

Do You Want to Know? Recommending Strangers in the Enterprise

Ido Guy*, Sigalit Ur*, Inbal Ronen*, Adam Perer**, Michal Jacovi*

*IBM Haifa Research Lab
Mt. Carmel
Haifa 31905, Israel
{ido, sigalit, inbal, jacovi}@il.ibm.com

** IBM Research
One Rogers Street
Cambridge, MA 02142, USA
adam.perer@us.ibm.com

ABSTRACT

Recent studies on people recommendation have focused on suggesting people the user already knows. In this work, we use social media behavioral data to recommend people the user is not likely to know, but nonetheless may be interested in. Our evaluation is based on an extensive user study with 516 participants within a large enterprise and includes both quantitative and qualitative results. We found that many employees valued the recommendations, even if only one or two of nine recommendations were interesting strangers. Based on these results, we discuss potential deployment routes and design implications for a stranger recommendation feature.

Author Keywords

People recommendation, social matching, social networks, social media.

ACM Classification Keywords

H.5.3. Group and Organizational Interfaces – Computer-supported cooperative work.

General Terms

Design, Experimentation, Human Factors, Measurement

INTRODUCTION

The Internet enables individuals to maintain existing social ties and develop new ones with people who share similar interests [33]. The emergence of the social web introduces new opportunities for people to interact and discover those with similar interests. As users of the social web join online communities and contribute content

(as in wikis and blogs) and metadata (such as tags, comments, and ratings), new ways of forming and maintaining relationships are becoming possible.

Social network sites (SNSs), such as Facebook, MySpace, Orkut, and Friendster, allow users to explicitly define their social network by sending and accepting invitations to connect. These explicit networks enable the sharing and diffusion of photos, music, applications, status updates, and more. Previous research on SNSs has found that people primarily connect to individuals they already know, and are less likely to approach strangers to initiate a connection [1,17].

SNSs have also emerged within enterprises. Research indicates that in addition to staying in touch with close colleagues, employees use enterprise SNSs to reach out to employees they do not know and build stronger bonds with their weak ties. Their motivations include connecting on a personal level with more coworkers, advancing their career within the company, and campaigning for their ideas [5]. The same study also recommends that “enterprise social software specifically supports users in discovering new colleagues through exploration and searching around common interests.”

In this work, we suggest a novel method for recommending strangers in the enterprise with whom the user shares similar interests. Our approach actively “brings” new people to the user, in contrast to the “exploration and search” approach, and can be viewed as an enterprise instance of a social matching system [32]. Connecting to strangers within the organization can be valuable for employees in many ways: get help or advice [4], reach opportunities beyond those available through existing ties [9], discover new routes for potential career development, learn about new projects and assets they can reuse and leverage, connect with subject matter experts and other influential people within the organization, cultivate their organizational social capital [25,29], and ultimately grow their reputation and influence within the organization.

Recently, leading SNSs such as Facebook and LinkedIn have added “People you may know” features to their homepages, suggesting new connections [23,24]. The recommendations are principally based on common

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW 2011, March 19–23, 2011, Hangzhou, China.

Copyright 2011 ACM 978-1-4503-0556-3/11/03...\$10.00..

friends and group co-memberships, and are mostly aimed at highlighting people the user knows well. Lately, people recommendation within enterprise SNSs has been explored [3,12]. These studies show that a people recommender can have a high impact on the number of connections within the site and the number of people who are initiating invitations. Yet, as pointed out [12], users exhaust the list of potential connections quite rapidly, and are interested in recommendations of people they do not yet know.

Our recommender infers mutual interests from shared activity in social media, such as bookmarking the same pages, use of the same tags, and membership in the same communities. According to previous studies [10], a user's similarity network mined this way often includes members of the user's familiarity network, i.e., the set of people the user knows well. As our goal is to recommend people who are not known to the user, we subtract the user's familiarity network from the user's similarity network to yield a list of recommended people who are similar yet strangers.

The task of recommending unfamiliar yet interesting people is quite different from existing people recommendation tasks, which focus on recommending individuals who are known to the user. Our recommender focuses more on discovery and exposure to new people and less on facilitating an action such as connecting on an SNS. It aims at satisfying two rather conflicting goals: on the one hand, the recommended person should not be familiar to the user, and, on the other hand, the person should be of some interest. While accuracy of recommendations that satisfy both goals might not be high, we argue that the potential serendipity and "surprise effect" in getting a fortuitous recommendation of an interesting new person in the organization may compensate for lower accuracy [21].

In our main experiment, we presented 11 recommended people to each of the 516 participants. Recommendations include two benchmarks that represent extreme cases – a strongly familiar person and a random person. For each recommended individual, participants were asked to indicate how familiar they were with that person and how interesting they found the recommendation. Previous studies on people recommendation, focusing on recommending familiar people, expected a "connect" action to indicate a successful recommendation [3,12]. In our case, it is unlikely that users would immediately connect with someone they do not know. We thus investigated a set of other specific actions that reflected interest in the recommended person, such as following the person's activities, or browsing his/her files or bookmarks.

Results indicate that not only does our recommender suggest many strangers (over two-thirds of the recommendations), but these strangers often arouse interest, significantly higher than a random stranger. For most participants, at least one or two recommendations suggested a stranger who was also interesting at some level. Many participants valued the scenario of stranger recommendation, and stated different motivations for their interest.

The rest of the paper is organized as follows. In the next section, we discuss related work. We then describe our

recommender engine, followed by a description of the experimental setup. Next, we present qualitative and quantitative results in detail. We conclude by discussing our findings and suggesting future work.

