about both annotations. We call a response credible if it falls in this 75% subset.

To verify our results, we examined only those credible responses. The results are shown in Figure 3. Note that the




Figure 4: Another view of the data shown in Figure 2, broken down by which contact was first. Circles indicate when the domain contact annotated the top result, and triangles indicate when the close contact did so.

results show a much stronger trend toward domain-contact annotations. Specifically, the “current news” and “netflix” queries show a statistically significant preference for domain contacts, in addition to the “HPMovies” and “Twitter” queries. This suggests that, when participants are presented with a scenario that they find unrealistic, results can be affected in unexpected ways. This highlights the need for data quality checks in experimental design, especially when asking participants about hypothetical scenarios.

Preference for domain contacts holds across information needs and domains

We can also aggregate responses over the information needs and domains identified in Table 1. The results of this, including only credible responses, are shown in Figures 5 and 6 for information needs and domains respectively. As expected from the trends in the per-query results, participants significantly prefer search results that are annotated by domain contacts over those annotated by close contacts. This is especially surprising for the informational-objective and navigational queries, where no effect was expected but a statistically significant result was found.

Explaining preferences

To try to understand the drivers of annotation preferences, we performed two logistic regressions on the study data, one each for domain and close-contact annotations. We take the selection of a domain or close-contact annotation as our outcome variable, and the other data we collected in our study as predictors. The results are given in Tables 2 and 3.

For selection of a domain-contact annotation, the significant predictors are relevance (positive) and user interest (negative). For close-contact annotations, the significant




Figure 5: Aggregate preferences over information needs for only credible responses.




Figure 6: Aggregate preferences over domain for only credible responses.

predictors are credibility (positive) and self-expertise (negative).

Expertise Note that expertise of contacts, or even expertise relative to the participant, is not a significant positive predictor of user preference. This suggests that users do not pay much attention to the expertise of their friends when evaluating social annotations from them. Further, in the case of close contacts, self-expertise is a significant negative predictor of selection. In other words, if the user has expertise, they are less interested in a close-contact annotation.

Nondomain-contact annotations

There is an alternate interpretation of our results. Remember that close contacts were chosen such that they were not

Variable	Log likelihood increase per unit	$Pr(> z)$
DC relevance	0.14	<0.01
DC credibility	0.32	0.09
User interest	-0.10	<0.01
Self expertise	-0.01	0.07
Connection type = co-worker	-12.77	0.82
Connection type = family	0.56	0.96
Connection type = friend	0.49	0.12
DC Closeness	-0.03	0.72
DC Expertise	-0.13	0.03
Relative DC Expertise to self	0.09	0.13

Table 2: Regression coefficients for whether the DC annotated results is picked. Significant coefficients are in boldface.

Variable	Log likelihood increase per unit	$Pr(> z)$
CC relevance	0.06	0.54
CC credibility	0.23	<0.01
User interest	-0.09	0.09
Self expertise	-0.17	<0.01
Connection type = co-worker	-0.76	0.53
Connection type = family	0.26	0.59
Connection type = friend	0.10	0.83
CC Closeness	0.05	0.43
CC Expertise	-0.01	0.90
Relative CC Expertise to self	0.08	0.27

Table 3: Regression coefficients for whether the CC annotated results is picked. Significant coefficients are in boldface.

on the user’s top-5 list of domain contacts. Rather than displaying a preference for domain-contact annotations, it is possible that participants are exhibiting a lack of preference, or even a dislike, of nondomain-contact annotations. This does not contradict our result. The general finding – that domain contacts are preferred for annotations – is still the same. However, this points toward the need for a better understanding of the influence of annotations in future research.

5 Limitations and Future Work

We realize that this study has several methodological limitations. We list them here, along with some indications for future work.

