

Predicting Reciprocity in Social Networks

Justin Cheng
Computer Science
Cornell University
Ithaca, NY 14853
jc882@cornell.edu

Daniel M. Romero
Center for Applied Mathematics
Cornell University
Ithaca, NY 14853
dmr239@cornell.edu

Brendan Meeder
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
bmeeder@cs.cmu.edu

Jon Kleinberg
Computer Science
Cornell University
Ithaca, NY 14853
kleinber@cs.cornell.edu

Abstract—In social media settings where users send messages to one another, the issue of *reciprocity* naturally arises: does the communication between two users take place only in one direction, or is it reciprocated? In this paper we study the problem of *reciprocity prediction*: given the characteristics of two users, we wish to determine whether the communication between them is reciprocated or not. We approach this problem using decision trees and regression models to determine good indicators of reciprocity. We extract a network based on directed @-messages sent between users on Twitter, and identify measures based on the attributes of nodes and their network neighborhoods that can be used to construct good predictors of reciprocity. Moreover, we find that reciprocity prediction forms interesting contrasts with earlier network prediction tasks, including link prediction, as well as the inference of strengths and signs of network links.

I. INTRODUCTION

As social media applications gain richer sets of features, they come to contain increasingly diverse connections among their users. The microblogging site Twitter is one example in which users share links to content, pass on messages, and initiate messages with an intended target. These directed messages, which we call @-messages, signal a communication link between two users and can represent many different forms of interaction. Indeed, earlier research has shown that Twitter contains a large amount of social activity between users who interact with each other as peers, as well as a large amount of information-seeking and information-sharing activity in which users interact with celebrities, news sources, and other types of high-visibility accounts [1], [2].

A challenge when studying an environment such as Twitter is that these types of connections are superimposed in a single communication network. Therefore, it is important to develop techniques capable of classifying the links in the underlying network according to the different activities that they represent. To this end, we formulate the problem of predicting *link reciprocity*, and develop a set of techniques for this purpose.

Reciprocity captures a basic way in which different forms of interaction on a site like Twitter take place. When two users v and w interact as peers, one expects that @-messages will be exchanged between them, passing in both directions — we consider this to be a *symmetric*, or *reciprocated*, interaction. On the other hand, if a user v sends multiple messages to a celebrity or news source w , it is likely that w will not send messages in return — this is an *asymmetric*, or *unreciprocated*

interaction. Which features characterize the difference between reciprocated and unreciprocated relationships? Can we tell them apart based on properties of the users involved, and properties of their network neighborhoods? Do the sub-networks of reciprocated and unreciprocated links have different structural properties? These are some of the questions we address in this paper; our approach also addresses the broader issue of isolating the different forms of interaction that take place on a complex social media site such as Twitter.

A. Summary of Results

We pursue two main approaches to analyzing reciprocity. The first is to study the problem of reciprocity prediction: we formulate several variants of the problem, all of them oriented around determining whether a link between users v and w is reciprocated or unreciprocated; and we identify a set of features based on the characteristics of v , w , and the nodes connected to v or w . (The precise definitions for these variants of the problem will be given in the next section.) Our analysis extracts features that have strong predictive power for this task. We find that differences in reciprocity can be related to the notion of *status*. Roughly speaking, people with similar status often participate in reciprocated interactions (e.g. messages between friends), while those with disparate status often participate in unreciprocated interactions (e.g. messages from fans to celebrities). In particular, we find that measures that formalize the relative “flow” of links from v to w , compared with the corresponding flow from w to v , constitute an important source of information for this task.

Our second approach involves comparing the structure of two subgraphs, one consisting of just the reciprocated links and the other consisting of just the unreciprocated links. We find that these structures exhibit important differences, including the presence of greater clustering and a larger giant component in the subgraph of symmetric links. Moreover, we find that almost all highly active Twitter users take part in at least some reciprocated interactions, while a non-trivial fraction take part in no unreciprocated interactions.

B. Related work

As noted above, several recent papers have discussed the heterogeneity of relationship types on Twitter, although they do not consider reciprocity as a measure [1], [2]. In the context of email, Tyler and Tang considered the dynamics of replying

to messages, which is the analogue of reciprocation in that domain [3]; however, the specifics of their analysis are quite different from what we pursue here.

Predicting reciprocity is related to other prediction tasks concerned with the links of an underlying network, but it is different in several important respects. Specifically, the *link prediction problem* seeks to identify links (v, w) that are currently missing in a network snapshot but are likely to form in the near future [4]. A key contrast between reciprocity prediction and link prediction is that the formation of any particular link is a rare event, whereas reciprocating an existing link (v, w) with the reverse edge (w, v) is common on sites such as Twitter. For this and other reasons, we find that features that have been observed to work well for link prediction are not the most effective for reciprocity prediction.