RELATED WORK

Social matching systems are recommender systems that recommend people to one another. Terveen and McDonald [32] define the scope and research agenda for social matching systems and explain how they are different from traditional recommender systems that suggest items such as books or movies. Online dating [6] is probably the most well-known social matching application, mating pairs by personal attributes and preferences [14], photos [27], and other parameters.

Apart from dating, various other social matching applications have been studied. The I2I [2] system aims to facilitate Web user awareness of online resources that are available in the context of their current task. Yenta [7] is a distributed agent-based matchmaker that clusters users with similar interests. Each user's agent infers topics of interest based on documents in that user's file system. The SocialNet application [31] uses patterns of physical collocation over time, retrieved through mobile devices, to infer shared interests between users. Klenk et al. [16] present a social matching system within the healthcare domain that helps detect similar patients based on similar symptoms and comparable diseases. Our work examines social matching in an enterprise setting, where a matching algorithm leverages social media activity to infer similar interests.

Social matching has also been studied in the context of community awareness. Sumi and Mase [30] provide personalized recommendations at a museum or an exhibition through a community-aware PDA application, which includes matchmaking of users with similar interests. McCarthy et al. [19] try to augment the social space at academic conferences by increasing interaction among attendees. Constant et al. [4] discuss the "kindness of strangers". They study information seeking within a large organization and find that strangers provide useful answers and advice, despite the lack of personal relationship with the information seeker.

Expertise location systems (e.g., [20,28]) are commonly approached as social matching systems with regard to a specific topic. Often, as a result of the expertise location task, the user is referred to an expert who is a stranger. However, as opposed to our scenario, this is an ad-hoc encounter with respect to a specific query that has been initiated by the user.

With the rise in popularity of SNSs, both on the web and within enterprises [1,5], various SNSs have studied recommendations for connecting people. The "Do You Know?" widget [12] recommends people to connect to within an enterprise SNS. Its deployment is shown to

increase the number of connections on the site and the number of people sending invitations. Chen et al. [3] study people recommendation on another enterprise SNS and show that algorithms based on social networks outperform ones based on similarity of user-created content. Quercia et al. [26] present a framework that recommends friends to mobile SNS users based on Bluetooth proximity data. Freyne et al. [8] show that people recommendations can be effective in increasing adoption and engagement of new social software users. All of these works focus on recommending people with whom the user is likely to be familiar, and are evaluated mainly by their ability to suggest people the user indeed knows. In contrast, in this work we aim to recommend people the user does **not** know.

RECOMMENDER ENGINE DESCRIPTION

We use Lotus Connections (LC) [15], an enterprise social software application suite, as the environment for our experiments. LC consists of a corporate directory with rich employee profiles, as well as different enterprise social media applications: blogs, wikis, social bookmarking, file sharing, online communities, an enterprise SNS, and a people tagging application that allows users to tag each other. More details on these applications and their level of usage within our organization can be found in [10].

To harvest social relationships, we use the SONAR system [11], which aggregates social network information across different data sources within the organization, in particular across the LC applications mentioned above. For each data source, SONAR computes a relationship score between two individuals in the range of [0,1], where 0 indicates no relationship and 1 indicates the strongest relationship. The different relationship scores are then aggregated to a unified single score using a weighted vector that defines the relative importance of the relationships. SONAR can distinguish between familiarity relationships and similarity relationships. Given a user u and the desired relationship type (familiarity or similarity) it returns a weighted list of people related to u and their unified relationship score with u , ordered by that score.

To extract the user's familiarity list (F), SONAR aggregates the following relationships: (1) explicit connection on the LC SNS, (2) connection via the organizational chart, which is part of the directory service in LC, (3) file sharing, (4) co-editing of wiki pages, (5) people tagging, and (6) co-authorship of patents, papers, and pages in a projects wiki. Calculation of scores for each of these familiarity relationships is based on different factors, such as the number of co-authored items. More details on the scoring calculation can be found in [11,13]. Ultimately, all relationships are aggregated with an equal weight. This configuration and scoring scheme of familiarity relationships was found to be effective in

extracting a representative ranked list of people the user knows [11,12].

Previous work on mining user similarity relationships [10] classifies them into three categories: *places* (e.g., co-membership in a community), *things* (e.g., co-usage of tags), and *people* (e.g., having a mutual friend). While *things* and *places* are found effective in mining user similarity, *people* is the least productive category, and yields a list that mostly overlaps with the user's familiarity list. We therefore opted to initially extract similarity relationships based solely on *things* and *places*. Thus, SONAR aggregates the following relationships to infer the similarity list (S): (1) membership in the same community, (2) commenting on the same blog entry, (3) reading the same file, (4) bookmarking the same page, (5) using the same tag, and (6) being tagged with the same tag. All scores for similarity relationships were calculated using *Jaccard's index*, i.e., by dividing the number of items in the intersection set by the number of items in the union set. All similarity relationships were aggregated with an equal weight. This method for aggregating and scoring similarity relationships has been shown effective in yielding a list of individuals who share common interests with the user and is described in detail in [10].

As for the *people* category, we opted to experiment with the following three alternatives: (a) include *people* relationships in S only, (b) include them in F only, and (c) exclude these relationships from both S and F. We expected that these alternatives would allow us to experiment with different levels of filtering of the similarity list, possibly trading between the likelihood of recommending a stranger and the likelihood of recommending someone of interest. The *people* category includes the following relationships: (1) having the same friend on the LC SNS, (2) tagging the same person, and (3) being tagged by the same person. Like the other similarity relationships, scores were calculated by Jaccard's index.

Our recommender engine, called *StrangerRS*, used SONAR to retrieve the user's top 150 similar people and top 150 familiar people. It then removed all individuals on the similarity list that also appear on the familiarity list, to yield the final list of recommended people who are presumably similar but unfamiliar to the user. We defined three experimental groups: (1) $S+p-F$ includes the *people* category in the similarity relationships; (2) $S-F+p$ includes *people* in the familiarity relationships; and (3) $S-F$, which does not include *people* at all.