The examination of only six search queries limits the generalizability of the results. A larger sample of queries should be examined. We found it surprisingly difficult to select queries for use in this study, as most common search queries can apply to multiple information needs. For example, the top search results for “netflix” indicate that this is often a navigational query (the top two results are the Netflix home page and login page), but it could also be used to search for information about the company. Refining the query to “netflix.com” does not solve the problem, and further refinements lead to unusual queries that can seem unrealistic.

The hypothetical nature of the study allowed us to isolate the effect of annotations. However, it should be confirmed

by research performed *in situ* with real social annotations. Muralidharan et al. attempted this and found that the occurrence of annotations for their participants was very sparse and often from unknown contacts (Muralidharan, Gyongyi, and Chi 2012). Once social search products become more familiar to users at large, experiments with social annotations should become more straightforward. Future research could also attempt to study only those users who already see a large number of annotations, for whom the effect of novelty would be negligible.

Finally, this work only examines preferences between two close contacts from each user. Annotations could come from sources that are not close to the user but still in their online social network, such as well-known experts that the user is following. Based on our results, such annotations might be highly valued by users.

6 Conclusion

In this paper, we present two novel findings: 1) that domain contacts are valuable for annotations, and 2) that “closeness” is not necessarily a valuable property of annotators, though it might have merit in cases where the user is not an expert. Domain contacts are those people with whom the user shares discussions related to the search query, and need not be experts.

We found a significant preference for domain contacts across three domains and all but the transactional informa-

tion need. That annotations have an effect on user choice for navigational and informational-objective queries is surprising. The effect of annotations can also be quite strong. In many cases, we observed that a domain-contact annotation on the second search result was selected at a higher rate than a close-contact annotation on the top result.

This work represents a first step forward in social annotations research, though there are still many questions that remain unanswered. As social search becomes more common, search engines will have more annotations to serve, and these questions will only increase in importance.

7 Acknowledgements

The authors thank Ed H. Chi for his extensive feedback during the planning stages of the project, Zoltán Gyöngyi for feedback during the initial stages, and Dan Russell for his feedback on the study design. We also thank Aditi Muralidharan for providing us with a copy of her paper while it was in submission.

