Who Should I Follow?
Recommending People in Directed Social Networks

Michael J. Brzozowski
Social Computing Lab
Hewlett-Packard Labs
Palo Alto, CA, USA
mike.brzozowski@hp.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.

Author Keywords
social networks, recommender systems, follower networks, attention

ACM Classification Keywords
H.5.3 Information Interfaces and Presentation: Group and Organization Interfaces—*Web-based interaction*; J.4 Social and Behavioral Sciences: Sociology

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 [1], so-called *weak ties* also provide valuable resources and information to a user [13], 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 [7]. Many also use them to build social capital within their organization [19].

On some social sites, the network plays a central role in propagating news [17]; having a diverse and well-connected network can improve the efficiency of this process and expose users to more diverse content.

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[1] 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.

Social network services are acutely aware that users need friends to enjoy a social network, and so services like Facebook[2] 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
Chen *et al*. compared similarity and network cues for recommending people in an enterprise social network at IBM,

[1] twitter.com
[2] facebook.com

Submitted for review to CSCW 2011.

and found that similarity was a stronger cue in recommending new contacts [4]. Guy *et al*. expand on this by proposing that user similarity might be derived from the "people, things, and places" that they share [14].

Golder and Yardi evaluated structural patterns on Twitter and found that structural paths involving reciprocated links were generally a strong signal in recommending users [11]. Our methodology is slightly different in that we chose to manipulate the *explanation* given for recommendations as well.

A variety of recent work goes in the other direction: determining the strength of social ties from empirical data exposed in an online social network. The mere presence of a network connection does not imply a strong relationship or interaction [15]. As a result, the frequency of interaction can be very predictive of the actual strength of a social network tie [9]. However, Wu *et al*. distinguish between the types of closeness that these interactions might imply in an enterprise social network [21].

Environment

To evaluate different types of recommendations, we performed an intervention experiment on an enterprise social network in use at HP, WaterCooler [3]. 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.

Prior to our experiment, the network had 6,844 members (nodes connected to other users) and 13,581 profiles, containing 52,407 person-tags.

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 [4, 11, 14]. 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

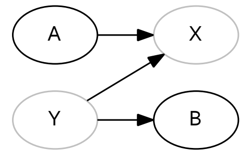

Figure 1. Collaborative filtering pattern.

The network itself provides clues as 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 [2], in an attention network, users don't feel socially obligated to reciprocate all links [11], 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 [10]. 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. This pattern is illustrated in Figure 1.

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 [12]. 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* [13]. This brings up a very interesting dilemma when trying to predict if a new friendship will occur between them in an online social network. We have two conflicting forces: 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* of each other). The 3×3 matrix of the nine possible types of directed closure are listed in Table 1.

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

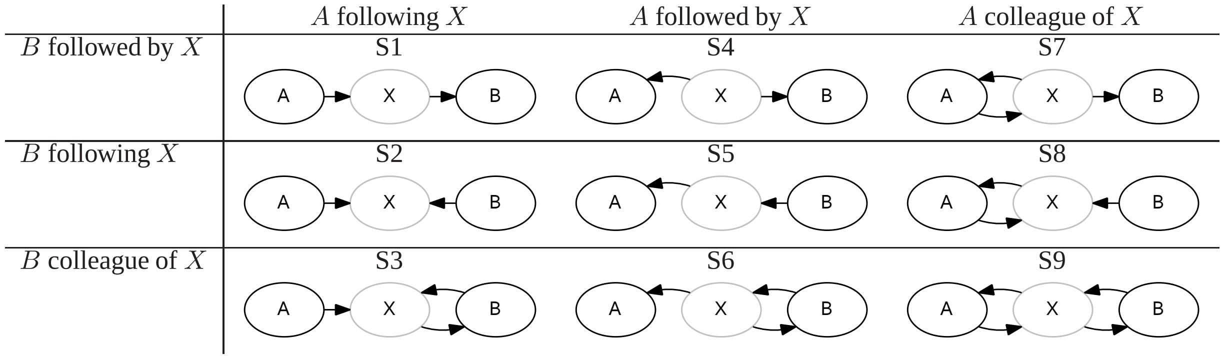

Table 1. *Structural* closure types.