SONAR was also used to retrieve the "evidence" for each recommendation, which included all things, places, and possibly people that the user shares with the recommended person. For example, evidence can include three communities in which both users are members (including links to these communities' homepages) and five tags with which both users have been tagged.

EXPERIMENTAL SETUP

Our evaluation is based on a user study where participants were asked to evaluate 11 recommended people. Nine of the recommendations were retrieved through *StrangerRS* in the following manner: we retrieved the top 30 recommended people, as explained in the previous section, and drew at random 3 recommendations out of the top 10, another 3 out of the recommendations ranked 11-20, and another 3 out of the ones ranked 21-30. This way, we could evaluate the quality of recommendations further down the list and compare it with the quality of the top ranked recommendations, without overwhelming the participants with too many recommendations. We also included two additional recommendations as benchmarks for two extreme cases: a person the user knows well and a random person. For the well-known recommendation (denoted *StrongFam*), we drew at random 1 of the top 10 people in the user's familiarity list. For the random individual (denoted *Random*), a random person was chosen from the corporate directory that contains all active employees in the organization. The order of the 11 recommendations was randomized per participant.

Figure 1 depicts the user interface for a recommended person in our survey. Existing widgets that focus on recommending familiar people [24] typically include the person's name, photo, and some title (e.g., job role) to describe the recommended individual. While this information is enough for a person you know, in order to arouse interest in a stranger, it can be valuable to include more details. Hence, we opted to present the profile page in LC, which is rich with details and to which LC users are already accustomed. As shown in Figure 1A, the profile includes the person's name and photo, and details such as their country, role, office address, phone number, and email. It also includes the person's management chain and a photo collage of the network of "friends", with whom s/he is reciprocally connected (on the right side), the person's "board" with recent status updates and messages from others and recent posts from different LC applications (bottom center), and the current status update and a list of tags applied by other employees (left). The user can scroll through the profile page to see all details.

Figure 1B details the "evidence" for the recommendation, and includes summary counts of the shared artifacts (e.g., 2 communities). The types of common artifacts that can appear as evidence reflect the different similarity relationships described in the previous section, and may include communities, blog entries, files, bookmarked pages, used tags, incoming tags, and people. In addition to the summary counts, the actual title of each artifact is presented (the community's title, the tag, the person's name, etc.) linked to the page of the artifact when appropriate (tags are the only type of artifact for which a link is not provided).

Figure 1C demonstrates the area where participants provided their feedback on the recommendation. In particular, they responded to six different questions on a five-point Likert scale, ranging from 'not at all' to 'very much': (Q1) Are you familiar with this person? (Q2) Is this person of interest to you? (Q3) Would you like to follow this person's tweets? (Q4) Would you like to watch this person's activity on LC? (Hovering over the adjacent question mark invokes a tooltip that explains that this means following this person's status updates, posts to blogs, files, wikis, etc.) (Q5) Would you ever wish to tag this person? (Q6) Would you browse this person's files, bookmarks, or blogs? Free-text comments could optionally be provided for each recommended person, as well as at the end of the survey.

Q1 examined the familiarity level with the person and Q2 asked for a measure of general interest in the recommended person. Since a question about general interest can be interpreted in many ways, we also examined specific interest indications in Q3-Q6. Q3 and Q4 both considered a "follow" scenario: Q3 referred to following on a micro-blogging system, while Q4 referred to following in LC. Q5 referred to tagging the person, and Q6 to browsing that person's artifacts (files, bookmarks, blogs). While Q3, Q4, and Q5 all refer to an action that would be publicly logged and persisted, Q6 reflects a one-time action that is not persisted or publicly exposed, and may thus be a "softer" interest indicator.

Our survey participants consisted of 1,885 LC users who were directly related to at least 30 other people, 30 tags, and 30 documents, as done in [13]. We note that this sample does not represent the entire population of our organizations' employees, but rather active users of enterprise social media, who are the target population for our recommender system. A link to the survey with an invitation to participate was sent by email to each of these 1,885 individuals. Each participant was randomly assigned to one of three groups, receiving recommendations based on $S-F$, $S+p-F$, or $S-F+p$.

EXPERIMENTAL RESULTS

A total of 516 participants completed our survey, originating from 33 countries and spanning the different units within our organization: 32% sales, 26% software, 21% services, 12% headquarters, 4% systems, 3% research, and 2% others. Of these, 172 participants were assigned to the $S-F$ group, 173 got $S+p-F$, and 171 $S-F+p$.

We first report the overlap inspected between the similarity and familiarity networks in order to verify the requisite for filtering familiar people. This was done by observing the position of the last (30th) recommendation in the original similarity list (i.e., before subtracting the

