

SearchBuddies: Bringing Search Engines into the Conversation

Brent Hecht¹, Jaime Teevan², Meredith Ringel Morris², and Dan Liebling²

¹Northwestern University, Evanston, IL, USA, ²Microsoft Research, Redmond, WA, USA

Abstract

Although people receive trusted, personalized recommendations and auxiliary social benefits when they ask questions of their friends, using a search engine is often a more effective way to find an answer. Attempts to integrate social and algorithmic search have thus far focused on bringing social content into algorithmic search results. However, more of the benefits of social search can be preserved by reversing this approach and bringing algorithmic content into natural question-based conversations. To do this successfully, it is necessary to adapt search engine interaction to a social context. In this paper, we present SearchBuddies, a system that responds to Facebook status message questions with algorithmic search results. Via a three-month deployment of the system to 122 social network users, we explore how people responded to search content in a highly social environment. Our experience deploying SearchBuddies shows that a socially embedded search engine can successfully provide users with unique and highly relevant information in a social context and can be integrated into conversations around an information need. The deployment also illuminates specific challenges of embedding a search engine in a social environment and provides guidance toward solutions.

Introduction

Recent research reveals that people are beginning to turn to online social networks with their information needs instead of using web search engines (Efron and Winget 2010; Morris et al. 2010a&b; Paul et al. 2011). Asking questions in the status message fields of social networks like Facebook and Twitter is a popular form of this information seeking behavior (Efron and Winget 2010; Morris et al. 2010b). For example, consider Figure 1, in which Elise has posted the question, “On Verizon, will I have to pay roaming charges if I use my cell phone in Hawaii?” to her Facebook friends.

The advantages of directing questions to online social networks include receiving trusted, personalized responses, reinforcing social ties, and avoiding the need to formulate a search query or triage a large set of search results (Efron and Winget 2010; Thom-Santelli et al. 2011; Morris et al. 2010a&b). However, this approach to information seeking also has important drawbacks: status message-based question asking is a slower means of information seeking than web search and is less likely to identify an answer (Morris et al. 2010a&b).



Elise Clayton

On Verizon, will I have to pay roaming charges if I use my cell phone in Hawaii?

8 minutes ago · Like · Comment



SearchBuddies Investigatore This page about “do cell phones work in Hawaii” may have relevant information: <http://bit.ly/h50Oul>

7 minutes ago · Like · 1 person



SearchBuddies Soshul Butterflie These friends have lived in or near Hawaii: Bo Kent (<http://on.fb.me/hXA>) and Otis Overcash (<http://on.fb.me/y77q2>).

7 minutes ago · Like

Figure 1. SearchBuddies responding to a user’s Facebook status message question. The names and pictures in all figures have been changed to preserve privacy and conversations have been truncated.

Major web search engines, recognizing the complimentary benefits of social and algorithmic search and the value of merging the two, have begun to incorporate limited information from social networks into their search results (e.g., Bing’s Facebook integration¹ and Google’s “Search Plus Your World” Google+ integration²). However, we believe that doing the inverse – bringing algorithmic search content into people’s online social information seeking activities – has substantial advantages. Namely, it preserves many more of the benefits of status message question asking and it supports people’s natural information seeking behavior.

In this paper we present a prototype system, *SearchBuddies*, which takes this inverted approach to integrating social and algorithmic search. SearchBuddies embeds itself in a person’s Facebook network and provides algorithmic answers to questions posed via Facebook status messages. For example, in Figure 1, SearchBuddies provided Elise with a pointer to a webpage that answers her question and a list of her friends who have lived in Hawaii in case she would like to follow up with them.

We developed and deployed SearchBuddies with a focus on understanding how question askers and their friends interact with algorithmic content inserted into their conversations. We describe how people responded when the system successfully provided useful information, and also when it provided undesirable responses. Through SearchBuddies, we show that embedding a search engine into a social network is a viable and useful approach to integrating social and algorithmic search. The deployment

¹ http://www.bing.com/community/site_blogs/b/search/archive/2011/05/16/news-announcement-may-17.aspx

² <http://www.google.com/insidesearch/plus.html>

of the system also illuminates the unique challenges and opportunities inherent to socially embedded search engines like SearchBuddies, such as adapting search engines' algorithms to the social norms and contexts of social network users.