Structure		Samples	Closed
S9	$A \leftrightarrow X \leftrightarrow B$	8,114	10.6%
S3	$A \rightarrow X \leftrightarrow B$	8,296	9.1%
S1*	$A \rightarrow X \rightarrow B$	15,331	6.6%
S7*	$A \leftrightarrow X \rightarrow B$	23,513	6.4%
S8	$A \leftrightarrow X \leftarrow B$	8,507	5.1%
S2	$A \rightarrow X \leftarrow B$	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	$A \leftarrow X \rightarrow B$	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	$A \rightarrow X \star B$	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	$A \star X \rightarrow B$	190,261	1.7%
At least one reciprocated link		73,136	5.6%
Unreciprocated links only		209,293	1.5%

Table 2. Empirical ranking of closure types. Unless otherwise noted, within a section, structures listed first had a higher closure probability than those listed later, with 95% confidence.

this, we looked for triads where each of the requisite links existed (for instance, $A \rightarrow X$ and $X \rightarrow B$), and counted the number of times that the closing link $A \rightarrow B$ 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.

Table 2 shows the empirical closure probabilities for each of the nine structural types, rank ordered by how likely they were to close.[3] Echoing similar findings in [11], it appears that reciprocated links are far more likely to imply a 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 classic "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 $A \rightarrow X \rightarrow B$. 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.

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. The empirical probabilities of closure, along with 95% confidence intervals, are shown in Figure 2. 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 [18], 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.

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

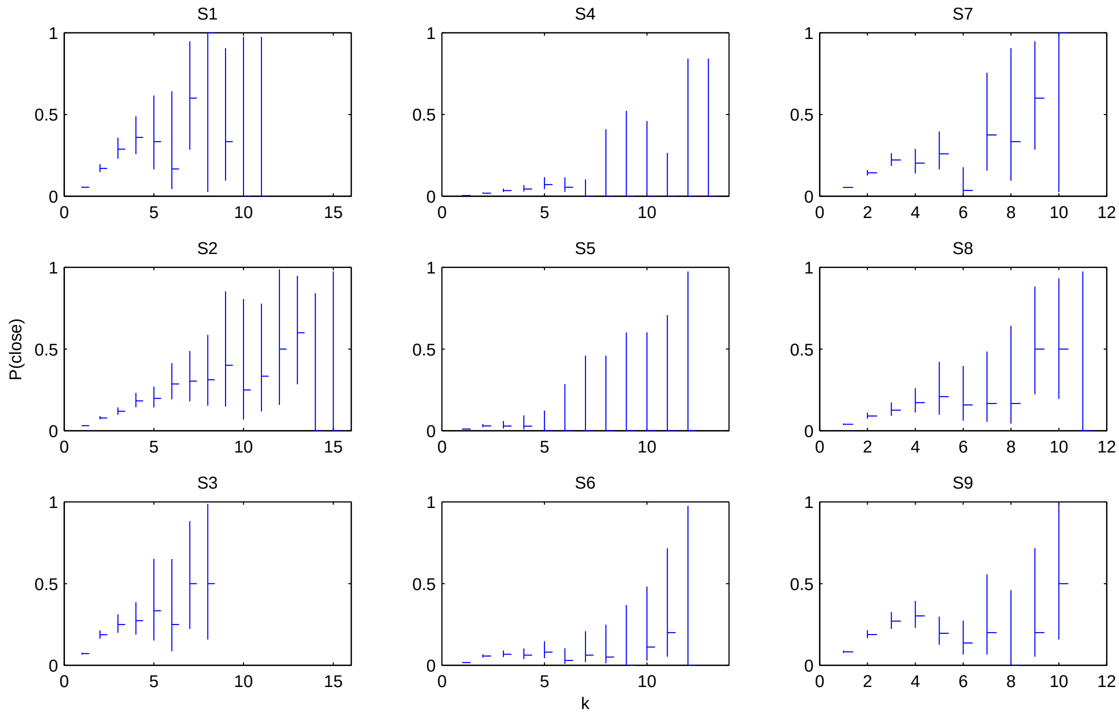

Figure 2. 95% confidence interval around probability of closure for each of the closure types, as a function of the k Xs participating.