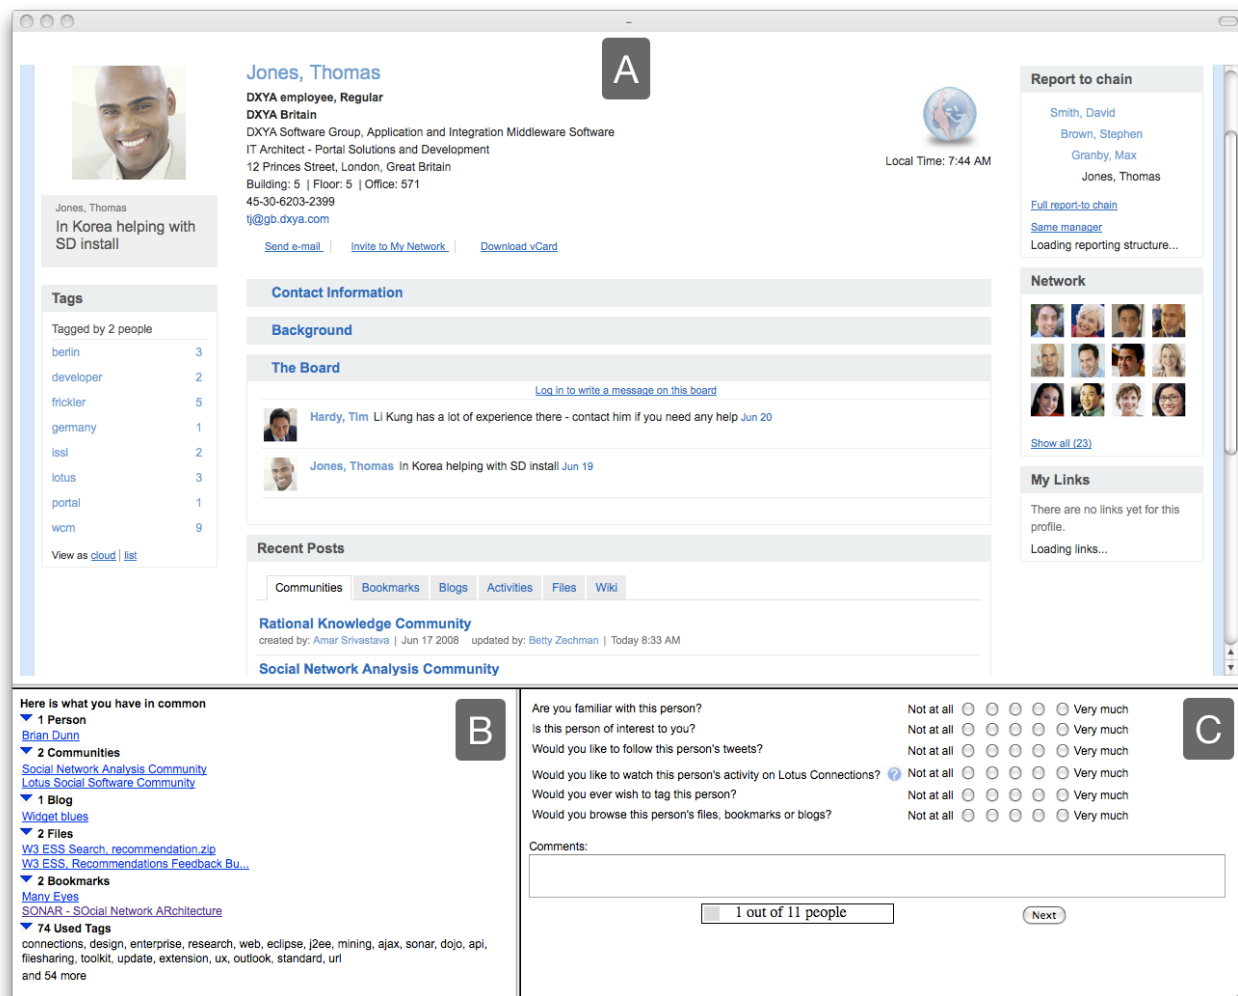


Figure 1. User interface of the stranger recommender system

familiar people). Over all participants, the average position of that recommendation was 39.9 (stdev 9.6, median 37, max: 102), indicating that on average, almost 10 familiar people were removed from the list of top similar people to yield the final recommendation list. This reinforces previous findings about the high overlap between the similarity and familiarity networks [10] and indicates that filtering familiar people is essential. For the *S-F* group, the average was the lowest – 37.4 (stdev 8, median 35, max 71). The addition of the *people* relationships to either side of the subtraction increases the overlap, as these relationships are on the border between similarity and familiarity. For the *S+p-F* group, the average was 43.4 (stdev 9.1 median 43 max 68) and for the *S-F+p* group it was 38.9 (stdev 10.6 median 36, max 102). We assume that the differences between these two groups are due to the higher overlap *people* have with familiarity than with similarity [10].

Many of the participants were excited about getting recommendations for interesting people they do not know. One participant commented “*VERY neat tool! It did a very*

good job of finding people I would be interested in,” and another described the recommender as “*Kind of a business relationship 'match.com'.*” Quite a few participants endorsed the stranger recommendation scenario. One wrote “*Would be great to be able to connect to others who may be doing similar work about whom I might not be aware,*” and another commented “*This experiment is interesting, because I'm sure that in [our organization] there are people with similar roles (and pains) that see similar customer expectation.*”

Several participants mentioned they would like to get more suggestions to further explore stranger discovery opportunities. Others expressed their desire to be exposed to new roles. One participant stated “[...] *my biggest interest would be connecting to people most related to new roles where I don't already have gravity*” and another even said “*Got the feeling that this was also helping [to] see what kind of other roles would be of interest to me.*”

Some recommendations were described as “weird” or “strange” and some participants articulated they would

expect a higher level of accuracy from the recommender, e.g., “only a few people that would be of interest at this time.” Also, not everyone was enthusiastic about the stranger recommendation scenario, and several mentioned they expected to see people they know. One participant wrote “Many recommendations are way off. I’m used to get better recommendations.” Another participant wrote “At the end of the day, I’m paid for working, not for browsing potential interesting people without having any need to contact them.”