Related Work

Social search is an active area of research. Evans and Chi (2010) presented a model of social search based on existing non-social search models. Evans et al. (2009) examined the usefulness of three types of online social search: directed asking, public asking, and searching repositories of user-generated content. They found that using all three strategies outperformed any single strategy alone. Through its approach of integrating web search results into online conversations, SearchBuddies makes it easier to employ all three strategies at once.

SearchBuddies embeds itself in people's social search processes by answering questions posed via the status message field in popular social networks. Morris et al. (2010b) found in a large survey that over half of respondents had asked a question using a status message, and approximately two-thirds had answered such a question. A similar characterization study focusing on Twitter suggests that the question asking and answering norms on different social networks may vary significantly (Paul et al. 2011). Yang et al. (2011) identified cultural factors that affect the kinds of questions people ask via status messages, and Teevan et al. (2011) reported on a controlled experiment that tested the effect of status message question phrasing on response quality.

While web search has been shown to identify relevant content faster and more successfully than status message question asking for certain information needs (Morris et al. 2010a), status message questions serve different informational purposes (e.g., they are often used to find opinions and recommendations) and have additional social benefits such as tie-building (Efron and Winget 2010; Morris et al. 2010b). Search engines like Bing and Google have made efforts to bridge social and algorithmic search by incorporating social data into their existing web search results. SearchBuddies is designed to serve as a bridge in the opposite direction.

There is an extensive literature on embodied and conversational agents that informs SearchBuddies (e.g., Cassell 2000; Isbister et al. 2000). Similarly, answering natural language questions and finding experts on topics are active areas of research (Bernstein et al. 2009; Dumais et al. 2002; Feng et al. 2006; Ferrucci et al. 2010). However, the focus of SearchBuddies is not on how to optimally answer questions but rather on providing an initial instantiation of a socially embedded search engine and understanding the social interaction that can take place around algorithmic search results.

Finally, while we are unaware of any attempts to fully embed a search engine within a social network, simple automated agents have been built to respond to questions posed on Twitter (Metcalf 2010; Horowitz 2009).

Design and Implementation

We now describe how we developed the SearchBuddies system, covering the mechanisms of how it participates in people's conversations, which questions it decides to answer, and what it says when it does. Our design strategy was to seek satisficing solutions to the system's significant algorithmic challenges. While our approach performed well in practice, any search algorithm is imperfect at times. As such, SearchBuddies' mistakes help us understand user interaction with a socially integrated search engine as much as its successes do.

Participating in the Conversation

SearchBuddies must be aware of people's conversations in order to participate in them. Conversations on social networking tools occur in many diverse fashions and forums. We implemented SearchBuddies on Facebook, and designed it to operate via Facebook's core interaction: the status message update. When the system identifies an answer to a status message question, the system responds the same way most people would respond, via a public reply. See Figures 1-4 for examples.

People register with the SearchBuddies system using a registration page that describes SearchBuddies, its terms of use and privacy policy, a FAQ, and information about unregistering. We developed two approaches for the SearchBuddies system to respond to people's questions. Each approach, described in greater detail below, has its own Facebook account which must be "friended" by the user (an architecture combining an app with an account was necessary to enable the SearchBuddies' replies to appear inline alongside "regular" replies). While Facebook's Terms of Service typically require accounts to belong to real people, we received a waiver to create several test accounts for the purpose of this research project.

When to Respond

Although people typically engage in many online conversations, not all of these conversations would benefit equally from the inclusion of algorithmic content. For instance, while a search engine is helpful in Figure 1, it might not be able to add much to a post thanking a user's friends for birthday wishes. SearchBuddies must therefore identify when it is useful for it to interject information. The system takes the straightforward and relatively conservative approach of only considering updates that contain question marks. This approach was also used by Paul et al. (2011), who note that it has been shown to have very high precision

(Cong et al. 2008). During initial trials, we found that more sophisticated heuristics like those used by Efron and Winget (2010) created too many false positives in the SearchBuddies context.

The system must also decide when to post a response if a question is detected. Because people are willing to wait up to a day for answers to status message questions (Morris et al. 2010b), responses do not need to appear immediately. However, to remove a source of variation, SearchBuddies responds as soon as it receives a status update notification.

What to Say

The content people find via social search has been shown to fall into two buckets (Morris et al. 2010a): 1) answers, such as what is found via Google or Bing, and 2) people who can provide answers, such as what is found via systems like Aardvark (Horowitz and Kamvar 2010) and IM-an-Expert (White et al. 2011). Within SearchBuddies, we implemented these two approaches: *Investigator*, which returns direct answers, and *Social Butterfly*, which returns pointers to relevant social contacts. Facebook requires account names to appear orthographically “human,” so the SearchBuddies system uses “Investigaetore” for Investigator, and “Soshul Butterflie” for Social Butterfly.