There has also been recent research predicting the *strengths* [5] and the *signs* [6] of links in on-line social networks. These works form interesting contrasts with reciprocity prediction. In particular, it is possible for a strong directed tie to be either reciprocated or unreciprocated. For example, an avid follower of the New York Times might regularly generate messages such as “@nytimes reports that ... ” without ever receiving a message from the @nytimes Twitter account. In the problem of sign prediction the notion of status is crucial [6], as it is in our work; but there are important distinctions in that reciprocated and unreciprocated links can easily exhibit either type of sign.

II. PROBLEM DEFINITION

We now formalize the problem of predicting link reciprocity. The communication network is represented as a directed graph $G = (V, E)$, with an edge (v, w) indicating that v has sent w at least one @-message. (In keeping with the focus on communication activities, we use @-messages to define all the networks in our analysis, rather than other relationships such as v following w 's account.) The input to our prediction problem is the graph G and a node pair $\{v, w\}$, where at least one of the edges (v, w) or (w, v) is present in G ; information about all edges of G is provided, except that the presence or absence of the two potential edges (v, w) and (w, v) has been hidden. Our task is to predict the direction of edges between v and w , and we note that this can be formalized in two distinct ways. First, we consider a formulation in which we decide whether a $\{v, w\}$ relationship is symmetric (that is, both (v, w) and (w, v) exist) or is asymmetric (only one of the directed (v, w) , (w, v) relationship exists). In the second formulation, we ask whether the (w, v) edge is present given that the (v, w) edge exists. Intuitively, predicting a symmetric relationship between nodes is a more difficult task than predicting reciprocation in a specific direction, and we show that this is indeed the case.

A. Notation

The number of messages produced by users on Twitter exhibits a long-tail distribution – many users produce only a small number of messages. In this work, we focus on users who have produced a large number of @-messages, so that we are studying a user population for whom Twitter is a

significant communication medium. We consider subgraphs of the form $G_n = (V_n, E_n)$, where $V_n = \{v \mid v \in V, v \text{ sent } \geq n \text{ @-messages}\}$ and $E_n = \{e = (v, w) \mid e \in E, v \text{ and } w \in V_n\}$. This subgraph thus captures the @-messaging interactions between individuals who are prolific Twitter users. We use the notation $v \xrightarrow{k} w$ to indicate that v sent at least k @-messages to w , and $v \xrightarrow{=k} w$ to indicate that v sent exactly k @-messages to w . From this definition we can parametrize reciprocity in terms of k . We say that an edge (v, w) is *reciprocated* if both $v \xrightarrow{k} w$ and $w \xrightarrow{k} v$, and is *unreciprocated* if $v \xrightarrow{k} w$ and $w \xrightarrow{=0} v$. Let the set of reciprocated edges be denoted E_k^r and the set of unreciprocated edges be denoted E_k^u . Finally, let $\text{deg}^-(v)$ and $\text{deg}^+(v)$ respectively denote the indegree and outdegree of node v , $\text{msg}^-(v)$ and $\text{msg}^+(v)$ be the numbers of messages received and sent by v , $\Gamma^-(v) = \{w \mid (w, v) \in E\}$ be the set of people who sent messages to v , and $\Gamma^+(v) = \{w \mid (v, w) \in E\}$ be the set of people to whom v sent messages.

B. Dataset description

We extracted the @-message graph from a large crawl of Twitter that took place between August 2009 and January 2010. More than three-billion messages from over 60 million users were collected in this data set [7]. The @-message graph is constructed by looking at messages a user v authors which mention user w at the beginning of the tweet. The graph G of users who authored at least one @-message contains 12,795,683 distinct users who sent a total of 819,305,776 @-messages, with 156,868,257 distinct directed interactions.

We focus our analysis on the subgraph G_{1000} induced by users who authored at least 1000 @-messages. In G_{1000} , which includes 181,033 users, we find that $|E_{10}^r| = 797,342$ and $|E_{10}^u| = 349,258$

III. METHODS FOR RECIPROCITY PREDICTION

Intuitively, features that capture whether v and w have similar status or a similar social circle should be potentially useful in predicting reciprocation. This section presents the various features that we use for predicting reciprocity in networks. Each feature corresponds to a method that assigns a value $\text{val}(v, w)$ to a node pair (v, w) , or a value $\text{val}(v)$ to a single node v . For each feature, we look at its value and whether the edge in question is reciprocated; this data is used for training models to predict reciprocity.

