Contact Recommendations from Aggegrated On-Line Activity

Abigail Gertner, Justin Richer, and Thomas Bartee

The MITRE Corporation
202 Burlington Road, Bedford, MA 01730
{gertner,jricher,tbartee}@mitre.org

Abstract. We describe a system for recommending people to contact based on similar interests and activities as part of a company-wide social networking site. Our contact recommendation service aggregates input from multiple on-line data sources and combines them using a Bayesian noisy-MAX function to generate a rating of the overall match between two users. The system is running as part of an experimental social networking site at MITRE. We present the results of a preliminary evaluation in which we compare the recommendations to existing friend relationships.

1 Introduction

The use of social networking platforms in Enterprise settings, including industry and government, is growing rapidly. One of the primary stated uses for these tools is to help workers connect with each other within their organizations. This is particularly attractive in the case of large corporations and government agencies whose organizational structure as well as geographic separation can make it very difficult to know who you should be talking to.

Modeled on the popular internet social networking sites, such as Facebook and LinkedIn, Enterprise social networking tools generally include contact lists, allowing users to connect with the people they know. These tools often also include a suite of social utilities such as blogs, wikis, bookmarks, tags and file sharing. By connecting with people via the contact list, a user can then keep track of their contacts' activities on the site, providing a *social filter* on the available information.

Social networking sites become more useful and more attractive the more connections their users make with each other. Contact recommendation is a feature that is intended to make it easier for users of a social network to create their online network. Facebook and LinkedIn both provide suggestions of people to connect to, primarily based on existing network connections in order to help users find people they know to connect to on the site.

In this paper we present a contact recommendation tool that is designed to help workers in a large distributed enterprise environment make connections with others who share similar interests or work activities. We generate contact recommendations by aggregating information about users from diverse data sources within the company. The contact recommendations are implemented as a standalone web-service which is designed to be integrated with a social networking user interface. We show how the recommendations appear in our company's internal social networking tool and then discuss two preliminary evaluations of the recommendations.

2 Related Work

The literature on recommender systems has primarily focused on recommending items to users. Much of this work is based on *collaborative filtering* [1] – a technique that clusters people according to their item preferences and then recommends items that other similar users have liked.

There are several research projects that have looked at recommending people to each other. ReferralWeb [2] had the goal of finding existing chains of relationships between people by mining on-line documents such as co-authored papers and organizational charts. The *Do You Know?* system from IBM [4] also attempts to find people who are already known to the user in order to suggest that they be added to their social network. *Do You Know?* is implemented using SONAR [5], a social network aggregation tool that is probably the most similar to our own contact recommendation tool, in that it brings together evidence from multiple data sources to form its recommendations. The primary differences are the way we do the aggregation, and the fact that the IBM tool is attempting to identify existing social relationships, while we are primarily concerned with recommending people who are *not* known but possibly should be.

Terveen and McDonald [3] coined the term "social matching" to refer to systems that try to recommend and connect people each other. They outline the problem space in a series of claims, such as the need for explicit user models and the application to on-line social networks.

Another related area of research is the field of expertise finding [6]. This typically involves keyword searches for an expert who can answer a question or help solve a problem. In contrast, our contact recommender is looking for people who may be appropriate to form a longer term connection with based on common interests, not necessarily based on their expertise on a single topic.

3 Contact Recommendation Implementation

Our contact recommendation service is implemented as part of MITREverse, an experimental social networking system that is deployed inside the MITRE firewall. MITREverse is built on the Elgg open source social networking platform [7], which includes the basic social network features of friend lists, activity streams and message boards, as well as providing additional social tools such as groups, blogs, bookmarks and file sharing. MITREverse was deployed as part of a research project two years ago and has about 400 members in spite of never having been officially publicized or advertised within the company. Figure 1 shows a screenshot of a profile page on MITREverse.

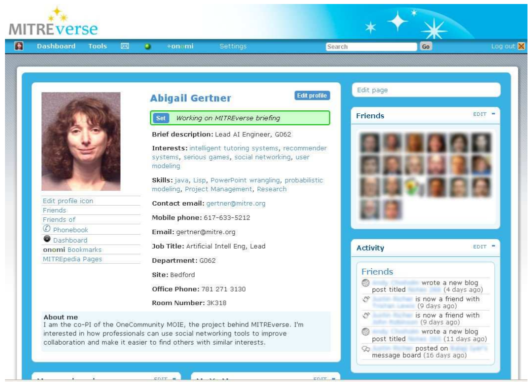

Fig. 1. A profile page on MITREverse

The original insight behind the MITREverse contact recommendation service was that by bringing together information about people from multiple places on the network, we could form a more accurate picture of what their interests are and what activities they are engaging in. As with many large companies, over the past five years or so MITRE has been making a number of social media tools available to its employees on our intranet. These include blogs, wikis, email lists, microblogging and social bookmarking tools. Some of these tools are official corporate supported offerings, and others are grassroots efforts started by individual employees. Many of these tools have built-in APIs that allow other programs to easily access and re-use their data. By aggregating the data from these different services together, the contact recommender creates a multi-dimensional view of what users have in common with each other.