General Rating

To examine whether most of the recommendations were indeed strangers, we inspected the distribution of answers to Q1 (familiarity) over all 516 participants for *StrangerRS* recommendations and the 2 benchmarks: *StrongFam* and *Random*, as depicted in Figure 2. For *StrangerRS*, 67.3% of the recommended individuals were completely unknown to the participants, as indicated by a rating of 1 to Q1 (denoted as Q1=1). About 12% were “weak ties” rated 2-3, and 21% were “strong ties” rated 4 or 5. Overall, *StrangerRS* succeeded in filtering familiar people for over two-thirds of the recommendations. For comparison, in the *StrongFam* benchmark, 78.9% of the recommended people were rated with Q1=5, an additional 10.8% were rated with Q1=4, and only 4.6% were rated with Q1=1. On the other hand, for *Random*, 97.7% of the recommended individuals were strangers with Q1=1. Differences among the ratings of *StrangerRS*, *StrongFam*, and *Random* are all significant¹.




Figure 2. Distribution of Q1 rating results over all participants for *StrangerRS* and the two benchmarks

Rating distributions of Q1 across the three groups are shown in Figure 3. The percentage of strangers was lowest for the *S+p-F* group at 61%, indicating that including common people in the similarity relationships decreases the likelihood of recommending a stranger. The strong tie percentage was also highest for this group, with over 27% rated with 4 or 5. On the other hand, when mutual people were included as part of the familiarity relationships, in the *S-F+p* group, the percentage of strangers was over 72%, while strong ties were only 17%. For the *S-F* group the percentages were in-between: 68.6% strangers (Q1=1), while 18.7% rated 4 or 5. Differences among the groups are statistically significant¹.




Figure 3. Distribution of Q1 rating results for *StrangerRS* across the 3 groups

Figure 4 shows the interest in recommended people by depicting the general rating distribution for Q2-Q6. For Q2 (general interest), 28.1% of the recommended people were rated as completely uninteresting (a score of 1). While this is a substantial percentage, it is much lower than the percentage of unfamiliar people (67%). On the other hand, 26.4% of the recommendations were rated 4 or 5. When inspecting the Q2 rating distribution by group, *S+p-F* had the lowest percentage of uninteresting recommendations (26.1%), while, surprisingly, *S-F+p* had a lower percentage (28%) of uninteresting persons than *S-F* (30.1%). In terms of recommendations rated 4 or 5, *S+p-F* had 28.4% such persons, *S-F+p* had 24.9%, and *S-F* had 24.1%. Differences are significant¹, apart from the difference between *S-F* and *S-F+p*. These results point to the trade-off between recommending a stranger and recommending an interesting person, but they also indicate that common people should be included in one side of the subtraction or the other, as the *S-F+p* group yields more strangers and also slightly more interesting individuals than *S-F*.




Figure 4. Distribution of Q2-Q6 rating results for *StrangerRS* over all participants

The ratings of Q3-Q6 were lower than those of Q2, as shown in Figure 4. It seems that boiling down interest to a specific scenario drops the interest rating, possibly since some people do not use the tools Q3-Q6 refer to and those questions reflect commitment such as following a person's activity. One participant mentioned “Usually I would like to follow people with whom I’ve already established some level of a relationship [...]. Following too many people dilutes value” and another commented on a recommendation “Interesting enough to have a look in the files or bookmarks, but not to follow regularly.” Indeed, the ratings for Q6 (browsing one's files, bookmarks, or blogs), which does not represent a long-lasting action such as follow or tag, were the highest of Q3-Q6, with 32.1% of the recommendations being uninteresting and 26.2% of the recommendations rated 4 or 5. The lowest

¹ All statistical significance tests are performed using one-way ANOVA with Games-Howell post-hoc comparisons.

interest percentages were for the tweet-following question (Q3), probably as many participants do not use microblogging. Following one's activity in LC (Q4) and tagging (Q5) get very similar interest ratings, in-between tweet-following and browsing. Some comments explained positive ratings to one or more of Q3-Q6, e.g., "There were people that were shown to me who I would follow, though I would not want to add to my network."

Interest Rating for Strangers

We next focus on analyzing the recommendations of strangers, i.e., those recommendations for which Q1=1 (67% of the *StrangerRS* recommendations, as mentioned before). Figure 5 shows the distribution of Q2 ratings when Q1=1 (denoted Q2|Q1=1), over all participants, for *StrangerRS* and *Random*. For *StrangerRS*, 40.1% of the recommendations were rated uninteresting (Q2=1). On the other hand, over 59% of the strangers aroused some level of interest (Q2 rated 2-5), and 11.9% yielded high interest (Q2 rated 4 or 5). These ratings are significantly¹ higher than ratings for strangers recommended by the *Random* benchmark, where 69.8% of the recommendations were rated uninteresting, and only 2.7% were rated 4 or 5. While the *Random* benchmark recommended almost only strangers, these recommendations were rarely interesting.




Figure 5. Distribution of Q2 rating results given Q1=1 for *StrangerRS* and the *Random* benchmark

Looking at the entire set of recommendations produced by *StrangerRS*, 40.3% of the recommended people were rated Q1=1, Q2≥2; 21.4% were rated Q1=1, Q2≥3; and 8% were rated Q1=1, Q2≥4. These numbers are substantially higher than for the *Random* benchmark (29.5%, 10.6%, and 2.6%, respectively), even though the percentage of strangers recommended by *Random* was significantly higher. In summary, these results show that *StrangerRS* suggests people that are likely to be strangers, yet have a decent chance of arousing interest.

Some of the comments reveal more specific reasons for interest in strangers within the organization:

- "Following this person might help me better understand the sales environment in his part of the world"
- "She works for a marquee customer in the Telecom sector I cover. Any lessons or best practices she shares I would be very much interested in"
- "I know of and have worked with a few people in RTN Netherlands, but not this person. Would be good to watch and see what he is involved in"
- "Seems to be the expert on SNA in his geo, and I wasn't

even aware of him by now"

- "Works with implementations of products I work with. Other key contacts are known by this person. Useful tags. Looks useful"
- "I should follow her because she is highly experienced in my field and I need to learn from her"
- "This is the right on the spot [...] I know for a few months that I should get in touch with this group and she is apparently their manager!"