Investigator

Investigator is implemented as an interface to a major web search engine’s public API. The underlying search engine is designed to handle natural language queries, and Investigator makes use of this by submitting the user’s entire status message as a query. It then selects the most relevant search result and posts a message using the result’s (shortened) link and title as its response (e.g., “This page about ‘Good Chinese Food in Charlotte’ may have relevant information: <http://bit.ly/i5IJZS>” in Figure 2b).

To increase precision, Investigator filters the results using a set of 31 whitelisted domains. If none of the top three results come from these domains (or no results are found at all), no answer is posted. The whitelist was developed using a dataset of status message questions from Morris et al. (2010b). Each of the questions was issued to a search engine, and the domains that returned the most relevant results as determined by a human judge were added to the whitelist. Whitelisted domains include yelp.com, cnet.com, and wikipedia.org.

Social Butterfly

While Investigator connects people with information, Social Butterfly connects people with other people who may have the desired information (Figure 3). Social Butterfly identifies topics in a status message question and finds friends of the question asker who have expertise on the mentioned topics.

There are many complex approaches to topic modeling, but most do not work well with short status message questions. While there are recent developments in this

space (e.g., Chen et al. 2011), Social Butterfly uses named entity recognition by leveraging three named entity extractors: one trained on Wikipedia pages, one based on OpenNLP, and one trained on Yahoo! Placemaker. Confidence is determined using a simple voting scheme. Once entities are recognized in status message questions, Social Butterfly attempts to match these entities with the expertise of the askers’ friends.

We consider two types of expertise: places and interests. For places, Social Butterfly mines Facebook’s profile fields for the geographic history of each friend (e.g., hometowns,

Ted Logan
since when did amazon VOD become free to Prime members? interesting...
14 hours ago · Like · Comment

SearchBuddies Investigaetore “Amazon to give Prime members Netflix” on news.cnet.com has some relevant info: <http://cnet.co/faAet3>
14 hours ago · Like · 1 person

Samantha Baker Last week
14 hours ago · Like

Andrew Tanner
Anyone know where to get good chinese food in charlotte?
8 minutes ago · Like · Comment

SearchBuddies Investigaetore This page about “Good Chinese Food in Charlotte” may have relevant information: <http://bit.ly/i5IJZS>
7 minutes ago · Like

Andrew Tanner So cool that this works.
6 minutes ago · Like

Figures 2a (top) and 2b (bottom). Two examples of Investigator responses to questions from Facebook users.

Samantha Baker
Any recommendation for a fun place to go dancing in Seattle on Saturday night?
11 hours ago · Like · Comment

SearchBuddies Soshul Butterflie Robbie Gould (<http://on.fb.me/AeR472Q>) and Elise Clayton (<http://on.fb.me/Tr42Z>) are fans of dancing.
11 hours ago · Like

Robbie Gould :)
11 hours ago · Like

Robbie Gould Depends on the type of dance. So what kind of dancing were you thinking about? Also, Happy Birthday!
about an hour ago · Like

Elise Clayton
Suggestions for a restaurant in San Francisco?
17 minutes ago · Like · Comment

SearchBuddies Investigaetore This page about “Restaurant Suggestions” may have relevant information: <http://bit.ly/gyjBxn>
16 minutes ago · Like

SearchBuddies Soshul Butterflie These friends have lived in or near San Francisco: Bobby Finstock (<http://on.fb.me/6qZp92>), Lisa Marconi (<http://on.fb.me/y887pQ>), Mick Macalester (<http://on.fb.me/fmXsoQ>), and Kirk Lolley (<http://on.fb.me/Z7qQW2>)
14 minutes ago · Like

Figure 3a (top) and 3b (bottom). Two Social Butterfly responses.

schools, or workplaces). For interests, we examine favorite movies, music, books, activities, and other interests. When relevant friends are identified, they are listed in a reply to the question. If Social Butterfly cannot identify any places or named entities in the question, or cannot find a relevant friend, it does not respond.

To generate the specific text of both the Investigator and Social Butterfly SearchBuddies' responses, we wrote a number of answer templates for each response type, one of which was randomly selected for each response.