Currently MITREverse uses seven data sources to compute the similarity scores: use of the same tags in onomi, our social bookmarking site, shared bookmarks in onomi, shared membership in internal email lists, co-editing of pages on our corporate-wide wiki, membership in groups on MITREverse, friend of a friend relationships, and use of the same tags on MITREverse. All of the data sources we use currently are inside the MITRE firewall and are accessed via public (to all MITRE users) APIs. We chose to avoid any privacy concerns by basing our recommendations only on data that would be accessible to anyone looking at the recommendations.

The contact recommender works by first generating a similarity score between each pair of users for each data source being considered. Since each data source may have a different type of user data, the data sources may use different algorithms for computing the similarity score. For instance, in the case of social bookmarking tags and MITREverse tags we use the cosine similarity of tag frequency vectors to compare two users' collections of tags. In the case of data sources in which there is a simple binary association between users and items, such as mailing list memberships, we use the Jaccard similarity coefficient (the size of the intersection divided by the size of the union) as the measure of similarity between two users. If there is not enough information about a user for one of the data sources, no scores will be computed for that data source for user pairs involving that user.

After the similarity scores are generated for the individual data sources, an overall score is generated for the match between each pair of users based on the combination of those scores. The overall match is represented as a rating from zero to five, which is displayed on the user interface as a set of zero to five stars, as shown in Figure 2. The icons beneath the stars represent the data sources that were used to generate the recommendation, and clicking on the recommendation will take the user to a detailed explanation page for the recommendation.

Fig. 2. The display of a single recommendation

There are several possibile approaches to combining the individual data source scores into an aggregated rating. The most straightforward solution is to use the average score of the data sources, possibly weighted according to which data sources are considered more important. This is the approach taken by the *Do You Know?* system [4]. However, there are often cases where there is a strong match between two users on one or two of the data sources and a weak match on the rest. Using an average over all data sources would cause the overall score in these cases to be low, whereas we believed that people with a strong match in even one area would be likely to benefit from knowing each other. Therefore we decided to use a model that works more like an OR relationship – if *any* of the data sources scores is high, the resulting aggregated score will be high.

The model we are using is a causal probabalistic model called the Noisy-MAX [8], which is used in Bayesian networks to model a multi-valued variable whose value depends on the maximum value of its causal influences. Figure 3 shows a graphical representation of the Noisy-MAX model for aggregating similarity scores. The definition of the Noisy-MAX says that the inferred value of the node representing the outcome variable (in this case, the strength of the match in question, from zero to five) is determined by the maximum value produced independently by that node's inputs. The "noise" built into the probabilistic relationship means that the more inputs there are with a high similarity score, the more likely the overall similarity is to have a high value.

The Noisy-MAX takes advantage of the fact that the causal influences on the effect node are considered to be independent of each other. In this case, the causal influences are the individual data sources and the effect is the aggregated rating. Using this independence assumption, it is possible to define the relationship between the causes and the effect with much fewer parameters than

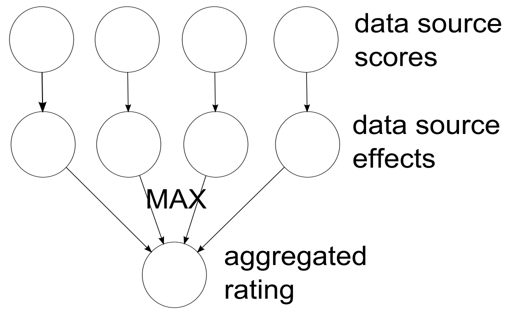

Fig. 3. The Noisy-MAX model for aggregating scores

a full conditional probability table. For each link from a data source to the combined rating, the parameters needed specify the probability that the combined rating will be equal to each possible value (0-5 stars), given the value of the data source and assuming that all other data sources are absent. Currently the parameters for the noisy-MAX are estimated subjectively but we are working to derive values from actual user judgments of matches between users.