Given the values corresponding to all node pairs (or nodes) in question, we can then choose threshold values or ranges where we predict reciprocity, and predict a lack thereof in the complementary region. We consider a simple threshold classification scheme which predicts that a node pair (u, v) is unreciprocated if the feature value is less than (or greater than) some threshold, and is reciprocated otherwise. For each feature, we determine the threshold value val_{OPT} and threshold direction (less than / greater than) to maximize prediction accuracy according to this threshold classifier. For example, a sufficiently high number of mutual neighbors for the nodes v and w could strongly indicate the existence of

a reciprocated link between them. The list of features we consider is summarized in Table I.

As previously mentioned, there are two ways that we can formulate the prediction problem, and we present four different formulations. The first addresses the question of symmetry, while the other three examine the problem of reciprocity in which the available information about the nodes in question is limited:

- 1) SYM (predicting symmetry): predict whether both (v, w) and (w, v) exist, or whether exactly one of (v, w) or (w, v) exists, using information about v and w but not about the presence or absence of communication between them.
- 2) REV (predicting a reverse edge): predict whether a reverse edge exists given that the forward edge (v, w) exists, using information about v and w .
- 3) REV- w (predicting a reverse edge using only w): predict whether a reverse edge exists given that (v, w) exists, but using only information about w .
- 4) REV- v (predicting a reverse edge using only v): predict whether a reverse edge exists given that (v, w) exists, but using only information about v .

Within this framework, we can compare the predictive power of specific features of the @-message graph, as well as more complex classifiers (such as decision trees) that utilize multiple features.

A. Degree and message features

It seems intuitive that the relative indegree or outdegree of nodes would indicate whether a pair of nodes are in a one-sided or two-sided relationship. If both have a similar indegree, this might indicate that they have similar social status in the network. In contrast, disproportionate indegrees could indicate that one user is a celebrity and the other is a non-celebrity, making it less likely that their relationship is reciprocated.

We now describe the features in our analysis that are based on degree and message counts. Both relative (e.g., the ratio of indegrees) and absolute (e.g., $\text{deg}^-(w)$) feature measures are considered:

- *Indegree and outdegree ratio* both measure the ratio of outdegrees or indegrees of two nodes, and we define $\text{val}(v, w) = \text{deg}^-(v)/\text{deg}^-(w)$ or $\text{deg}^+(v)/\text{deg}^+(w)$, respectively.
- *Incoming message and outgoing message ratio* are similar, but uses the total number of messages that a node receives or sends, regardless of the nodes to which messages are sent or from which messages are received. Specifically, we use the analogous measures discussed above for degree, but with $\text{msg}^-(v)$ and $\text{msg}^+(v)$ playing the roles of $\text{deg}^-(v)$ and $\text{deg}^+(v)$, respectively.
- *Incoming message/indegree ratio and outgoing message/outdegree ratio* compares the ratio of two nodes' incoming message to indegree ratio or outgoing message to outdegree ratio. A high incoming message to indegree ratio might characterize users who have a small group

of friends with which they exchange many messages. Alternatively, a low incoming message to indegree ratio could characterize highly connected people in a network, since the messages they receive are distributed over many more users.

- *Outdegree/indegree ratio* is a heuristic that characterizes the messaging activity of a single node. A large outdegree/indegree ratio might indicate a user of celebrity status because she receives many messages from many followers but sends relatively few messages. Given a pair of nodes we can compute the outdegree/indegree ratio as $\text{val}(v, w) = \frac{\text{deg}^+(v)}{\text{deg}^-(v)} / \frac{\text{deg}^+(w)}{\text{deg}^-(w)}$.

B. Link prediction features

It is not intuitive whether methods that work well for link prediction would work well in predicting reciprocity. While link prediction asks whether an edge between two nodes exists, reciprocity prediction asks whether a pair known to have at least a directed edge in fact has a bi-directional pair of edges. The following are some measures used for link prediction:

Newman [8] showed that the number of common neighbors in a collaboration network can be a predictor of future links.

Mutual neighbors calculates the number of people to whom both v and w send messages ($|\Gamma^+(v) \cap \Gamma^+(w)|$), or the number of people from whom both v and w receive messages ($|\Gamma^-(v) \cap \Gamma^-(w)|$).