SearchBuddies Deployment

During a 67-day sign-up period, 122 people registered with the SearchBuddies system. We initially invited colleagues, friends, and family to sign up. Once SearchBuddies began responding to their questions, we observed members of the initial users' social networks register as well. Relationships initiated via personal invitation and word of mouth account for 82% of users. The use of "snowball" sampling can result in a bias toward tech-savvy users. To reach a more representative population, we also invited people via a Facebook advertisement and an email to a usability study recruitment mailing list, resulting in the remaining 18% of users. All users in our analysis were friends with the SearchBuddies for at least one week prior to data collection cutoff, and 89% were friends with them for over a month.

When people registered with SearchBuddies, we collected basic information via a brief opt-in survey. Among the 57 users (47%) who reported gender, 39% were female and 61% male. The mean age was 40. 23% percent of respondents reported asking status message questions at least once a week. Looking at participants' Facebook account data, we found that median network size was 277.5 friends.

In addition to collecting status updates and responses, we also collected natural user feedback about SearchBuddies, much in the way a web search engine collects search result clicks. We did this by leveraging the rich social feedback native to Facebook. For instance, people can "like" status messages and their responses (a form of positive feedback), or delete them (a form of negative feedback). We also employed bit.ly to track link clicks whenever the system provided a link as part of an answer.

Deployment Statistics

The SearchBuddies system registered 1,692 status updates during its deployment. Of these, 262 updates from 78 users were identified as questions by SearchBuddies. The median number of questions per user was 2, and the maximum was 17. Of these 262 status updates, two judges determined 190 to actually be questions (97% agreement, with ties referred to a third judge). Most of the 72 false

Description of Status Message Subtypes	Median # Responses	Human Only
All status messages	0	0
Status message questions (SMQ)	2	2
SMQs (automatically identified)	2	1
SMQs answered by SearchBuddies	2	1
SMQs answered by Butterfly	3	2
SMQs answered by Investigator	2	1

Table 1. The median number of responses to status message questions posted by our users. Medians rather than means are reported since the distribution is heavily skewed.

positives were rhetorical questions, which is not surprising given the results of previous work (Morris et al. 2010b; Paul et al. 2011). We also saw a handful of false negatives, most due to the fact that some status question messages are phrased as statements or have non-standard grammar. For instance, SearchBuddies did not classify "Anyone wanna mayb go see a movie or sumthin" (*sic*) as a question.

Investigator answered 58 questions, or 22.1% of the automatically identified questions (see Figures 1 and 2). Social Butterfly answered 70, or 26.7% (see Figures 1 and 3). Thirty of Social Butterfly's responses were place-based and 40 were interest-based. Twenty-two questions were answered by both Investigator and Social Butterfly.

The median number of responses posted by humans to a status message was zero, with the maximum being 25 (see Table 1). Using a Mann-Whitney test, we found status messages that were questions had significantly more human responses (median = 2) than other status messages ($p < .01$). However, over 36% of status message questions still had no responses. As such, there is a potential to provide answers to unaddressed status message questions as well as a chance to engage in conversations with question askers and their friends around their information needs.

Useful, Complementary Information Provided

Our primary goal in implementing SearchBuddies was to provide question askers and their friends with algorithmic insight in the context of their natural online conversations. Observing the conversations in which SearchBuddies participated, we saw that SearchBuddies was able to contribute supplementary and relevant content as predicted. Users gave explicit feedback to this effect, both using natural language and by "liking" SearchBuddies' posts.

For example, in response to Investigator, one question asker wrote, "Yes, that's what I was looking for," and another posted, "by golly, that's a great recommendation." The latter was in response to an Investigator answer to a rhetorical question ("Can poop qualify as a character in a book?"); the Investigator post contained a link pointing to the popular children's book *Everybody Poops* (Gomi 2001) on Amazon.com, suggesting that algorithmically retrieved information may even be useful in the case of rhetorical questions. Figure 2b contains another example of Investigator contributing valuable information to a conversation. There were also several instances of Social Butterfly receiving similar feedback from question askers

when it suggested relevant friends. For instance, in response to a Social Butterfly recommendation, one user wrote “Thanks Butterflie. [Friend] might be interested.”

Question askers were not the only ones to give SearchBuddies positive feedback; their friends did as well. For instance, after the asker commented that Investigator had given a great suggestion with *Everybody Poops*, one of the asker’s friends noted that she had read the book and found that it was “a delight.” Moreover, three of Investigator’s 58 responses received a total of four “likes,” all for highly relevant responses. Social Butterfly also received four “likes,” two of which were for relevance. We discuss the reasons for the remaining two below.

