Who Should I Follow? Recommending People in Directed Social Networks*

Michael J. Brzozowski
HP Labs, Social Computing Lab
Palo Alto, CA, USA
zozo@google.com

Daniel M. Romero
Cornell University, Center for Applied Mathematics
Ithaca, NY, USA
dmr239@cornell.edu

Abstract

A variety of social networks feature a directed attention or "follower" network. In this paper, we compare several methods of recommending new people for users to follow. We analyzed structural patterns in a directed social network to evaluate the likelihood that they will predict a future connection, and use these observations to inform an intervention experiment where we offer users of this network new people to connect to. This paper compares a variety of features for recommending users and presents design implications for social networking services. Certain types of structural closures significantly outperform recommendations based on traditional collaborative filtering, behavioral, and similarity features. We find that sharing an audience with someone is a surprisingly compelling reason to follow them, and that similarity is much less persuasive. We also find evidence that organic network growth is very different from how users behave when they are prompted to connect to new people.

Introduction

A wide variety of Web services feature publicly articulated social networks. Often the value users derive from these services depends on the quality and diversity of users' networks. While *friendship* is a socially loaded concept often used to establish a sense of community on personal social network sites (boyd 2006), so-called *weak ties* also provide valuable resources and information to a user (Granovetter 1983), suggesting that users may benefit from connecting to new people and expanding their networks. Indeed, users of enterprise social networks are particularly motivated to cultivate a network of weak ties and to seek out new people (DiMicco et al. 2008).

While traditional social network sites represent friendship as reciprocated links in an undirected graph (*i.e.* if you're my friend, I'm also your friend), services such as Twitter are popularizing a directed graph: a *follower network* or *attention network*. This allows users to *follow* people of interest without requiring them to reciprocate, thus lowering the cost of expanding one's network. Furthermore, undirected social networks naturally allow for some users to be followed by many people without following many themselves, effectively becoming "celebrities" or "stars". Usually these types of relationships are based on information interest by the follower rather than purely social interactions.

*The full version of this paper is at http://goo.gl/v3zZn.

Social network services are acutely aware that users need friends to enjoy a social network, and so services like Facebook often recommend users to people. Typically this is done by looking for people who are connected to many of a user's friends, reasoning that the user is likely to already know them. This intuition works well in an undirected social network. However, in a directed social network setting where both social and information relationships are found, the question of who to recommend a user to follow becomes more complicated.

Furthermore, many social network services have access to much more information than the network structure: users supply details about their interests and background, and they provide clues about who interests them by choosing whose posts to read and respond to. This led us to ask what cues in a user's profile, behavior, and network might be most effective in recommending people.

Related Work

Golder and Yardi evaluated structural patterns on Twitter with an intervention experiment that asked whether users would consider following specific users selected by a black box algorithm. Users responded via a Likert scale indicating their potential interest. Their analysis found that structural paths involving reciprocated links were generally a strong signal in whether those recommendations were accepted (2010).

We seek to extend work by Golder and Yardi in an enterprise social network, where the motivations and costs of usage vary significantly from Twitter (Ehrlich and Shami 2010). To better understand the factors behind users' decision to follow someone, we chose to measure whether users actually chose to follow the people we presented, forcing them to evaluate the costs of admitting them to their attention stream. We also wondered whether the types of connections formed as this network grows "organically" would be different from the types users might form when prompted to.

Environment

To evaluate different types of recommendations, we performed an intervention experiment on an enterprise social network in use at HP, WaterCooler (Brzozowski 2009). WaterCooler allows users to build profiles using tags to indicate their personal and professional interests, skills, customers they support, and projects and teams they work on. It also has a directed follower network similar to that of Twitter. Unlike Twitter, though, all people are publicly visible to all

HP employees, and do not need to grant permission for users to follow them.

Recommendation Types

We set out to compare the effectiveness of different types of features in recommending new people for users to follow. We compare three classes of features that have been used to recommend people in social networks (Chen et al. 2009; Golder and Yardi 2010; Guy et al. 2010). We call these *behavioral*, *network*, and *similarity*. Throughout this paper, we'll call the user we're making recommendations for A, and a potential user he or she may want to follow B.