Jaccard's coefficient [9] is a similarity measure that we apply to the concept of mutual neighbors. We calculate the similarity between two sets by taking the ratio of the cardinality of their intersection and their union:

$$\text{val}(v, w) = \frac{|\Gamma^-(v) \cap \Gamma^-(w)|}{|\Gamma^-(v) \cup \Gamma^-(w)|}.$$


Adamic and Adar [10] defined the similarity between Web sites v, w to be $\sum_{\{x|v,w \text{ share feature } x\}} \frac{1}{\log \text{frequency}(x)}$, and we similarly define $\text{val}(v, w)$ to be

$$\sum_{\{x|x \in \Gamma^-(v) \cap \Gamma^-(w)\}} \frac{1}{\log \text{deg}^-(x)}.$$

Preferential attachment is another popular heuristic in modeling network growth, where the probability that an edge forms with a specific node is proportional to its existing indegree. Newman [8] and Barabasi et al. [11] showed that the product of the in-degrees of two nodes in a co-authorship network can be a predictor of a future link between the nodes. We apply preferential attachment in a slightly different way and define $\text{val}(v, w) = \text{deg}^-(v) \cdot \text{deg}^+(w)$ or $\text{deg}^+(v) \cdot \text{deg}^-(w)$. Notice that taking the ratio of these two values is equivalent to the outdegree/indegree ratio between two nodes.

Two-step paths (ratio) is a simplification of Katz's [12] measure of status by calculating the number of paths between two nodes. In this work, we only consider paths of length 2, and define $\text{val}(v, w) = |\text{paths}^2(v, w)|$, where $\text{paths}^2(v, w)$ is the set of paths from v to w of length 2. The two-step paths ratio is simply the ratio of the number of directed two-step paths from v to w to that from w to v .

Fig. 1: Two-step hops



C. Different sets of features

The features above have been shown to work well for the related but different task of predicting the existence of a link. Since our task is to predict whether a link is reciprocated we examine several additional features. For convenience, we further break them down into four sets:

- 1) Absolute degree/message features - degree, messages, message-degrees, outdegree-indegrees
- 2) Relative degree/message features - degree ratios, message ratios, message-degree ratios, and outdegree-indegree ratios
- 3) Two-step hop features - mutual neighbors (in and out), and two step paths (v to w and w to v)
- 4) Link prediction features - all other link prediction features not mentioned

D. Two-step paths

The importance of “friends of friends,” or people two links away from a given node, lends itself to exploring features that directly arise from the directed @-message graph. There are essentially four types of two-step hops (shown in Figure 1) corresponding to either the number of common in-neighbors or out-neighbors (mutual neighbors), or the number of directed paths from v to w or from w to v (two-step paths).

If both v and w send messages to many common people, it is likely that they are in the same social circle, or that they mention the same celebrities. If v and w receive many messages from the same group of people, it could be that both v and w are in the same community, or that they are celebrities with overlapping fan-bases.

As the number of paths from v to w increases, there are two conflicting forces: v has a stronger source of connections to w , but at the same time w is more popular and hence less likely to reciprocate the (v, w) edge. The reverse case is simpler—intuitively, as the number of paths from w to v increases, the likelihood that w will communicate with v grows.

IV. RESULTS AND DISCUSSION

A. Individual properties

To calculate the accuracy of the individual heuristics, we calculated val for each feature on the subset $E_{10}^I \cup E_{10}^U$ of the graph G_{1000} , where equal numbers of edges were taken from the sets of reciprocated and unreciprocated edges. This gives a baseline accuracy of 0.500, achievable by predicting that all edges are of one type. We applied the SYM and REV

TABLE I: Reciprocity Prediction Features

Feature	$\text{val}(v)$ or $\text{val}(v, w)$
<i>Absolute degree/message features</i>	
Indegree or outdegree	$\text{deg}^-(v)$ or $\text{deg}^+(v)$
Incoming or outgoing messages	$\text{msg}^-(v)$ or $\text{msg}^+(v)$
Message-degree (in or out)	$\frac{\text{msg}^-(v)}{\text{deg}^-(v)}$ or $\frac{\text{msg}^+(v)}{\text{deg}^+(v)}$
Outdegree-indegree	$\frac{\text{deg}^+(v)}{\text{deg}^-(v)}$
<i>Relative degree/message features</i>	
Indegree ratio	$\text{deg}^-(v) / \text{deg}^-(w)$
Outdegree ratio	$\text{deg}^+(v) / \text{deg}^+(w)$
Incoming message ratio	$\text{msg}^-(v) / \text{msg}^-(w)$
Outgoing message ratio	$\text{msg}^+(v) / \text{msg}^+(w)$
Message-degree ratio (in)	$\frac{\text{msg}^-(v)}{\text{deg}^-(v)} / \frac{\text{msg}^-(w)}{\text{deg}^-(w)}$
Message-degree ratio (out)	$\frac{\text{msg}^+(v)}{\text{deg}^+(v)} / \frac{\text{msg}^+(w)}{\text{deg}^+(w)}$
Outdegree-indegree ratio	$\frac{\text{deg}^+(v)}{\text{deg}^-(v)} / \frac{\text{deg}^+(w)}{\text{deg}^-(w)}$
<i>Link prediction features</i>	
Mutual neighbors (in)	$ \Gamma^-(v) \cap \Gamma^-(w) $
Mutual neighbors (out)	$ \Gamma^+(v) \cap \Gamma^+(w) $
Jaccard's coefficient (in)	$\frac{ \Gamma^-(v) \cap \Gamma^-(w) }{ \Gamma^-(v) \cup \Gamma^-(w) }$
Jaccard's coefficient (out)	$\frac{ \Gamma^+(v) \cap \Gamma^+(w) }{ \Gamma^+(v) \cup \Gamma^+(w) }$
Adamic/Adar	$\frac{1}{\sum_{\{x x \in \Gamma^-(v) \cap \Gamma^-(w)\}} \log \text{deg}^-(x)}$
Preferential attachment (v to w)	$\text{deg}^+(v) \cdot \text{deg}^-(w)$
Preferential Attachment (w to v)	$\text{deg}^+(w) \cdot \text{deg}^-(v)$
Two-step paths (v to w)	$ \text{paths}^2(v, w) $
Two-step paths (w to v)	$ \text{paths}^2(w, v) $
Two-step paths ratio	$\frac{ \text{paths}^2(v, w) }{ \text{paths}^2(w, v) }$