Class	Name	Label
Behavioral	*MostRead*	People you read the most
	MostReplied	People you reply to the most
Network	*Collaborative*	People following your contacts also follow
	Structural	Your network neighborhood
Similarity	*Similar*	People with similar tags

Table 3. People recommendation categories, as seen by users.

After experimenting with several metrics, we chose Dice's coefficient [6] to compare the similarity of two users A and B, which is defined as:

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

[20], 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, labeled as shown in Table 3, with the ordering of the sections randomized. Figure 3 shows the user interface, which provides a brief explanation of the evidence for each recommendation. For example, Wayne is recommended "because" he follows Alex and Arumugam, two people the user might recognize. We hypothesize that this introspectability may make the recommendations more useful, since users can see why they were provided. 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 (see Figure 1 for their roles). The explanation given is a list of two or three Xs whom A follows, ordered by how many Ys follow both X and B.

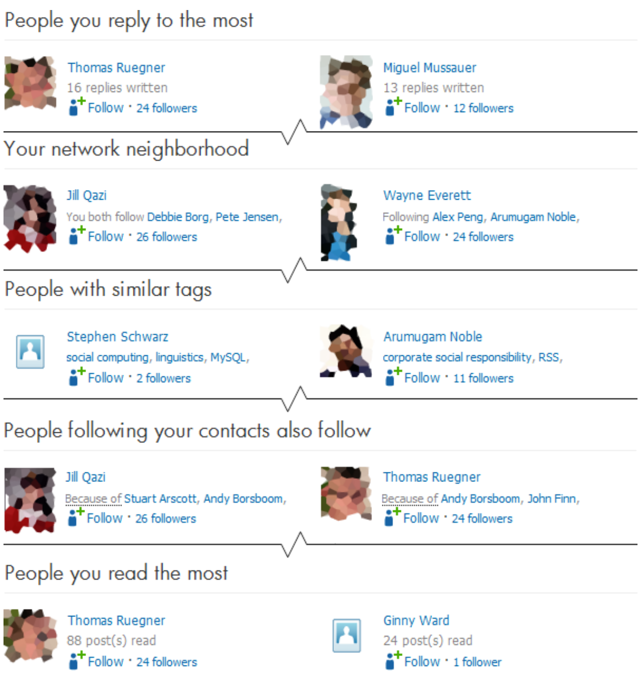

Figure 3. The user interface for the experiment. Users are given up to 50 recommended people to follow. (Screenshot is truncated for space.)

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 shown in Figure 2, 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. The explanations used come from the leftmost column of Table 1.[4] Since the score of each structure is chosen randomly inside the confidence interval, every time a user utilizes 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 following them; we attribute this to correcting unintended clicks rather than dissatisfaction with that user's posts.

In this section we discuss the results obtained from the experiment. First we show that the percentage of or accepted recommendations is independent on the number of followers of the users before the experiment, we discuss which types of recommendations were more successful and possible explanations, and finally we show the results obtained from the surveys taken by the users.

Users' Out-Degree

An interesting question that arises is who is more likely to find "Build Your Network" useful: users who already follow many people or users who only follow a few? One hypothesis is that users who already follow many others are more likely to find the tool useful because they are more active and engaged with the network and therefore are more likely to be open to extending their network even more. Furthermore, the recommendations made to these users may be of better quality since we have more information about them. On the other hand, one could argue that the tool is not as useful for those who follow many people because they have already found the ones they are interested in, whereas users who are following only a few people are more likely to find

[4] S2 ($A \rightarrow X \leftarrow B$) is described as "You both follow X".