We observed some evidence that Social Butterfly was able to elicit additional participation in a conversation using its algorithmic approach of recommending people. For instance, Figure 3a gives an example of a friend who was mentioned by Social Butterfly chiming into the discussion, both with informational content and with utterances of purely social intent (e.g., “:-)”). Another user wrote, “you caught me online :-)” after Social Butterfly recommended her (Social Butterfly includes information about online status in its posts). Another friend noted that he thought being mentioned was “kinda creepy,” but entered the conversation with information anyway. More generally, when Social Butterfly replied, we observed a higher median number of human responses (Table 1), although this difference is not significant ($p = 0.14$).

Finally, users also gave implicit feedback that they were interested in what Social Butterfly and Investigator had to say. 57.4% of Investigator’s bit.ly links were clicked at least once, as were 61.0% of Social Butterfly’s.

Bad Responses Deleted, Discussed, and Ignored

All search algorithms return undesirable results at least on occasion, and understanding interaction with these results in a social setting was an important goal of our deployment. We saw four types of reactions to undesired responses by SearchBuddies: deleting, explicit commenting, joking, and ignoring.

Investigator made 11 (18.9%) posts that were deleted by users. Qualitative post-hoc analysis suggests most deletions were due to low relevance. An example of an Investigator response deleted for relevance was a post that pointed to the Wikipedia page on dim sum following a request for a San Francisco dim sum restaurant recommendation.

Seventeen (24.3%) Social Butterfly posts were deleted. While some of these deletions were clearly due to an undesired response to a rhetorical question or entity disambiguation issues, the reasons for most deletions could not be determined without more social context, an issue we address in the following section.

We also found, however, that users often eschewed deleting undesirable posts in favor of commenting about them (sometimes quite harshly) in the same thread as their

question. When Investigator responded to the question, “Where should Matt and I go April 21 - 25?” with an irrelevant pointer to the Wikipedia page about a famous North Pole explorer (Matthew Henson), the question asker replied to the thread with “[Investigator]! No! Unrelated.” Another user posted, “I might need to do something about [SearchBuddies], it's acting like the annoying friend you silently delete.” Similarly, one user wrote “Oh hell. Die ... bot” in response to an irrelevant Social Butterfly result.

On the other hand, irrelevant posts made by SearchBuddies were received quite positively if they could be interpreted humorously. Two of the four “likes” Social Butterfly received were almost certainly due to humorous errors. For instance, in one of these humorous “liked” cases, when a user asked a question about the Queen Anne neighborhood in Seattle, Social Butterfly responded with a friend recommendation that began with “Some of your friends like queen: ...” (the friends were fans of the 80s rock band rather than the neighborhood).

Occasionally users and their friends would just ignore irrelevant posts by SearchBuddies (at least in the context of the Facebook conversation). For instance, when Investigator posted a link to an article about iPhones in general in response to a question about the iPhone autocomplete functionality, an entire conversation ensued without reference to (or deletion of) Investigator’s error.

SearchBuddies Anthropomorphized

An additional interesting trend we saw in our deployment was the extent of anthropomorphization of the SearchBuddies, a result that might be predicted by Nass et al. (1994). Examples included, “No, wrong, [SearchBuddy]!” and “Well, hello, [SearchBuddy]! What a surprise to hear from you on this topic.” Often when SearchBuddies was anthropomorphized the context was derogatory. These derogatory responses were usually posted by the original question asker and done when the SearchBuddies post had low relevance. This may have been a way of signaling to other friends – who did not necessarily know that a SearchBuddy was not a real person – how to respond to and understand its undesired replies.

Social Context is a Vital Consideration

Relevance is the predominant evaluation metric for search engines, and our discussion of deployment results thus far has examined interaction with SearchBuddies through this lens. However, because SearchBuddies is deployed in a social context, we observed that it was also important to understand how users interacted with the system based on purely social dimensions.

Our results indicate that for search engines in a social setting, understanding social context may be as important as identifying relevant information. The stories behind two

deleted Social Butterfly posts are particularly informative. In the first case, one user wrote to us:

“I don't like that [Social Butterfly] posts names of some of my Facebook friends on my wall as an answer to my restaurant query. I don't know them that well and feel they may not appreciate their names up there. I will probably take that comment down to get their names off my wall.”