Behavioral

This is perhaps the most intuitive of reasons to follow someone: if you already seem to pay attention to someone, why not follow them? We define two recommendation features:

MostRead. The people who've written the most posts that A clicked on. WaterCooler has recorded 60,566 clicks on posts by 4,397 authenticated users.

MostReplied. The people whom A has written the most replies to. WaterCooler has 182,866 replies to posts, written by 22,567 users.

Network

The network itself provides clues to who a user might want to connect to, as his or her peers have already done some filtering. In contrast to an undirected social network (boyd 2004), in an attention network, users don't feel socially obligated to reciprocate all links (Golder and Yardi 2010), making a directed link a stronger statement of interest.

Collaborative Filtering Collaborative filtering is a common means of recommending content to users based on other users with apparently similar tastes (Goldberg et al. 1992). Often this manifests intuitively to users in the form "People who liked X also liked B", as commonly seen on e-commerce sites. Formally, there is at least one user Y who shares an interest in following X with A. So we can then suggest that A may also be interested in another user that Y follows, namely B.

Structural Closures In a social network, suppose two members A and B each know a third person X. Granovetter suggests that in most of these structures, a *triadic closure* occurs: A and B are likely to know each other, or are more likely to meet one another the more they associate with X (Granovetter 1973). Intuitively this happens in real-life social networks: the more mutual friends you have in common with someone, the more likely you are to know them. Therefore, in an online social network, if A and B have many friends in common, they may be very likely to become friends as well.

However, if they have many friends in common and do *not* friend each other, they may be making a conscious decision to stay disconnected. This pattern is called a *forbidden triad* (Granovetter 1983). This brings up an interesting dilemma: triadic closure says they they are very likely to eventually meet or to know each other already, but on the other hand, there may be a good reason this triad will not close.

In a directed network, we have more information because the direction of the links tell us something about the relationship between A and X, and between B and X. A could be following X, X could be following A, or A and X could be following each other (we call them *colleagues*).

We analyzed 282,429 triads on the WaterCooler network prior to our experiment to determine the empirical probability that triads of each type would eventually close. To do this, we looked for triads where each of the requisite links existed (for instance, $A \to X$ and $X \to B$), and counted the number of times that the closing link $A \to B$ was added *afterward*. A triad that satisfies, *e.g.*, $S9$, also satisfies each of the other structures, so we decided to count each triad only once, under the most restrictive structure that it satisfies.

Overall, less than 3% of the triads closed, highlighting the sheer scale of possible triads a user A might be connected to (on average, each user plays the role of A in 41 of these triads, despite following only 7 people). But the structure of the triad has a significant impact.

Structure		Samples	Closed
S9	$A \leftrightarrow X \leftrightarrow B$	8,114	10.6%
S3	$A \to X \leftrightarrow B$	8,296	9.1%
S1⋆	$A \to X \to B$	15,331	6.6%
S7⋆	$A \leftrightarrow X \to B$	23,513	6.4%
S8	$A \leftrightarrow X \leftarrow B$	8,507	5.1%
S2	$A \to X \leftarrow B$	26,810	4.3%
S6	$A \leftarrow X \leftrightarrow B$	24,706	2.2%
S5	$A \leftarrow X \leftarrow B$	14,735	1.2%
S4	$A \leftarrow X \to B$	151,417	0.5%
⋆ S1 and S7 are virtually tied.			
S7,8,9	$A \leftrightarrow X \star B$	40,134	7.0%
S1,2,3	$A \to X \star B$	50,437	5.8%
S4,5,6	$A \leftarrow X \star B$	191,858	0.8%
S3,6,9	$A \star X \leftrightarrow B$	41,116	5.3%
S2,5,8	$A \star X \leftarrow B$	51,052	3.5%
S1,4,7	$A \star X \to B$	190,261	1.7%
At least one reciprocated link		73,136	5.6%
Unreciprocated links only		209,293	1.5%

Table 1: Empirical ranking of closure types. Unless otherwise noted, within a section, types listed first had a higher closure rate than those listed later, with 95% confidence.