mechanisms to feature sets 2-4, and REV- v and REV- w to set 1.

As described earlier, we picked a threshold value val_{OPT} to optimize prediction accuracy: we predicted reciprocity above the threshold, and non-reciprocity below (or vice versa depending on which performed better). Tables III and IV summarize the performance of each heuristic on the subgraph G_{1000} , $k = 10$, while table II summarizes the different mechanisms of prediction for a single heuristic.

In tables III and IV, a star (*) indicates that reciprocity was predicted when val was below the threshold, and a lack thereof indicates reciprocity was predicted when val was above the threshold.

In table II, SYM⁺ refers to the prediction mechanism where we aim to predict symmetry and predict all edges with values above val_{OPT} to be reciprocated, and REV⁻ refers to the mechanism where we aim to predict whether a reverse edge (w, v) exists given (v, w) and predict all edges with values below val_{OPT} to be reciprocated.

1) *Comparison of prediction mechanisms:* We observe higher accuracy for the REV task than SYM, as REV is “easier” than SYM since we know more information about the edge (v, w) .

Comparing REV- v to REV- w , we see REV- w obtains higher accuracy, suggesting that when trying to predict the existence

(w, v) of given (v, w) , knowing about properties of w is more valuable than knowing properties of v .

Note that SYM^- , REV^- , $\text{REV}-w^+$ and $\text{REV}-v^-$ are such poor predictors that simply predicting that everything was reciprocated (or unreciprocated) would have been better.

2) *Comparison of methods of prediction:*

a) *Trends:* On the whole, outdegree-indegree ratio and the two-step paths ratio are the best indicators of reciprocity. In fact, outdegree-indegree ratio alone already achieves accuracy to within $\pm 5\%$ of a decision tree using every feature.

b) *Sending and receiving:* When we look at features using one of the four mechanisms, we find that for the majority of the features larger values indicate reciprocity is more likely to occur. However, the *smaller* the outdegree-indegree ratio, the more likely reciprocation occurs. In other words, a large denominator and small numerator in $\frac{\text{deg}^+(v)/\text{deg}^+(w)}{\text{deg}^-(v)/\text{deg}^-(w)} = \frac{\text{deg}^+(v)\text{deg}^-(w)}{\text{deg}^-(v)\text{deg}^+(w)}$ is a good indicator of reciprocation. A large denominator and small numerator indicate that v has many in-links and few out-links and that w has many out-links and few in-links. This suggests that v has higher “status” than w and hence increases the probability that w links to v .

Interestingly, separating the numerator and denominator from the outdegree-indegree ratio above, which corresponds to our two preferential attachment features, leads to very different results. While a small numerator does reasonably well (preferential attachment (v to w)), a large denominator does not (preferential attachment (w to v)) and performs only marginally better than chance. It is reasonable that the numerator and denominator would individually perform worse than the ratio, since the ratio takes into account both of them. However, the fact that the denominator performs so poorly is surprising because a large denominator suggests that v has a higher status than w ; this could increase the chance that w links to v even if w were to randomly link to others. On the other hand, a small numerator provides some information about status but not about increased reciprocation under random linking. The fact that a small numerator is more important than a large denominator suggests that status, as measured by the number of in- and out-links, is a powerful predictor of reciprocity.