Another user indicated that Social Butterfly had referred to two of her friends in the same post who had recently been through a difficult break-up. This user subsequently deleted the post to avoid upsetting them. In both cases, relevance was not a factor in the perceived undesirability of the post. Instead, it was social considerations that caused these users to delete the Social Butterfly responses.

Implications for Design

SearchBuddies represents a new way to support information seeking by integrating algorithmically identified content into people's pre-existing social information seeking behavior. We refer to algorithmic tools that participate in social search as *socially embedded search engines* and are unaware of a full-fledged socially embedded search engine other than SearchBuddies. While our results demonstrate that a socially embedded search engine can be a viable and useful tool for integrating social and algorithmic search, they also highlight that this approach requires tackling unique challenges not present when considering either search paradigm individually. However, if these challenges can be overcome, there is the opportunity to merge the benefits of each paradigm to answer people's questions in a way that neither would be able to do alone. Below, we highlight these novel challenges and opportunities.

New Challenges

Novel Search Result Quality Metrics Needed

Evaluating the quality of search results in a social setting is problematic. Our findings indicate that while traditional relevance-focused search metrics (e.g., precision and recall) are important for algorithmic search in a social context, they must be balanced with metrics that account for social appropriateness, which we call *conformance* metrics. The people who used SearchBuddies cared not only about what information was provided to them by the system, but also about how that information was perceived by others, who was mentioned, and how mentioned people were related. These issues are captured in conformance metrics, which measure the extent to which an algorithmic response conforms to the social norms of the questioner, the questioner's social circle, and the social network. Even a system that returns perfect answers at every possible opportunity can be harmful if it does not conform.

Developers of socially embedded search engines will need to design algorithms that display good conformance, but doing so is likely to be difficult. Consider the two examples in which Social Butterfly returned answers that were deleted for the entirely social reasons of not wanting to bother acquaintances and avoiding an awkward situation involving the ending of a romantic relationship. If SearchBuddies had a means of detecting sensitive relationships and identifying shyness in its users, it could have done its job better, perhaps by offering a different answer or by answering via a private message (i.e., as in Krulwich and Burkey 1996). Although accessing some information about a user's friends is permissible for Facebook applications (contingent on a user's privacy settings), understanding how to utilize such information appropriately is important.

SearchBuddies also likely missed other categories of conformance signals. For example, many of the additional deleted Social Butterfly responses seemed quite relevant in post-hoc analysis (although we cannot be sure without full knowledge of social context), and thus the deletions were likely due to factors other than search result quality issues. Algorithms built into socially embedded search engines need to be designed to detect as many of these conformance signal types as possible. Active research in tie strength detection and other types of social network mining may provide a valuable starting point (e.g., Gilbert and Karahalios 2009).

Experience from our deployment suggests that relevance and conformance are interrelated. For example, the need for conformance can result in a greater need to avoid posting results with low relevance. Wrong answers can be the source of massive and public critique in a social environment. Algorithmic search engines often fail, but rarely are they *publicly* told to “die”.

However, just as relevant results do not necessarily conform to social contexts and norms, sometimes irrelevant results *do* conform. We saw this phenomenon with SearchBuddies when Social Butterfly was lauded by its users and their friends for returning irrelevant results that were interpreted humorously. Indeed, humor may be an important tool when designing for good conformance.

Higher Relevance Thresholds Needed

Where humor is not possible, the deployment suggests that users of socially embedded search engines require results of very high relevance. Although Investigator only returned links that were in the top three results of a query issued to a major search engine, users sometimes found those results of low enough quality to delete them and publicly criticize them. This is true even though Investigator included an additional quality filter in the form of its domain whitelist. Our findings suggest that socially embedded search engines should adopt a “when in doubt, leave it out” approach to returning results. This is different from the standard web search experience, where returning

no results is considered a significant failure. How high the relevance bar must be for a socially embedded search engine to return content is currently unknown. A Wizard of Oz study where answer quality is human-controlled may be a good way to identify realistic targets.

New Opportunities

Although this new area of socially embedded search engines faces the many challenges listed above, it also enables new opportunities. Namely, our results suggest that friends and search engines may be able to work together to help people find information better than they could with either approach (or both approaches) individually.

Ability to Frame the Discussion

Our results indicate the potential for socially embedded search engines to frame the discussion that takes place around a question in addition to providing direct answers to the question. This is not something that traditional web search engines can do. Consider the *Everybody Poops* conversation, for instance. Investigator's response steered the conversation from being rhetorical in nature to becoming a conversation about a specific children's book. Designers of socially embedded search engines could take advantage of this capability to make conversations more informative, more social, or both. For example, it has been shown that if the first response to a Facebook question contains an answer, future replies are more likely to be social and less likely to answer (Morris et al. 2010b). If increasing social interaction was considered important to its designer, a socially embedded search engine could choose to answer quickly, possibly encouraging human responses to focus on the social tie-building.