Table 1 shows the empirical closure probabilities for each of the nine structural types, ordered by how likely they were to close. [1] Echoing similar findings (Golder and Yardi 2010), it appears that reciprocated links are far more likely to imply closure than un-reciprocated links; these reciprocal ties between two users indicate mutual interest in one another, which may strengthen that tie. Indeed, the most likely structure to close is $S9$, the "colleague of my colleague", followed by $S3$, the "colleague of someone I follow".

In general, users are more likely to trust the "taste" of people they're following: the top four structures all fulfill at least $A \to X \to B$. By contrast, users seem relatively unlikely to care about the people who are connected to their followers ($S4, 5, 6$). The worst-performing triad is the "mutual follower" $S4$. This may reflect the diverse reasons people have for following each other; X may have completely orthogonal interests in A and B, making B less likely to be of interest to A.

[1] We make no claims about how normative these closure probabilities are for similar social networks, only the relative ranking between them in WaterCooler.

We also evaluated each structure by the number of intermediate X connectors there are between A and B, to see if having more mutual contacts makes a closure more likely; we call this parameter k. We calculate the empirical probabilities of closure, along with 95% confidence intervals. The margins of error for higher values of k reflect the relative scarcity of triads with, for instance, ten mutual connectors. For most structures, the probability of closure increases with higher values of k—at first. However, in some cases we observe a decrease in closure probability as k increases. One possible explanation for this decrease is that it is driven by the error; indeed, the margin of error increases dramatically as k becomes large. Where significant, though, this decrease could also be explained by the forbidden triad mechanism. The fact that A and B have so many X's in common could mean A already knows about B but has decided B is not interesting to follow.

Similarity

In real life, people tend to befriend others with similar sociodemographic characteristics and interests (McPherson, Smith-Lovin, and Cook 2001), a phenomenon called *homophily*. It stands to reason, then, that users might be interested in attending to people with similar interests. As a proxy for this, we use user-defined "person tags", which list users' interests, hobbies, and skills. Shared tags also reflect a variety of offline ties: working on the same team, supporting the same product or customer, or attending the same event or school. All of these provide a rich source of information about the similarity of two users' interests and experiences.

After experimenting with several metrics, we chose Dice's coefficient (Dice 1945) to compare the similarity of two users A and B, which is defined as:

$$\frac{2\,|T_A \cap T_B|}{|T_A| + |T_B|}$$

(van Rijsbergen 1979), where T_A and T_B are the sets of tags on A and B's profiles, respectively. We found this metric seemed to be resilient to "tag spam attacks" by some WaterCooler users to associate themselves with every possible tag because potential user Bs with very large $|T_B|$ will be penalized as matches for A.

Experiment

To evaluate these different types of recommendations, we created a tool that recommended people for WaterCooler users to follow. Recommendations were presented to users in five sections,with the ordering of the sections randomized. Users could also hover over the people recommended to them for a sample of their three most recent posts, allowing them to judge the frequency and relevance of their postings.

For each section we selected the ten highest-ranked users (Bs) according to the recommendation criteria, and provided an explanation for each. As a result, users could receive up to 50 recommendations at a time. If there wasn't enough information to recommend people in a section, a message informed the user of this and provided concrete steps (*e.g.*, adding tags to their profile, clicking through to read others' posts) that would enable that section, inviting them to return to the tool for updated recommendations.

For behavioral recommendations *MostRead* and *MostReplied*, the ranking and explanation are simply based on the number of posts read or replied to.

For *Collaborative* recommendations, we rank candidate Bs by how many $\langle X, Y \rangle$ pairs there are connecting A and B. The explanation given is a list of two or three Xs whom A follows, ordered by how many Ys follow both X and B.