We also received comments that reveal particular reasons for **not** being interested in a recommended stranger. The gap between a business role and a technical role was highlighted by a few participants, e.g., "Seems a little too technically-focused for me to follow/connect with regularly" or "The only reason I didn't rate him 'very much' is that he is very technical and I more highly value business relevant information about our solutions." Language barriers were also mentioned by several participants, as one noted, "I can see that he uses Japanese characters, so it is likely that I won't be able to read his content." Others stated they are "not interested in those doing my same role in a different country." Or as another participant commented: "Unfortunately I am not in the habit of comparing notes with my counterparts in other geo." These kinds of comments typically came from sales people, indicating that for business people it might be less important to be aware of or connect with individuals who carry similar roles in other locations.

We next inspected Q2|Q1=1 by group in order to compare the interest rate each produced. For *S-F*, 42.3% of the unknown people were uninteresting, while 9.9% of them were rated 4 or 5. For *S+p-F* the percentages were 40.3% and 14.1%, respectively, while for *S-F+p* they were 37.6% and 11.9%. These results indicate, again, that *S-F* was the least effective group, while the other two groups were comparable; one had a higher percentage of strangers rated 4-5 and the other had a lower percentage of strangers rated 1. Given that *S-F+p* yields significantly more strangers, it might be the preferred configuration.

We also compared the rating results of recommended people ranked 1-10, 11-20, and 21-30 in the original list of recommendations, to examine whether quality of recommendation changes along the list. The percentage of strangers increased from 60% for people ranked in the top 10, to 68% for people ranked 11-20, to 73.2% for people ranked 21-30. These differences are all significant¹. Rating of Q2|Q1=1 slightly decreased from the top tenth to the second tenth and then remained stable (and even slightly increased) at the third tenth. These differences are all insignificant¹. Very similar trends were observed for Q3-Q6. These results indicate that the effectiveness of the recommendations remains steady along the top 30; the interest rating remained quite stable and the likelihood of recommending a stranger even increased. This may imply that the potential pool of effective recommendations can be large when recommending strangers.

Participant-Level Analysis

To examine the success of stranger recommendations per participant, we inspected the percentage of participants who received at least one “good” recommendation from *StrangerRS*, i.e., a recommendation of an interesting stranger for whom $Q1=1$ and $Q2 \geq q2$, where $q2 \in \{2,3,4,5\}$. The $q2$ threshold allows us to examine different levels of interest. Table 1 depicts these percentages for all 516 participants. For example, 85.3% had at least 1 recommendation rated $Q1=1, Q2 \geq 2$ and 67.6% had at least 1 with $Q1=1, Q2 \geq 3$. Table 1 also shows these results for at least 2 and 3 recommendations. For example, 47.3% had at least 2 recommendations rated $Q1=1, Q2 \geq 3$ and 61.6% had at least 3 recommendations rated $Q1=1, Q2 \geq 2$. These results indicate that out of 9 *StrangerRS* recommendations, a user is likely to receive a few strangers in whom s/he has some level of interest. While this is not the accuracy observed for recommendation of familiar people [12], we believe these are encouraging results for this exploratory scenario.

	$Q1=1, Q2 \geq 2$	$Q1=1, Q2 \geq 3$	$Q1=1, Q2 \geq 4$	$Q1=1, Q2 \geq 5$
≥ 1 rec	85.27%	67.64%	36.82%	10.27%
≥ 2 recs	73.06%	47.29%	17.64%	1.94%
≥ 3 recs	61.63%	31.01%	9.3%	0.39%

Table 1. Percentages of participants for whom at least 1, 2, or 3 *StrangerRS* recommendations rated $Q1=1, Q2 \geq \{2,3,4,5\}$

Some comments demonstrate that participants appreciated discovery of even just one or two interesting strangers. For example, one participant responded “*Great idea! Certainly there are one or two people that I will follow up with*” and another wrote “*Found a new contact in China - that made this experiment worth my time.*”

We also compared heavier and lighter users of LC. We defined “avid users” as users who are directly related to at least 100 other people, 100 tags, and 100 documents. Out of the 516 participants, 256 were avid users by this definition; we refer to the other 260 as regular users. The percentage of recommended strangers was lower for avid users at 64.2%, compared to 70.3% for regular users. However, the interest rating in strangers was higher for avid users, as shown in Figure 6. Differences were found to be significant¹. This trend is also consistent across all four interest scenario questions (Q3-Q6). Moreover, for 89.4% of the avid users (vs. 81.1% of the regular users), at least one recommendation was rated $Q1=1, Q2 \geq 2$. For 73.8% (vs. 61.5%), at least one was rated $Q1=1, Q2 \geq 3$, and for 41.1% (vs. 32.3%), at least one was rated $Q1=1, Q2 \geq 4$. These results demonstrate that the capability of producing interesting strangers increases for users who are more active in enterprise social media.

This point was also reflected in participants' comments. One wrote: “*I think your system is great and will become more valuable as more people generate social metadata.*”

Some participants thought they need to be more active on social media in order to get better recommendations. One wrote: “*I probably need to join more communities to be related to better 'matches'*” and another stated “*Seemingly my current, infrequent tags do not suffice to identify many colleagues with real, shared interests. I'll tag more going forward.*” These comments also imply that stranger recommendation can serve as an incentive for increasing engagement in enterprise social media [8].