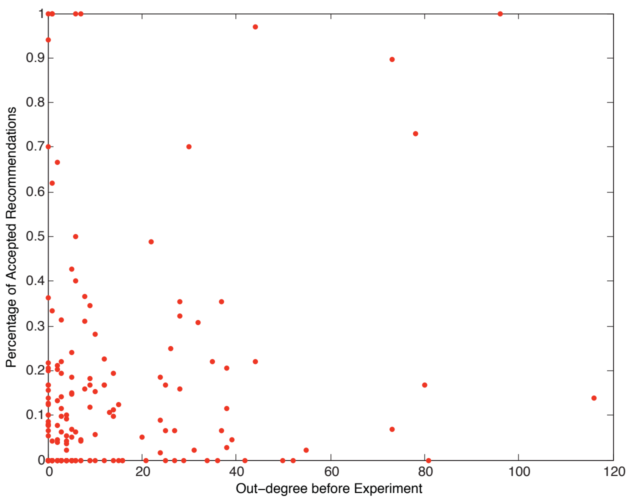

Figure 4. Out-degree of users before the experiments vs. the fraction of offers they accepted. There is no clear increasing or decreasing trend even when partitioning the horizontal axis into bins.

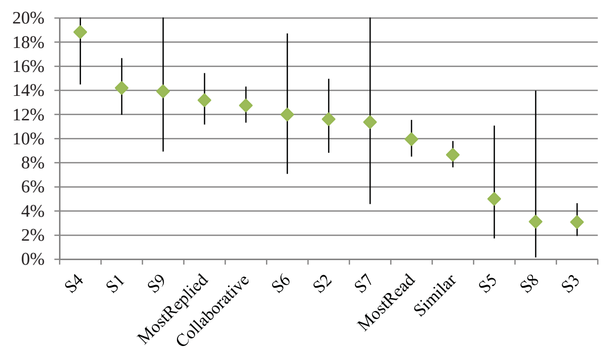

Figure 5. Acceptance rates for offers of each type, with 95% confidence intervals.

other encounter users they would like to follow but did not know about.

In order to answer this question we looked at the user's out-degree (the number of people they were following) before using the tool and compared it with the number of recommendations they accepted. Figure 4 shows that there is no evidence that users with many followers are more likely to accept a larger percentage of the offers, or vice-versa. Overall, we do not find a definite domination of the accepted recommendation by users following very many or very few people. We should note that even users who weren't following anyone could still get recommendations from $S4$, $S5$, and $S6$, if they already have followers in the network, and isolated users in the network could still get *Similarity* recommendations if they have tags on their profile.

Which Recommendations Were Accepted

Figure 5 shows the acceptance rate of each type of recommendation as well as the 95% confidence intervals; rates

(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 4. Confidence intervals by which recommendation classes on the left outperformed classes on the top.

range from 3-19%. Overall, about 11% of recommendations were accepted.

Comparing this with the empirical closure probabilities predicted by Table 2 shows some stark differences. $S4$ ($A \leftarrow X \rightarrow B$) vaunts from dead last in empirical closure likelihood to being the most-accepted recommendation type. This suggests that these "shared audience" links may be compelling, but it may be relatively difficult for A to discover B "naturally". This result also contradicts Golder and Yardi's finding that $S1$ and $S4$ had no significant effect on subjects' preferences [11], but since their experiment didn't explain the recommendations it provided, we suspect that revealing an implicit "endorsement" by X may be persuasive. Indeed, Table 4(a) shows that recommendations where $X \rightarrow B$ outperformed those where $X \leftarrow B$ by 50% ($p < .07$) or those where $X \leftrightarrow B$ ($p < .01$). When broken down by how recommendations were presented to users, "followed by X" dominates "you both follow X", which dominates "colleague of X" and "following X". Having a common interest in X may also be a strong motivation to follow someone.

Meanwhile, $S3$ ($A \rightarrow X \leftrightarrow B$) plummets to the bottom three, perhaps suggesting that, while A is likely to discover X's colleagues in the course of building their network, A might not be all that interested in "new" colleagues of the people they're already following. However, "colleagues of colleagues" $S9$ ($A \leftrightarrow X \leftrightarrow B$) are significantly more likely than $S3$ recommendations to be accepted ($p < .01$).