If we consider only edges $E_{=} = \{(v, w) \mid \text{deg}^-(v) = \text{deg}^-(w)\}$ (edges connecting nodes of equal degree), the accuracy of the outdegree ratio feature increases to 0.811. This result is comparable to what we get for the outdegree-indegree ratio feature.

c) *REV-v vs. REV-w:* $\text{REV}-w$ performs better than $\text{REV}-v$ on almost all features, and where $\text{REV}-v$ performs better, the difference is not as great. The fact that information about w is more useful than information about v suggests a contrast for various domains of potential application: for example, if we think of v as sending information to w via the (v, w) communication link (consider for example a marketer v contacting a potential customer w), then we find that knowledge of the recipient (w) tells us more about the probability of a response than knowledge of the sender (v).

TABLE II: Indegree performance - different methods

Mechanism	val _{OPT} (Percentile)	Accuracy
<i>Indegree ratio</i>		
SYM ⁺	0.256 (40)	0.702
SYM ⁻	-	-
REV ⁺	0.414 (46)	0.759
REV ⁻	-	-
<i>Indegree of v or w</i>		
REV- w^+	-	-
REV- w^-	74 (61)	0.731
REV- v^+	61 (60)	0.582
REV- v^-	-	-

TABLE III: Reciprocity Prediction Feature Performance: Individual (REV)

Feature	val _{OPT} (Percentile)	Accuracy
Indegree ratio	0.414 (46)	0.759
Outdegree ratio	0.667 (43)	0.628
Incoming message ratio	0.333 (48)	0.772
Outgoing message ratio	0.905 (46)	0.547
Incoming message-indegree ratio	0.650 (39)	0.569
Outgoing message-outdegree ratio	0.791 (33)	0.615*
Outdegree-indegree ratio	1.72 (53)	0.820*
Mutual neighbors (in)	10 (61)	0.552
Mutual neighbors (out)	8 (51)	0.580
Jaccard's coefficient (in)	0.0345 (48)	0.684
Jaccard's coefficient (out)	0.0637 (55)	0.660
Adamic/Adar	1.94 (55)	0.561
Two-step paths (v to w)	6 (59)	0.517*
Two-step paths (w to v)	5 (51)	0.657
Two-step paths ratio	0.556 (52)	0.760
Preferential attachment (v to w)	10230 (58)	0.687*
Preferential attachment (w to v)	2610 (37)	0.534*

TABLE IV: Reciprocity Prediction Feature Performance: Individual (REV- v ,REV- w)

Feature	val _{OPT} (Percentile)	Accuracy
Indegree (v)	61 (60)	0.582
Indegree (w)	148 (61)	0.731*
Outdegree (v)	25 (14)	0.506*
Outdegree (w)	105 (60)	0.647*
Incoming messages (v)	619 (53)	0.637
Incoming messages (w)	1802 (54)	0.733*
Outgoing messages (v)	906 (51)	0.542
Outgoing messages (w)	506 (17)	0.524*
Incoming message-indegree (v)	9.4 (41)	0.596
Incoming message-indegree (w)	9.12 (30)	0.535
Outgoing message-outdegree (v)	13.2 (50)	0.523
Outgoing message-outdegree (w)	8.14 (36)	0.661
Outdegree-indegree (v)	1.28 (53)	0.679*
Outdegree-indegree (w)	0.747 (50)	0.777

B. *Decision tree analysis*

We can also combine subsets of features and evaluate their performance by randomly splitting the edges in $E_{10}^u \cup E_{10}^v$

TABLE V: Decision Tree Accuracy

Set	Accuracy	Top-level attribute
Degree/message (1)	0.832	Outdegree-indegree (w)
Degree/message ratio (2)	0.861	Outdegree-indegree ratio
Two-step hops (3)	0.796	Two-step paths (w to v)
Link prediction (4)	0.739	Two-step paths ratio (directed)
<i>Combined</i>		
All ratio (2,3,4)	0.861	Outdegree-indegree ratio
All absolute (1,3,4)	0.832	Outdegree-indegree (w)
All (1-4)	0.862	Outdegree-indegree ratio

into two sets and performing 2-fold cross-validation. We use the ID3 algorithm to train the decision tree classifiers, and because the val features are continuous, we quantize each feature into deciles (dividing the data equally into tenths) to reduce computation time.

We consider the following combined sets of features, as well as each set individually:

- 1) **All** (sets 1-4) – every single feature was considered.
- 2) **All ratio** (sets 2,3,4) – all features that used ratios were considered.
- 3) **All absolute** (sets 1,3,4) – this allows us to see how using only “absolute” features affects accuracy.