References

- Ackerman, M. S., and McDonald, D. W. 1996. Answer garden 2: merging organizational memory with collaborative help. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work*, CSCW '96, 971105. New York, NY, USA: ACM.
- Bailey, S., and Marsden, P. 1999. Interpretation and interview context: examining the general social survey name generator using cognitive methods. *Social Networks* 21(3):287309.
- Bao, S.; Xue, G.; Wu, X.; Yu, Y.; Fei, B.; and Su, Z. 2007. Optimizing web search using social annotations. In *Proceedings of the 16th international conference on World Wide Web*, WWW '07, 501510. New York, NY, USA: ACM.
- Borgatti, S. P., and Cross, R. 2003. A relational view of information seeking and learning in social networks. *Management Science* 49(4):432–445. ArticleType: research-article / Issue Title: Special Issue on Managing Knowledge in Organizations: Creating, Retaining, and Transferring Knowledge / Full publication date: Apr., 2003 / Copyright 2003 INFORMS.
- Broder, A. 2002. A taxonomy of web search. *SIGIR Forum* 36(2):310.
- Carmel, D.; Zwerdling, N.; Guy, I.; Ofek-Koifman, S.; Har'el, N.; Ronen, I.; Uziel, E.; Yogev, S.; and Chernov, S. 2009. Personalized social search based on the user's social network. In *Proceedings of the 18th ACM conference on Information and knowledge management*, CIKM '09, 12271236. New York, NY, USA: ACM.
- Cornwell, E. Y., and Cornwell, B. 2008. Access to expertise as a form of social capital: An examination of race- and Class-Based disparities in network ties to experts. *Sociological Perspectives* 51(4):853–876. ArticleType: research-article / Full publication date: Winter 2008 / Copyright 2008 Pacific Sociological Association.
- de Fockert, J. W.; Rees, G.; Frith, C. D.; and Lavie, N. 2001. The role of working memory in visual selective attention. *Science* 291(5509):1803–1806.
- Evans, B. M., and Chi, E. H. 2008. Towards a model of understanding social search. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, CSCW '08, 485494. New York, NY, USA: ACM.
- Golbeck, J., and Fleischmann, K. R. 2010. Trust in social Q&A: the impact of text and photo cues of expertise. In *Proceedings of the 73rd ASIS&T Annual Meeting on Navigating Streams in an Information Ecosystem - Volume 47*, ASIS&T '10, 77:177:10. Silver Springs, MD, USA: American Society for Information Science.
- Golbeck, J. 2006. Filmtrust: movie recommendations from semantic web-based social networks. In *3rd IEEE Consumer Communications and Networking Conference, 2006. CCNC 2006*, volume 2, 1314–1315. IEEE.
- Heymann, P.; Koutrika, G.; and Garcia-Molina, H. 2008. Can social bookmarking improve web search? In *Proceedings of the international conference on Web search and web data mining*, WSDM '08, 195206. New York, NY, USA: ACM.
- Heymans, M., and Viswanathan, M. 2009. Official google blog: Introducing google social search: I finally found my friend's new york blog! <http://googleblog.blogspot.com/2009/10/introducing-google-social-search-i.html>.
- Hong, L.; Chi, E. H.; Budiu, R.; Pirolli, P.; and Nelson, L. 2008. SparTag.us: a low cost tagging system for foraging of web content. In *Proceedings of the working conference on Advanced visual interfaces*, AVI '08, 6572. New York, NY, USA: ACM.
- Horowitz, D., and Kamvar, S. D. 2010. The anatomy of a large-scale social search engine. In *Proceedings of the 19th international conference on World wide web*, WWW '10, 431440. New York, NY, USA: ACM.
- Mamykina, L.; Manoim, B.; Mittal, M.; Hripcsak, G.; and Hartmann, B. 2011. Design lessons from the fastest q&a site in the west. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, 28572866. New York, NY, USA: ACM.
- Mehdi, Y. 2011. Bing | facebook friends now fueling faster decisions on bing. *Bing Community*.
- Muralidharan, A.; Gyongyi, Z.; and Chi, E. H. 2012. Social annotations in web search. In *Proceedings of the 2012 annual conference on Human factors in computing systems*, CHI '12.
- Nelson, L.; Held, C.; Pirolli, P.; Hong, L.; Schiano, D.; and Chi, E. H. 2009. With a little help from my friends: examining the impact of social annotations in sensemaking tasks. In *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, 17951798. New York, NY, USA: ACM.
- Simons, D. J., and Chabris, C. F. 1999. Gorillas in our midst: sustained inattention blindness for dynamic events. *Perception* 28(9):1059–1074. PMID: 10694957.

- Simons, D., and Rensink, R. 2005. Change blindness: Past, present, and future. *Trends in Cognitive Sciences* 9(1):1620.
- Singhal, A. 2012. Official google blog: Search, plus your world. <http://googleblog.blogspot.com/2012/01/search-plus-your-world.html>.
- Straits, B. 2000. Ego's important discussants or significant people: An experiment in varying the wording of personal network name generators. *Social Networks* 22(2):123140.
- Teevan, J.; Dumais, S. T.; and Horvitz, E. 2010. Potential for personalization. *ACM Trans. Comput.-Hum. Interact.* 17(1):4:14:31. ACM ID: 1721835.
- Yanbe, Y.; Jatowt, A.; Nakamura, S.; and Tanaka, K. 2007. Can social bookmarking enhance search in the web? In *Proceedings of the 7th ACM/IEEE-CS joint conference on Digital libraries, JCDL '07*, 107116. New York, NY, USA: ACM.
- Zanardi, V., and Capra, L. 2008. Social ranking: uncovering relevant content using tag-based recommender systems. In *Proceedings of the 2008 ACM conference on Recommender systems, RecSys '08*, 5158. New York, NY, USA: ACM.