Table 4(b) compares the type of relationships with X that are most effective in recommendations. In contrast with the categories in Table 4(a), the distinctions between these classes are not made explicit to users (though they may recall their relationship with Xs when presented as an explanation). It appears $A \leftarrow X$ is significantly more likely to be accepted than $A \rightarrow X$, perhaps indicating that users care more about

	Rate	S1	S9	*Repl*	*C*	S6	S2	S7	*Read*	*Sim*	S5	S8	S3
S4: $A \leftarrow X \rightarrow B$	18.8%	66%	50%	80%	90%	66%	80%	50%	90%	99%	90%	90%	99%
S1: $A \rightarrow X \rightarrow B$	14.2%		<50%	<50%	<50%	<50%	50%	<50%	90%	99%	90%	80%	99%
S9: $A \leftrightarrow X \leftrightarrow B$	13.9%			<50%	<50%	<50%	<50%	<50%	66%	80%	80%	66%	99%
MostReplied	13.2%				<50%	<50%	<50%	<50%	80%	90%	90%	80%	99%
Collaborative	12.8%					<50%	<50%	<50%	80%	99%	90%	80%	99%
S6: $A \leftarrow X \leftrightarrow B$	12.0%						<50%	<50%	<50%	50%	66%	66%	90%
S2: $A \rightarrow X \leftarrow B$	11.6%							<50%	<50%	80%	80%	66%	99%
S7: $A \leftrightarrow X \rightarrow B$	11.4%								<50%	<50%	<50%	<50%	80%
MostRead	9.9%									50%	66%	66%	99%
Similar	8.7%										50%	50%	99%
S5: $A \leftarrow X \leftarrow B$	5.0%											<50%	<50%
S8: $A \leftrightarrow X \leftarrow B$	3.1%												<50%
S3: $A \rightarrow X \leftrightarrow B$	3.1%												

<50% 50% 66% 80% 90% 99% confidence

Table 5. Matrix showing the confidence intervals by which each recommendation type (on the left) performed better than each other type (on the top).

networking with their audience's connections than those of the people they follow.

In Table 5 we show the types of recommendations that dominate others as well as the statistical confidence. $S4$ emerges as a likely favorite. The recommendation types break up into three rough clusters: those that do significantly better than *MostRead*, those that tie *MostRead*, and those that do worse than *MostRead*. *MostReplied* outperforms *MostRead*, reflecting that the greater effort required to reply to someone is a stronger endorsement of their content than clicking through to read their posts. Traditional collaborative filtering performs relatively well, in a virtual tie with most of the top two tiers, and on par with similar $S1$ ($A \rightarrow X \rightarrow B$).

The broad categories of recommendation types are compared in Table 4(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 [5] and that homophily does not have a significant effect in triadic closure [16]. 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 [3].

User Reactions

Fifty-six users completed our follow-up survey, providing additional feedback on their experience with our recommendations. While the sample size is relatively small, we can still draw a few conclusions and observe some patterns.

We found that 79% of respondents reported the *MostReplied* recommendations were useful, significantly more than the 63% who found the *Structural* recommendations useful ($p < .32$); see Table 6. Curiously, though, users are more likely to perceive *MostRead* as useful than *Collaborative* ($p < .41$),

		Read	*Sim*	*C*	*Str*
MostReplied	79%	<50%	<50%	66%	66%
MostRead	77%		<50%	50%	50%
Similar	74%			<50%	50%
Collaborative	64%				50%
Structural	63%				

Table 6. User-reported usefulness of each type of recommendation ($N = 47$).

		Read	*C*	*Str*	*Sim*
MostReplied	63%	<50%	80%	90%	90%
MostRead	56%		66%	80%	90%
Collaborative	40%			<50%	50%
Structural	36%				<50%
Similar	25%				

Table 7. Portion of users who were likely to recognize most or all of the people recommended to them ($N = 46$).

even though they're more likely to actually accept *Collaborative* recommendations ($p < .2$). Perhaps these behavioral recommendations are "useful" in that they reflect users' past behavior, but less likely to be novel.

Indeed, Table 7 shows that most respondents reported being already familiar with "most or all" of the people recommended to them in the *MostReplied* and *MostRead* conditions (unsurprisingly, since of course, they had to reply to or read posts by these people to receive them as recommendations). They were significantly less likely to recognize the majority of their *Collaborative* recommendations ($p < .17$). *Similar* seemed the most likely to provide mostly novel recommendations to users, although they were seldom accepted.