Table V shows the accuracy of the trees and the most important attribute for the different sets of features. We find that using only degree/message features (set 1) performs as well as using all absolute features (sets 1, 3, 4). The two-step paths ratio alone obtains an accuracy of 0.760, while the decision tree for link prediction only manages 0.739. This can be attributed to inaccuracies introduced while quantizing the continuous features. Furthermore, features commonly used for link prediction yield a tree of lower accuracy than other features, providing evidence that the problem of reciprocity prediction is different from link prediction. Whenever the outdegree-indegree value or ratio was included in the feature vector, it was the single most important variable.

If we only consider $E_{=}$, node pairs with equal indegree ($|E_{=}| = 16,311$), the accuracy of All ratio drops to 0.776. This suggests that predicting reciprocity becomes considerably more difficult as we lose the ability to differentiate between nodes of different status or indegree. Again, as the indegrees of v and w were equal in every pair (v, w) , it is not surprising that the outdegree ratio is the most important feature.

1) *Performance*: Classifying features based on their computation time, the two-step hop and link prediction features (excluding preferential attachment) take more than two orders of magnitude longer to compute than all other features (500 times as long for $k = 10$). If we only use the other features in prediction, we still obtain an accuracy of 0.862, similar to what we obtained above when we used all features. Therefore, it appears sufficient to use only these features in practical applications.

2) *Effect of k on accuracy*: As k increases, the proportion of edges that are defined as reciprocated increases, and natu-

TABLE VI: Logistic regression – relative degree/message-based features

Feature	β	p value
Indegree ratio	0.0101903	$< 2 \times 10^{-16}$
Outdegree ratio	0.0005775	0.2545
Incoming messages ratio	0.0230161	$< 2 \times 10^{-16}$
Outgoing messages ratio	-0.0047152	$< 2 \times 10^{-16}$
Incoming messages-indegree ratio	-0.0005545	0.0798
Outgoing messages-outdegree ratio	-0.0049387	$< 2 \times 10^{-16}$
Outdegree-indegree ratio	-0.0562983	$< 2 \times 10^{-16}$

TABLE VII: Logistic regression – two-step hop features

Feature	β	p value
Mutual neighbors (in)	-0.0117269	$< 2 \times 10^{-16}$
Mutual neighbors (out)	0.0180579	$< 2 \times 10^{-16}$
Two-step paths (v to w)	-0.1193624	$< 2 \times 10^{-16}$
Two-step paths (w to v)	0.1296081	$< 2 \times 10^{-16}$

TABLE VIII: Logistic regression – All ratio

Feature	β	p value
Indegree ratio	0.0120256	$< 2 \times 10^{-16}$
Outdegree ratio	-0.0015554	0.005739
Incoming messages ratio	0.0145437	$< 2 \times 10^{-16}$
Outgoing messages ratio	-0.0043189	$< 2 \times 10^{-16}$
Incoming messages-indegree ratio	0.0048525	$< 2 \times 10^{-16}$
Outgoing messages-outdegree ratio	-0.0046674	$< 2 \times 10^{-16}$
Outdegree-indegree ratio	-0.0301592	$< 2 \times 10^{-16}$
Mutual Neighbors (in)	-0.0279290	$< 2 \times 10^{-16}$
Mutual Neighbors (out)	0.0147103	$< 2 \times 10^{-16}$
Two-step paths (v to w)	-0.0530463	$< 2 \times 10^{-16}$
Two-step paths (w to v)	0.0182572	$< 2 \times 10^{-16}$
Two-step paths ratio	0.0394657	$< 2 \times 10^{-16}$
Jaccard (in)	-0.0238541	$< 2 \times 10^{-16}$
Jaccard (out)	0.0572358	$< 2 \times 10^{-16}$
Adamic-Adar	-0.0001424	0.881637
Preferential attachment (v to w)	0.0010837	0.000627
Preferential attachment (w to v)	-	-

rally accuracy also increases (0.8818 for $k = 20$ and 0.9032 for $k = 50$). If we instead take equal numbers of reciprocated and unreciprocated edges, giving a baseline accuracy of 0.500, accuracy gradually increases from 0.836 for $k = 10$ to 0.846 for $k = 30$.

C. Regression analysis

We also used a logistic regression model on subsets of features, where $f(z) = \frac{e^z}{e^z + 1}$, $z = \beta_0 + \beta F$, $f(z)$ is binary (1 when an edge is reciprocated, 0 otherwise) and F is the vector of features. The results are shown in tables VI–IX, where struck-out p-values indicate insignificant features. Here, the two-step paths (v to w), two-step paths ratio, and Jaccard (out) features are most significant when simultaneously using all features for classification.