We expect that it is the case that some data sources are more influential than others in determining a good match between users. Furthermore, each user may prioritize the various data sources differently. Therefore it is important to be able to weight the inputs in order to adjust their effect on the aggregated rating, as was done in [4] with the weighted average of the input scores. We have modified our implementation of the Noisy-MAX to include a weight parameter for each input data source, so that the influence of the individual data sources on the aggregated score can be adjusted according to the corresponding weights. Since different users may have different priorities for the data sources, we plan to allow them to adjust these weights via the user interface, although this feature is not yet implemented.

The contact recommender runs nightly to update its recommendations. With 396 users on MITREverse, the update takes about twelve minutes. However much of this time is spent downloading the full user data from the external (outside of MITREverse) data sources and so that time will not increase as MITREverse gains additional users.[1] The noisy-MAX computation to combine the data source scores currently takes about 1 minute and forty seconds. Extrapolating that value to apply it to all possible pairs of MITRE's approximately 7000 employees, the update would take 8.35 hours, so it will still be possible to update the recommendations once a day.

[1] Several of the web services we connect to include API calls that allow us to retrieve all user data in a single call. We do this even though many of the users are not current MITREverse users because we include non-MITREverse users in our recommendations in order to encourage current users to invite their recommended contacts onto the site.

4 Evaluation

As a preliminary evaluation of the accuracy of the contact recommendation ratings we have looked at the correspondence between the ratings and the actual friend relationships that currently exist in MITREverse. We hypothesized that for pairs of users who are connected to each other in the social network, the recommendation rating should be higher than for pairs of people who are not connected. This is not a perfect measure because, first, there may be people who know each other who have not yet connected on the site and, second, people may have friends on the site who they don't have much in common with. However, both of these disadvantages actually make it less likely that a difference would be detected. If we can see a difference in the ratings between the friend pairs and the non-friend pairs, it would give us an initial confirmation that our recommendations are doing the right thing.

There are 396 total user accounts on the MITREverse site, making 156,420 possible friend relationships (friend relationships in MITREverse are unidirectional, so a relationship of person A to person B is treated separately from a relationship of person B to person A). Out of these possible friend relationships there are 1,752 actual friend relationships in the site. The contact recommender found enough information in at least one data source to generate recommendations for 29,609 of the possible friend relationships. 1316 of these recommendations were for pairs of users who are already connected as friends, and the remaining 28293 are for pairs who are not yet connected. Table 1 summarizes the number of recommendations generated for existing friend and non-friend pairs.

Table 1. Total number of recommendations generated for friend and non-friend pairs

	Friends	Not Friends
Recommendation	1316	28293
No Recommendation	436	126374

We then compared the recommendations that were generated for the friend pairs with the recommendations that were generated for the non-friend pairs. The mean recommendation rating (taken from the 1-5 star ratings) for the friend relationships is 1.94, while the mean for the non-friends is 1.38. While this is a small difference, it is encouraging given our caveats about the friend relationships not being a completely reliable correlate of match strength. Figure 4 shows the percentage of the friend and non-friend pairs that were given each of the five possible ratings by the contact recommender. The non-friend pairs are 1.7 times as likely to be given a rating of one star as compared to the friend pairs. Conversely, the friend pairs are 5.15 times as likely to get a four star rating and 3.73 times as likely to get a five star rating.

To evaluate the use of the noisy-MAX function for combining the data source inputs, we generated recommendations using an alternative method of simply

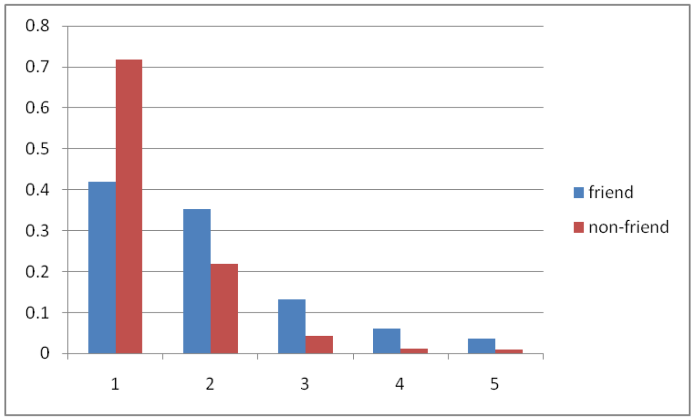

Fig. 4. Percentage of friend and non-friend user pairs assigned each rating, from one to five stars, using the noisy-MAX function

taking the average of the input scores, which is the approach taken in SONAR [5]. Figure 5 shows the percentage of pairs associated with one to five star ratings using this method. In this case, since we are not using the max function, both the friend and non-friend pairs are more likely to be given a rating of one star, and there are very few four or five star ratings, even for the friend relationships. Based on these results, it appears that the Noisy-MAX model does a better job of assigning a high rating to people who are actually friends than the straight average does.