A number of users reported surprise at being able to recall people whom they've engaged with or paid attention to in the

I follow people who...		2	3	4	5	6
1. post interesting things	85%	<50%	90%	90%	99%	99%
2. I work with	77%		66%	80%	99%	99%
3. I've met	57%			<50%	80%	90%
4. have similar interests	53%				66%	90%
5. follow me	34%					<50%
6. are high up in the org	26%					

<50% 66% 80% 90% 99% confidence

Table 8. Reasons users follow people, with confidence intervals that these are higher than other responses.

past and reconnect with them. For example, P48 reported:

> The 'replied to' section led me back to someone I found interesting way back when, but didn't follow...

Some described following behavioral recommendations out of convenience. P5 remarked, "I read his blog posts all the time so following him will make it easier to see what he is up to." Several others described behavioral recommendations as a check on who they might have forgotten:

> It was really useful to see who I was NOT following. I "talk" to [him] all the time but I had not realized I had never hit the "follow" button for him. —P15

Most users appreciated being able to preview recent posts for their recommendations; 71% reported this as "somewhat" or "very" useful. Still, about 57% visited people's profiles before deciding whether to follow them. 85% of users reported the recommendation explanations were helpful. Several cited the connecting users (Xs) as helpful in evaluating potential connections. "Nice to know what kinds of posts I might expect to read from the suggested person based on who follows them," said P26. In describing why she chose to follow one user B, P29 wrote: "[B] had a great pedigree (linked to [X])". However, the *Structural* recommendations caused some confusion; some of the structures were not as intuitive to users, and it was less clear how they were derived.

Table 8 lists respondents' reasons for following people on WaterCooler; users were far more likely to follow people who "post interesting things" or who work directly with them than because they're "high up in the organization" ($p < .01$). Likewise, users reported relatively little obligation to reciprocate by following people who follow them.

Many users said having "interesting" or "useful" posts was a primary criterion in deciding whether to follow someone, reflecting the value users place on allowing others into their attention stream. Most (83%) said they'd have no problem unfollowing people if they post "too much junk". A few people reported deliberately not following recommended users because they felt disconnected from existing social circles:

> I don't like to follow the folks from the [group] except in a few instances, because they seem to think it's instant messaging and I am not interested that the soup is cold today. —P29

Some recommended people seem to direct their postings at their close colleagues, informing their team of events like weddings, births, or "colleague hi-fives"—content of little value to people outside their own network. A few respondents avoided following users who they feel are too personal or social in their postings:

> Some folks are a bit "chatty" on Chatter or need to install the "personal filter". My favorite is the person who announced they are bored (for [many] employees to read). Second favorite example of filtering is some of [the] [location] folks who talk about non-work stuff fairly often to each other (lots of social niceties). Great that folks are social (World Cup chats are great examples), but we are at work... —P15

A handful of users zeroed in on people close to them in the organization: teammates, managers, and people in the same business unit or job role. Indeed, a few users suggested adding a section to recommend people from within the same group. Some ended up mainly following familiar people:

> Most of the people that I chose to follow (going from 25 to 42 after visiting this page) were names I recognized as those either active in IT or social media or communications and had positive associations with. —P6

Several users looked for people to follow whom they seemed to share things in common with beyond the org chart.

> I've discovered a guy that's posting constantly and actually found out he worked at the same site as me. His posts are interesting because they relate to my actual role within HP and he shared best practices with me. —P50

A few reported following users because of specific tags they had in common with people: being alumni of the same school, speaking the same language, or being part of the same internal community.

DISCUSSION

We compared a variety of different mechanisms for recommending users to one another on online social networks. We combined empirical analysis, an intervention experiment, and a followup survey to illuminate different facets of how social network services can help users to grow their networks. This area is drawing renewed interest, with LinkedIn, Facebook, and now Twitter using recommendations to stimulate network growth.

Limitations

We should note that this experiment was conducted on an enterprise social network, which has somewhat different expectations of privacy and motivations for use than external social media sites [3, 7]. Indeed, people use microblogging differently in the workplace than externally on Twitter [8].