Figure 6. Distribution of Q2 rating results given $Q1=1$ for *StrangerRS* recommendations for avid and regular users

DISCUSSION AND FUTURE WORK

The challenge of stranger recommendation is twofold: on the one hand recommend people that the user does not already know and on the other hand recommend individuals in whom the user is interested. Compared to the benchmarks, *StrangerRS* performs reasonably well. First, it recommends mostly strangers (over 67%), as opposed to a typical recommendation of a familiar person. Second, while not a strong benchmark, we show that even though *Random* recommends more strangers (almost 98%), *StrangerRS* ultimately recommends more **interesting** strangers. Overall, for over 67% of the participants at least 1 recommendation out of 9 yielded a stranger who was rated 3 or above for general interest, and for almost 37% there was at least 1 stranger recommendation rated 4 or above. We find these results encouraging for the stranger recommendation scenario.

Comments from participants highlight various scenarios in which a recommended stranger can be of interest – someone with a similar technical expertise in another location or unit; someone who is well connected to other individuals with whom you are interested to get in touch; someone who has high centrality or influence within the organizational network (“*the main thing that is attractive for me is that she is a VP*”); or someone who creates relevant content in blogs, files, wikis, e.g., “*Definitely would like to understand her activities [...] because she is ready to externalize content (basically publish stuff within the intranet).*” Boosting the recommendation of individuals with influential roles or social positions, or those who are particularly active in social media, can contribute to enhancing stranger recommendations. Future research should examine how this can be accomplished, while maintaining recommendation diversity.

Our experimentation indicated that it is worth including the *people* relationships on either side of the subtraction equation – either as part of the familiarity or the similarity

relationships. Including *people* as part of familiarity yields the highest percentage of interesting strangers. Including them on the similarity side yields fewer strangers and more weak ties. While these differences are not substantial, they can help fine-tune the recommender according to the specific requirements.

Our work relates to the concept of *homophily* [18] – the tendency to associate and bond with similar others. McPherson et al. [22] discuss homophily in social networks, arguing that people's tendency to connect with others who are similar to them leads to very homogenous networks. One could claim that our stranger recommender encourages homophily and connection between employees who are similar. We believe that from an organization's perspective such connections can be highly valuable, as demonstrated in some of the comments quoted in the previous section. Yet, it is also desirable to maximize recommended people's diversity in terms of location, unit, or job title, as well as in terms of the evidence that yields the recommendation. One of the participants wrote *“Many of the recommendations came from the same communities and tags, which gradually made them less compelling.”*

Some participants mentioned they would like to use such a tool regularly. One wrote *“I'd almost opt in to having something like this appear once a quarter to get me re-evaluate my network”* and another added *“I think it's good to do an experiment like this every 2 weeks, just to be aware of other people who have interests in common with me.”* As opposed to the case of friend recommendations, which exhausts quite rapidly [12], it seems that stranger recommendations can retain good quality for longer, since the potential number of similar people within large organizations is higher [10]. Our results show that the quality of recommendations does not significantly change along the top 30, but this needs to be validated for larger numbers of recommendations and also along time.

In this work, similarity is computed based on common activity in social media and does not take into account demographic attributes. Several participants suggested that profile attributes, such as location, division, job title or description be taken into account. It can be interesting to examine whether inclusion of profile-based similarity in our recommender can further improve the results.

One limitation of our evaluation is that it was conducted through a one-time user survey. Real deployment of such a feature can help to further evaluate its usefulness, also taking into account the time factor, as users' interests change and develop, and people join or leave the organization, or change positions over time. It would also be interesting to examine whether recommendations lead to stronger relationships over the long run.

Using stranger recommendations in combination with or as a substitute for friend recommendations can also be productive. For example, friend recommendations can be

suggested to new social media users who are building their initial network. Once established, stranger recommendations can help extend social circles and expand reach. Another option is to mix both friend and stranger recommendations in parallel, integrating both the higher accuracy of friend recommendations [12] and the serendipity, or “surprise effect”, of stranger recommendations. Further research needs to examine in detail how to interleave both types of recommendations.

In the era of information overload and social spam, some people may argue that they are already connected well enough and do not have the time or will to follow or even become aware of more individuals. A few participants highlighted this point. One of them explained *“It's really a matter of time and not interest, if that makes sense.”*

A potential enhancement may suggest “matchmakers” who can introduce a stranger to the user. The familiarity network can be used to suggest middlemen who are connected both to the end user and to the recommended stranger, in case such exist. Proposing a potential matchmaker can make a stranger recommendation more attractive, as suggested by previous studies [31,32].

We believe this work may inspire a similar feature outside the firewall, for social media users in general or for members of large communities in particular. The core engine can work analogously, suggesting similar users based on common social media activity and filtering out familiar people. Obviously, more exploration must be done regarding the effectiveness and potential impact of such a feature in web environments, as well as other challenges such as larger scales and multiple identities.

CONCLUSION

We suggested a novel system that recommends strangers within the organization, based on subtraction of the user's familiarity network from the user's similarity network. Both networks are mined from activity within enterprise social media. The results show that our recommender is capable of suggesting people who are strangers, but may be interesting nonetheless. Most users received at least one recommendation of an interesting stranger out of nine attempts. Many participants appreciated the recommendations and explained their value through different examples, such as learning from the experience of other individuals in their field; becoming aware of others with similar expertise, projects, or roles, in another location or division; or locating a new “bridge” to a department or community they do not have contact with.

This paper adds to the literature on people recommendation in social media, suggesting a more serendipitous scenario that would yield less expected recommendations. While these recommendations may be of lower accuracy, their contribution is in expanding the user's social capital, and ultimately reputation and influence within the organization. Future work should

examine how to further enhance the accuracy of recommendations, towards deployment of a stranger recommendation feature as part of enterprise social media or on the web. The longer-term effects of stranger recommendations should also be thoroughly studied.