For *Structural* recommendations, we calculated, for each candidate B, how many closures existed of each of the nine types (that is, how many distinct Xs could form that closure); we call this parameter k. While we have empirical closure probabilities, we have relatively low confidence in many scenarios, particularly with large k. So to introduce some randomness, we score each structure by randomly (with uniform distribution) selecting a value t within the 95% confidence interval for that structure and value of k. We then use the sum of all the $t_1 ... t_9$ values to rank each candidate B. To explain the recommendation, we take the structure s that contributed the highest t_s value. Since the score of each structure is chosen randomly inside the confidence interval, every time a user uses the tool he or she is recommended different people to follow. This property gives users a slightly different experience every time they use the tool, potentially making it more likely that they will eventually find people they are interested in following.

For *Similar* recommendations, we ordered Bs by their Dice coefficient, and explained the recommendations using the most rare tags that A and B have in common.

We recorded the timestamps of each person "offered" to users and the details surrounding each recommendation, and tracked which recommendations led to users being followed. We considered a recommendation given to user A to follow user B *accepted* if A followed B when the recommendation was given, or if A first visited B's profile and then decided to follow him/her. A user B might be recommended in multiple sections; in this case, we count the section where the user clicked a "Follow" link to be the offer that was ultimately accepted. We avoid double-counting of the same recommendations offered over multiple sessions by only counting each $\langle A, B, Method \rangle$ triad once and counting it as accepted if A eventually chose to follow B using the tool.

We invited users to try out our experimental "Build Your Network" tool by posting announcements on a variety of HP internal media, by linking to it from several places in WaterCooler, and by making it part of the flow for new WaterCooler users. Users who had recommendations in most (at least three) of the sections were invited to take a short follow-up survey, where we asked them questions about how they reacted to the tool. Respondents who completed the survey were eligible to win a USD 10 Amazon.com gift card.

Results

We collected data over a 24-day period in July 2010. During this time, 227 users tried the tool; 45% of them used it at least twice, perhaps suggesting that our enticement to "unlock" enough sections to win a gift card was effective. However, 19% of our users returned on a subsequent day, a surprising level of engagement. In all, 110 users followed 774 new people with our tool. While this is a minority, most of the users who never accepted any offers received 20 or fewer recommendations out of the possible 50. Of the users who received 50 or more recommendations, 76% of them followed at least one new user.

Over the course of the month since the experiment opened, we've only observed three users unfollow someone after accepting a recommendation to follow them, and in each case it happened within 30 seconds of originally

(a)	Rate	2	258	369	58
S1,4,7:[a] $A \star X \rightarrow B$	15.1%	66%	90%	99%	99%
S2: "you both follow X"	11.6%		N/A	90%	80%
S2,5,8: $A \star X \leftarrow B$	9.8%			80%	N/A
S3,6,9:[b] $A \star X \leftrightarrow B$	6.0%				<50%
S5,8: "following X"	4.5%				

[a] "followed by X" [b] "colleague of X"

(b)	Rate	789	123
S4,5,6: $A \leftarrow X \star B$	14.2%	<50%	90%
S7,8,9: $A \leftrightarrow X \star B$	11.5%		<50%
S1,2,3: $A \rightarrow X \star B$	9.8%		

(c)	Rate	Beh	Sim
Network	11.6%	<50%	90%
Behavioral	11.2%		90%
Similarity	8.7%		

N/A · <50% · 50% · 66% · 80% · 90% · 99% confidence

Table 2: Confidence intervals by which recommendation classes on the left outperformed classes on the top.

following them; we attribute this to correcting unintended clicks rather than dissatisfaction with that user's posts.

Overall, about 11% of recommendations were accepted, ranging by category from 3 to 19%.

For each recommendation class R and confidence level α we can construct a binomial proportion confidence interval of the success of R using the number of times R was offered and accepted, calculated as in (Clopper and Pearson 1934). For each pair of recommendation classes R_1 and R_2 we find the highest confidence level α^* such that the binomial proportion confidence intervals of R_1 and R_2 are disjoint. If the confidence interval of R_1 is higher than that of R_2 we say that R_1 outperforms R_2 with α^* confidence. In keeping with statistical convention we report the significance of outperforming classes with $p = 1 - \alpha^*$.