However, users of any social network have limited attention and limited time to explore the periphery of their network, and are likely to perceive only a small number of people as worth following. Also, since everyone in the network is an employee of the enterprise we do not have to worry about the biases that certain accounts such as spammers or robot accounts present when studying open social networks.

Furthermore, we note that given the small size of they network, the sample size used for the experiment is necessarily small. However, even though many of the confidence intervals are large, we are still able to draw definite conclusions based on the clear trends we observed.

Future Work

One open question is whether different types of network recommendations might be more appropriate at different stages of a user's network growth, or for different user motivations. Should a user who explicitly wants to network with and meet new people branch out to people who are relatively disconnected from their network, or seek natural triadic closures?

Another area that could extend this work is to try to distinguish between people who are locally interesting (*i.e.*, to their peers) versus globally interesting (to a broader audience). A number of users noted that they rejected recommended people because they seemed to post mainly for the benefit of a smaller audience they didn't feel a part of (*e.g.*, a team or community). Users who are likely to be connected to these communities may well be interested in these "local" recommendations, while users who are more distant may prefer connections who have more "global" appeal.

We have analyzed the quality of the different types of recommendations by simply checking if they were accepted or not. However, many questions about the quality of the formed edges remain open. For example, it would be interesting to see if the edges created through the experiment are more or less likely to be deleted form the network later on than edges that were spontaneously created. In other words, will the users change their mind about following the people they chose during the experiment? Also, measuring the amount of interaction facilitated by the new edges could be another important measure of their quality. Another open question is how the edges created during the experiment compare to other edges in terms of the effect they have on the whole social network. Do the experiment edges improve graph theoretical measures of the network such as diameter or average clustering coefficient more or less than other edges?

Conclusions

We draw four main conclusions from this work.

Empirical behavior is not necessarily predictive.

A caveat of purely empirical analysis of this graph data is that the relative likelihood of structures to close might be more reflective of users' offline social networks than their online connections. Imagine a team that is destined to all follow each other and form a strongly-connected clique. Then the formation of those network links might purely be to bring the online network "in synch" with the real-world network. As seen in the relative probabilities of structural closure, empirical observation "in the wild" is an imperfect predictor of how users will react in an intervention. Site operators should take care when drawing inferences between past behavior and the results of interaction with a new tool.

Homophily isn't everything.

In contrast to [4], we find that similarity is a relatively poor reason to recommend that a user follow someone. Perhaps for a more robust definition of *similarity*, recommendations might be more compelling, but homophily by itself doesn't necessarily imply that someone's updates are typically of interest to a user.

Tell users something they don't know.

Another plausible explanation for the discrepancies between the empirical and experimental closure rates is that some connections, like $A \leftarrow X \rightarrow B$, are difficult for users to learn about organically, but potentially very interesting because they signal a sort of commonality between A and B. Conversely, closures like $A \rightarrow X \leftrightarrow B$ may be relatively easy for users to encounter in the offline world, so suggesting them to users may not provide additional value. In fact, these may be *forbidden triads*. Users may even be consciously avoiding those connections—well aware of them but thoroughly uninterested.

Users care about their audiences too.

Some of the strongest recommendations (*e.g.*, $A \leftarrow X \rightarrow B$) came via connections the user *hadn't* made—just by virtue of followers they had in common with other people. A user's audience is a reflection of how their contributions are perceived and who pays attention to them; perhaps sharing the same audience is the strongest type of homophily in an attention economy.

Design Implications

We offer these implications for designers implementing social network services.

Provide introspectability for system behavior.

The vast majority of our users appreciated the explanations our system provided for the recommendations they received—and in a few cases the explanations meant the difference between accepting and ignoring the suggestion. It's difficult for users to visualize the entire social graph, but short, simple descriptions of the connections they share with potential contacts proved compelling in explaining why a user might want to take a chance on following someone new.

Leverage directed links.

The nature of the link between two people says a lot about their relative status and interest in one another. By providing directed links, users encode this information in the network without the social baggage that comes from requesting a mutually approved friendship. As an added benefit, the directed network provides information even to users who don't follow many people, as users can benefit from knowing what sort of people are following them.