REFERENCES

1. Boyd, D.M. and Ellison, N.B. Social Network Sites: Definition, History, and Scholarship. *JCMC*, (2007), 13, 1.
2. Budzik, J., Bradshaw, S., Fu, X., and Hammond, K. J. 2002. Clustering for Opportunistic Communication. *Proc. WWW '02*, 726-735.
3. Chen, J., Geyer, W., Dugan, C., Muller, M., and Guy, I. 2009. Make New Friends, but Keep the Old: Recommending People on Social Networking Sites. *Proc. CHI '09*, 201-210.
4. Constant, D., Sproull, L., and Kiesler, S. 1996. The kindness of strangers: the usefulness of electronic weak ties for technical advice. *Organization Science* 7 (2), 119-135.
5. DiMicco, J., Millen, D. R., Geyer, W., Dugan, C., Brownholtz, B., and Muller, M. 2008. Motivations for Social Networking at Work. *Proc. CSCW '08*, 711-720.
6. Ellison, N., Heino, R. and Gibbs, J. 2006. Managing Impressions Online: Self-P Processes in the Online Dating Environment. *JCMC*, 11 (2), 415-441.
7. Foner, L. N. and Crabtree, I. B. 1997. Multi-Agent Matchmaking. *Software Agents and Soft Computing: Towards Enhancing Machine Intelligence, Concepts and Applications*. Springer-Verlag, London, 100-115.
8. Freyne, J., Jacovi, M., Guy, I., and Geyer, W. 2009. Increasing Engagement through Early Recommender Intervention. *Proc. RecSys '09*, 85-92.
9. Granovetter, M. The Strength of Weak Ties. 1973. *American Journal of Sociology* 78(6), 1360-1380.
10. Guy, I., Jacovi, M., Perer, A., Ronen, I., and Uziel, E. 2010. Same Places, Same Things, Same People? Mining User Similarity on Social Media. *Proc. CSCW '10*, 41-50.
11. Guy, I., Jacovi, M., Shahar, E., Meshulam, N., Soroka, V., and Farrell, S. 2008. Harvesting with SONAR: the value of aggregating social network information. *Proc. CHI '08*, 1017-1026.
12. Guy, I., Ronen, I., and Wilcox, E. 2009. Do You Know?: Recommending People to Invite into Your Social Network. *Proc. IUI '09*, 77-86.
13. Guy, I., Zwerdling, N., Carmel, D., Ronen, I., Uziel, E., Yogev, S., and Ofek-Koifman S. 2009. Personalized Recommendation of Social Software Items based on Social Relations. *Proc. RecSys '09*, 53-60.
14. Hancock, J. T., Toma, C., and Ellison, N. 2007. The Truth about Lying in Online Dating Profiles. *Proc. CHI '07*, 449-452.
15. IBM Social Software for Business – Lotus Connections: <http://www.ibm.com/software/lotus/products/connections/>.
16. Klenk, S., Dippon, J., Fritz, P., and Heidemann, G. 2010. Determining Patient Similarity in Medical Social Networks. *MedEx Workshop, WWW '10*.
17. Lampe, C., Ellison, N., and Steinfield, C. 2006. A Face(book) in the Crowd: Social Searching vs. Social Browsing. *Proc. CSCW '06*, 167-170.
18. Lazarsfeld, P. F. and Merton, R. K. Friendship as a social process: A substantive and methodological analysis. *Freedom and Control in Modern Society* (1954), 18-66.
19. McCarthy, J. F., McDonald, D. W., Soroczak, S., Nguyen, D. H., and Rashid, A. M. 2004. Augmenting the Social Space of an Academic Conference. *Proc. CSCW '04*, 39-48.
20. McDonald, D.W and Ackerman M.S. 1998. Just talk to me: a field study of expertise location. *Proc. CSCW '98*, 315-324.
21. McNee, S. M., Riedl, J., and Konstan, J. A. 2006. Being accurate is not enough: how accuracy metrics have hurt recommender systems. *Proc. CHI '06 EA*, 1097-1101.
22. McPherson, M., Lovin, L. S. and Cook, J. M. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415-444, 2001.
23. Official LinkedIn Blog: <http://blog.linkedin.com/2008/04/11/learn-more-about-2>.
24. Official Facebook Blog: <http://blog.facebook.com/blog.php?post=15610312130>.
25. Putnam, R. D. 2000. Bowling Alone: the Collapse and Revival of American Community.
26. Quercia, D. and Capra, L. 2009. FriendSensing: Recommending Friends using Mobile Phones. *Proc. RecSys '09*, 273-276.
27. Rao, S., Hurlbutt, T., Nass, C., and JanakiRam, N. 2009. My Dating Site Thinks I'm a Loser: Effects of Personal Photos and Presentation Intervals on Perceptions of Recommender Systems. *Proc. CHI '09*, 221-224.
28. Shami, N.S., Ehrlich, K., Gay, G., Hancock T.J. 2009. Making sense of strangers' expertise from signals in digital artifacts. *Proc. CHI '09*, 69-78.
29. Steinfield, C., DiMicco, J. M., Ellison, N. B., and Lampe, C. 2009. Bowling Online: Social Networking and Social Capital within the Organization. *Proc. C&T '09*, 245-254.
30. Sumi, Y. and Mase, K. 2000. Supporting Awareness of Shared Interests and Experiences in Community. *SIGGROUP Bull.* 21, 3 (Dec. 2000), 35-42.
31. Terry, M., Mynatt, E. D., Ryall, K., and Leigh, D. 2002. Social Net: Using Patterns of Physical Proximity over Time to Infer Shared Interests. *Proc. CHI '02 EA*, 816-817.
32. Terveen, L. and McDonald, D. W. 2005. Social Matching: A Framework and Research Agenda. *ACM Trans. Comput.-Hum. Interact.* 12, 3 (Sep. 2005), 401-434.
33. Wellman, B. 2001. Computer networks as social networks. *Science*, 293 (5537), 2031-2034.