Recommendations where $X \rightarrow B$ outperformed those where $X \leftarrow B$ ($p < .07$) and those where $X \leftrightarrow B$ ($p < .01$); see Table 2(a). When broken down by how recommendations were presented to users, "followed by X" significantly dominates "colleague of X" and "following X" ($p < .01$). Having a common interest in X may be a strong motivation to follow someone.

Table 2(b) compares the type of relationships with X that are most effective. It appears $A \leftarrow X$ is significantly more likely to be accepted than $A \rightarrow X$, perhaps indicating that users care more about networking with their audience's connections than those of the people they follow.

As shown in Table 2(c), we find that *Similar* recommendations are dominated by most others, suggesting that homophily alone does not produce good quality recommendations. This result agrees with the previous findings that similarity of interests are not always a good predictor of future behavior in online communities (Crandall et al. 2008) and that homophily does not have a significant effect in triadic closure (Kossinets and Watts 2006). In this case, having similar tags was not a good enough signal for predicting when a user decides to follow someone. It's also possible that tags are not highly predictive of what users regularly post about, as suggested by (Brzozowski 2009).

Conclusions

We compared a variety of different mechanisms for recommending users to one another on online social networks. It's clear that the directionality of links in such networks has a significant impact on the likelihood of organic closure and the success rate of people recommendations. For more details, see the full version of this paper at http://goo.gl/v3zZn.

Acknowledgments

We thank Gabor Szabo, Wojciech Galuba, Bernardo Huberman, and Jon Kleinberg for helpful conversations and our experiment participants for their feedback.

References

boyd, d. m. 2004. Friendster and publicly articulated social networking. In *CHI '04: Ext. abstr. on Human factors in computing systems*, 1279–1282. ACM.

boyd, d. 2006. Friends, friendsters, and top 8: Writing community into being on social network sites. *First Monday* 11(12).

Brzozowski, M. J. 2009. Watercooler: exploring an organization through enterprise social media. In *GROUP '09: Proc. of the int. conf. on supporting group work*, 219–228. ACM.

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. In *CHI '09: Proc. of the int. conf. on Human factors in comput. syst.*, 201–210. ACM.

Clopper, C., and Pearson, E. S. 1934. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika* 26:404–413.

Crandall, D.; Cosley, D.; Huttenlocher, D.; Kleinberg, J.; and Suri, S. 2008. Feedback effects between similarity and social influence in online communities. In *KDD '08: Proc. of the 14th int. conf. on Knowledge discovery and data mining*, 160–168. ACM.

Dice, L. R. 1945. Measures of the amount of ecologic association between species. *Ecology* 26:297–302.

DiMicco, J.; Millen, D. R.; Geyer, W.; Dugan, C.; Brownholtz, B.; and Muller, M. 2008. Motivations for social networking at work. In *CSCW '08: Proc. of the ACM conf. on Comp. supported coop. work*, 711–720.

Ehrlich, K., and Shami, S. 2010. Microblogging inside and outside the workplace. In *ICWSM '10: Proc. of the Int. Conf. on Weblogs and Social Media*. AAAI.

Goldberg, D.; Nichols, D.; Oki, B. M.; and Terry, D. 1992. Using collaborative filtering to weave an information tapestry. *Comm. of the ACM* 35(12):61–70.

Golder, S. A., and Yardi, S. 2010. Structural predictors of tie formation in twitter: Transitivity and mutuality. In *SocialCom '10: Proc. of the 2nd IEEE International Conference on Social Computing*. IEEE.

Granovetter, M. 1973. The strength of weak ties. *American Journal of Sociology* 78:1360–80.

Granovetter, M. 1983. The strength of weak ties: A network theory revisited. *Sociological Theory* 1:201–233.

Guy, I.; Jacovi, M.; Perer, A.; Ronen, I.; and Uziel, E. 2010. Same places, same things, same people?: mining user similarity on social media. In *CSCW '10: Proc. of the 2010 conf. on Computer supported cooperative work*, 41–50. ACM.

Kossinets, G., and Watts, D. J. 2006. Empirical analysis of an evolving social network. *Science* 311(5757):88–90.

McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27(1):415–444.

van Rijsbergen, C. J. 1979. *Information Retrieval*. London: Butterworths.