Provide preview.

Following a new user carries the additional overhead of being subscribed to his or her future updates. Most of our users didn't take the decision to follow a new person lightly—many scanned a sample of their recent posts, and some even visited their profiles. However, many current user recommendation interfaces ask users to blindly follow new people. Providing some sort of means of previewing the repercussions of this decision helps users make better-informed decisions.

Mix it up.

Despite the high accept rate of our most popular recommendation feature, 95% of the links formed during our experiment came from other, "inferior" features. As poorly as similarity-based recommendations performed, we still found users who appreciated them—and many even *perceived* them to be more useful than demonstrably better-performing features. All this suggests that perhaps user experience is best served by an ensemble method where users might be recommended for a variety of reasons. Further, we found a significant number of users returning to the "Build Your Network" tool repeatedly to get new recommendations. Injecting an element of randomness may help engage users to keep them coming back.

ACKNOWLEDGMENTS

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

REFERENCES

1. d. boyd. Friends, friendsters, and top 8: Writing community into being on social network sites. *First Monday*, 11(12), December 2006.
2. d. m. boyd. Friendster and publicly articulated social networking. In *CHI '04: Ext. abstr. on Human factors in computing systems*, pages 1279–1282. ACM, 2004.
3. M. J. Brzozowski. Watercooler: exploring an organization through enterprise social media. In *GROUP '09: Proc. of the 2009 int. conf. on Supporting group work*, pages 219–228. ACM, 2009.
4. J. Chen, W. Geyer, C. Dugan, M. Muller, and I. Guy. Make new friends, but keep the old: recommending people on social networking sites. In *CHI '09: Proc. of the 27th int. conf. on Human factors in computing systems*, pages 201–210. ACM, 2009.
5. D. Crandall, D. Cosley, D. Huttenlocher, J. Kleinberg, and S. Suri. 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*, pages 160–168. ACM, 2008.
6. L. R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26:297–302, 1945.
7. J. DiMicco, D. R. Millen, W. Geyer, C. Dugan, B. Brownholtz, and M. Muller. Motivations for social networking at work. In *CSCW '08: Proc. of the ACM conf. on Comp. supported coop. work*, pages 711–720.
8. K. Ehrlich and S. Shami. Microblogging inside and outside the workplace. In *ICWSM '10: Proc. of the Int. Conf. on Weblogs and Social Media*. AAAI.
9. E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *CHI '09: Proc. of the 27th int. conf. on Human factors in computing systems*, pages 211–220. ACM, 2009.
10. D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Comm. of the ACM*, 35(12):61–70, 1992.
11. S. A. Golder and S. Yardi. Structural predictors of tie formation in twitter: Transitivity and mutuality. In *SocialCom '10: Proc. of the 2nd IEEE International Conference on Social Computing*. IEEE, 2010.
12. M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–80, May 1973.
13. M. Granovetter. The strength of weak ties: A network theory revisited. *Sociological Theory*, 1:201–233, 1983.
14. I. Guy, M. Jacovi, A. Perer, I. Ronen, and E. Uziel. 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*, pages 41–50. ACM, 2010.
15. B. A. Huberman, D. M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1), January 2009.
16. G. Kossinets and D. J. Watts. Empirical analysis of an evolving social network. *Science*, 311(5757):88–90, January 2006.
17. K. Lerman. Social networks and social information filtering on digg. In *ICWSM '07: Proc. of the Int. Conf. on Weblogs and Social Media*.
18. M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
19. C. Steinfield, J. M. DiMicco, N. B. Ellison, and C. Lampe. Bowling online: social networking and social capital within the organization. In *C&T '09: Proc. of the 4th int. conf. on Communities and technologies*, pages 245–254. ACM, 2009.
20. C. J. van Rijsbergen. *Information Retrieval*. Butterworths, London, 1979.
21. A. Wu, J. M. DiMicco, and D. R. Millen. Detecting professional versus personal closeness using an enterprise social network site. In *CHI '10: Proc. of the 28th int. conf. on Human factors in computing systems*, pages 1955–1964. ACM, 2010.