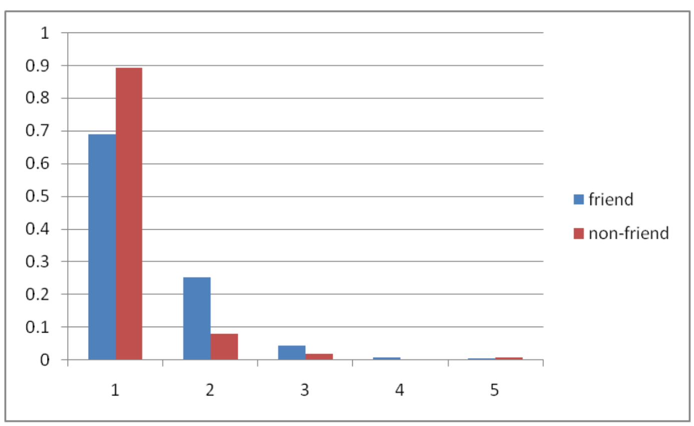

Fig. 5. Percentage of friend and non-friend user pairs assigned each rating, from one to five stars, using the average of data source inputs

4.1 Human Evaluation

An additional evaluation of the contact recommendations is underway, in which we are asking people to judge the similarity of pairs of people based on the same types of inputs that the recommender uses, to get a better understanding of what algorithms people use to make those decisions. This will help us to evaluate the use of the Noisy-MAX model as well as to determine the optimal parameter values. Initial results indicate that many people do consider the maximum of the individual data source inputs to be important in determining the final rating.

In this study we are also asking subjects to weight the individual data sources according to their importance to the overall recommendation. Eight out of the ten subjects so far have said that they consider the differences between data sources to be important or very important when computing the overall rating, but their ordered rankings of the data sources are all very different from each other. This supports our belief that it will be important to allow users to adjust the data source weights themselves to deliver optimal recommendations.

5 Discussion and Future Work

We have described a contact recommendation tool that looks at data available about users based on their on-line activities and uses that information to generate recommendations for other people with similar interests. This capability will soon be deployed company wide on a new social networking site called Handshake.

Our initial evaluation of the contact recommendations based on existing friend relationships shows that the ratings are at least able to distinguish between people who are connected to each other in the social network and those who aren't. It also suggests that the noisy-MAX model provides an advantage over taking the average of the inputs because it leads to higher ratings in general, especially for the existing friend pairs. However, as we noted earlier, we are primarily interested in the ability of the recommender to identify relevant contacts that are *not* already known to the user. Therefore we are planning a followup evaluation in which we will ask users to judge the actual recommendations that the system generates for them personally. In that case, we will be able to see whether the system is able to recommend novel connections that users would be likely to follow up on.

References

1. Herlocker, J. L., Konstan, J. A., Borchers, A., Riedl, J.: An Algorithmic Framework for Performing Collaborative Filtering. In: Proceedings of Research and Development in Information Retrieval, ACM, New York (1999)
2. Kautz, H., Selman, B., Shah, M.: Referral Web: combining social networks and collaborative filtering, Communications of the ACM 40(3), pp 63–65 (1997)

3. Terveen, L. and McDonald, D. W.: Social Matching: A Framework and Research Agenda. ACM Transactions on Computer-Human Interaction, 12(3), pp. 401–434, (2007)
4. Guy, I., Ronen, I., Wilcox, E.: Do You Know? Recommending People to Invite into Your Social Network, In: Proceedings of the 13th international conference on Intelligent user interfaces, pp. 77–86. ACM, New York (2009)
5. Guy, I., Jacovi, M., Shahar, E., Meshulam, N., Soroka, V., Farrell, S.: Harvesting with SONAR: the value of aggregating social network information. In: Proceedings of the twenty-sixth annual SIGCHI conference on human factors in computing systems (CHI '08), pp. 1017–1026. ACM, New York (2008)
6. Maybury, M., D'Amore, R., House, D.: Automating Expert Finding. International Journal of Technology Research Management. 43(6): 12-15 (2000)
7. Elgg open source social networking platform, `http://www.elgg.org/`
8. Díez, F., Druzdzel, M.: Canonical probabilistic models for knowledge engineering, Technical Report CISIAD-06-01. UNED, Madrid, (2006)