New Modes of Interaction Have Benefits for Search

The conversational nature of the environment in which socially embedded search engines operate presents the opportunity for these search engines to engage users in new modes of interaction. For example, in the case when available contextual information is insufficient for an algorithm to effectively contribute, socially embedded search engines could ask for feedback and iterate. If a user inquires about a good place to experience traditional Irish cuisine in Dublin, a system like SearchBuddies might ask, "Are you looking something cheap or expensive?" Explicit user feedback has the potential to drastically improve result quality, but in practice search engine users rarely provide such information (Anick 2003). In contrast, clarifying dialog is common for conversational questions.

There is also the opportunity to more fully engage in the conversation with all of the participants. A socially embedded search engine does not only need to answer questions directly, by, for example, recommending a Dublin restaurant to a user. It also could supplement the asker's human friends' responses by providing related content. For example, if a friend were to first suggest The



Figure 4. An example of how socially embedded search engines could be used to post supplementary information in addition to answering questions directly.

Pig's Ear restaurant, the system could add, "The Pig's Ear gets 4 stars on Yelp." (Figure 4)

Further, as we observed with SearchBuddies, conversations around status message questions occur at speeds orders of magnitude slower than that typically considered by web search engines (Morris et al. 2010b). Most search algorithms are limited by the need to provide responses at near instantaneous speeds. Within a social context there is an opportunity for a search engine to take more time and devote more resources to the questions than can be done in the few milliseconds typically allotted following a web search query. For example, a socially embedded search engine could employ complicated and slow algorithms, crawl new Web content, use crowdsourcing to find answers (e.g., Bigham et al. 2010), or poll a person's friends – and still provide a sufficiently timely answer. This could greatly increase the relevance of the output relative to that of traditional web search (and could help conformance as well).

Additional Rich Feedback Mechanisms

As noted above, sources of feedback in traditional web search can be limited (e.g., link clicks and query modifications). We saw in the SearchBuddies deployment that socially embedded search engines have a large variety of strong feedback signals available to them. While clicks can still be measured at scale (using, for example, a bit.ly strategy), so can the number of post deletions, post "likes," and the quantity and sentiment of follow-up comments by the questioner and friends. We observed that each of these feedback mechanisms was commonly used during SearchBuddies' deployment. Importantly, some of these mechanisms provide negative feedback, which can be quite difficult to acquire at a large scale for web search engines.

The SearchBuddies deployment also provides guidelines for obtaining feedback on conformance, not just relevance. "Likes," deletions, and the types of social signals present in the follow-up conversation can all contain conformance feedback signals. For instance, we saw positive conformance feedback in both "likes" and the follow-up conversation, although in the case of "likes" this came at the expense of the signal quality for relevance (such as in the cases of humor). Negative feedback can come both in the form of deletions and the character of the follow-up responses (e.g., "oh hell, die...bot").

Conclusion and Future Work

In this paper, we presented the SearchBuddies system, a prototype socially embedded search engine that identifies and replies to status message questions on Facebook. Through a three-month deployment of the system with 122 users, we observed that socially embedded search engines enable friends and search tools to work together to find answers to people's questions in a way that neither would be able to do alone. We also identified the unique challenges posed by this new type of search experience. For instance, we introduced the notion of conformance metrics and showed how they must be considered in addition to the traditional relevance metrics of web search.

Finally, we noted specific ways that socially embedded search engines can help search systems break out of the box of instantly ranked "ten blue links" search results. We discussed how this new type of search engine can make use of social context, extend response time windows, and iterate with explicit and implicit feedback from the asker and other parties in a conversation. In doing so, socially embedded search engines may be able to provide answers that would be more relevant than those provided through typical web search interactions.

Going forward, we are using what we learned to target improvements to the SearchBuddies' question answering algorithms. To determine how relevant and conforming these responses need to be, we plan to conduct a Wizard of Oz study. We are also particularly interested in exploring question answering approaches that make use of or augment other people's replies to an initial question. Finally, we are studying how search algorithms might make use of additional time to devote more resources or take slower approaches to coming up with better answers.

Acknowledgements

We thank Darren Gergle and our reviewers for their valuable suggestions. We also thank users of SearchBuddies for their participation. This work was partially supported by an NSF GRFP award.

References

- Anick, P. "Using terminological feedback for web search refinement." *SIGIR '03*.
- Bernstein, M., Ackerman, M., Chi, E., Miller, R. "The trouble with social computing systems research." *CHI '11 Ext. Abstracts*.
- Bernstein, M., Tan, D., Smith, G., Czerwinski, M., and Horvitz, E. "Collabio: a game for annotating people within social networks." *UIST '09*.
- Bigham, J., Jayant, C., Ji, H., Little, G., Miller, A., Miller, R., Tatarowicz, A., White, B., White, S., and Yeh, T. "VizWiz: Nearly Real-time Answers to Visual Question." *UIST '10*.
- Cassell, J. 2000. "More than just another pretty face: Embodied conversational interface agents." *CACM* 4 (43): 70-78.
- Chen, J., Nairn, R., and Chi, E. "Speak Little and Well: Recommending Conversations in Online Social Streams." *CHI '11*.
- Cong, G., Wang, L., Lin, C., Song, Y., and Sun, Y. "Finding question-answer pairs from online forums." *SIGIR '08*.
- Dumais, S., Banko, M., Brill, E., Lin, J., and Ng, A. "Web question answering: is more always better?" *SIGIR '02*.
- Efron, M., and Winget, M. "Questions are Content: A Taxonomy of Questions in a Microblogging Environment." *ASIS&T '10*.
- Evans, B., Kairam, S., and Pirolli, P. "Exploring the Cognitive Consequences of Social Search." *CHI '09*.
- Evans, B., and Chi, E. 2010. "An elaborated model of social search." *Information Processing & Management* 46 (6).
- Ferrucci, D., Brown, E., Chu-Carroll J., Fan, J., Gondeck, D., Kalyanpur, A. et al. "Building Watson: An Overview of the DeepQA Project." *AI Magazine* (Fall 2010): 59-79.
- Gilbert, E., and Karahalios, K. "Predicting tie strength with social media." *CHI '09*.
- Feng, D., Shaw, E., Kim, J. and Hovy, E. An intelligent discussion-bot for answering student queries in threaded discussions. *IUI '06*.
- Gomi, Taro. 2001. *Everyone Poops*. Kane/Miller Book Pub.
- Horowitz, Damon. "Send questions to Aardvark from Twitter." *Aardvark Blog*. <http://blog.vark.com/?p=107>.
- Horowitz, D., and Kamvar, S. "The anatomy of a large-scale social search engine." *WWW '10*.
- Isbister, K., Nakanishi, H., Ishida, T. and Nass, C. Helper agent: designing an assistant for human-human interaction in a virtual meeting space. *CHI '00*.
- Krulwich, B. and Burkey, C. The ContactFinder agent: Answering bulletin board questions with referrals. *AAAI '96*.
- Metcalfe, John. 2010. "Climate Cyborg brings global-warming debate to Twitter." *ABC StormWatch* 7. <http://www.wjla.com/blogs/weather/2011/02/climate-cyborg-brings-global-warming-debate-to-twitter-8599.html>.
- Morris, M., Teevan, J., and Panovich, K. 2010a. "A Comparison of Information Seeking Using Search Engines and Social Networks." *ICSWM '10*.
- Morris, M., Teevan, J., and Panovich, K. 2010b. "What do people ask their social networks, and why?: a survey study of status message Q&A behavior." *CHI '10*.
- Nass, C., Steuer, J., and Tauber, E. "Computers are social actors." *CHI '94*.
- Paul, S., Hong, L., and Chi, E. "Is Twitter a Good Place for Asking Questions? A Characterization Study." *ICWSM '11*.
- Teevan, J., Morris, M., and Panovich, K. "Factors Affecting Response Quantity, Quality, and Speed for Questions Asked via Social Network Status Messages." *ICWSM '11*.
- Thom-Santelli, J., Helsey, S., Matthews, T., Daly, E., Millen, D. "What Are You Working On? Status Message Q&A in an Enterprise SNS." *ECSCW '11*.
- White, R., Richardson, M., and Liu, Y. "Effects of community size and contact rate in synchronous social Q&A." *CHI '11*.
- Yang, J., Morris, M., Teevan, J., Adamic, L., and Ackerman, M. "Culture Matters: A Survey Study of Social Q&A Behavior." *ICWSM '11*.