TABLE IX: Logistic regression - All

Feature	β	p value
Indegree ratio	0.0041791	6.09×10^{-8}
Outdegree ratio	0.0046914	1.28×10^{-13}
Incoming messages ratio	0.0029794	0.000125
Outgoing messages ratio	0.0033361	3.10×10^{-10}
Incoming messages-indegree ratio	0.0040884	1.11×10^{-14}
Outgoing messages-outdegree ratio	-0.0015075	0.006283
Outdegree-indegree ratio	-0.0057958	$< 2 \times 10^{-16}$
Indegree (v)	0.0050520	4.30×10^{-11}
Indegree (w)	-0.0089197	$< 2 \times 10^{-16}$
Outdegree (v)	-0.0063247	$< 2 \times 10^{-16}$
Outdegree (w)	0.0035881	3.58×10^{-7}
Incoming messages (v)	0.0063390	$< 2 \times 10^{-16}$
Incoming messages (w)	-0.0179975	$< 2 \times 10^{-16}$
Outgoing messages (v)	-0.0070869	$< 2 \times 10^{-16}$
Outgoing messages (w)	0.0095572	$< 2 \times 10^{-16}$
Incoming message-indegree (v)	-0.0023250	1.11×10^{-5}
Incoming message-indegree (w)	-0.0004044	0.42781
Outgoing message-outdegree (v)	0.0007430	0.175454
Outgoing message-outdegree (w)	0.0024155	2.54×10^{-5}
Outdegree-indegree (v)	-0.0110324	$< 2 \times 10^{-16}$
Outdegree-indegree (w)	0.0218874	$< 2 \times 10^{-16}$
Mutual Neighbors (in)	-0.0194635	$< 2 \times 10^{-16}$
Mutual Neighbors (out)	0.0050245	$< 2 \times 10^{-16}$
Two-step paths (v to w)	-0.0462950	$< 2 \times 10^{-16}$
Two-step paths (w to v)	0.0167156	$< 2 \times 10^{-16}$
Two-step paths ratio	0.0440107	$< 2 \times 10^{-16}$
Jaccard (in)	-0.0398243	$< 2 \times 10^{-16}$
Jaccard (out)	0.0561815	$< 2 \times 10^{-16}$
Adamic-Adar	0.0111504	$< 2 \times 10^{-16}$
Preferential attachment (v to w)	0.0009537	0.003002
Preferential attachment (w to v)	-	-


V. TWITTER AS A SUPERPOSITION OF NETWORKS

A. (Un)reciprocated subgraph analysis

We also analyze how various properties of the subgraphs G_n , as well as the edge sets E_k^r and E_k^u , vary as we adjust n and k .

Reciprocated and unreciprocated edges: We observe that the frequency of reciprocated edges is approximately 2 to 3 times that of unreciprocated edges, and the proportion of reciprocated edges increases as n and k increases (Fig. 2, 3). While reciprocated communication is the dominant form of interaction, we also see a significant number of unreciprocated interactions, indicating that a significant number of relationships on Twitter are unbalanced. This could occur when a user of lower status repeatedly invokes the name of a more influential user (of higher status) through @-mentioning, as suggested in the introduction.

Reciprocated and unreciprocated nodes: A significant proportion of nodes take part in both reciprocated and unreciprocated interactions, and while a majority of nodes have reciprocated interactions, only a small proportion have purely unreciprocated interactions. This indicates that while there are

Fig. 2: Proportion of nodes or edges (varying n)

two distinct types of relationships occurring on Twitter, they do not correspond to two distinct types of users. The fact that it is rare to find active Twitter users taking part in only unreciprocated interactions suggests that social (reciprocal) relationships are associated with an active and continued use of the site.

We can also see this in a scatter plot of the number of users v with each of three types of interaction (Fig. 6): (i) reciprocated interactions with both (v, w) and (w, v) present; (ii) unreciprocated “out-going” interactions with only (v, w) present; and (iii) unreciprocated “in-coming” interactions with only (w, v) present. We thus differentiate between both ends in an unreciprocated edge ($v \xrightarrow{k} w$ and $w \xrightarrow{=0} v$), where a user could play the role of v if she’s not replied to, or the role of w if she doesn’t reply. From the plot, we can see that the types with the greatest numbers of associated nodes are those with only reciprocated interactions.

Clustering coefficient remains relatively stable as n and k vary: The clustering coefficient is much larger in the subgraph of reciprocated edges, which corresponds to the natural notion that reciprocated edges represent more social activity with a larger density of triangles. The fact that these quantities are stable as we change n and k suggests that the network properties of these subgraphs do not change significantly even if we sample from a relatively smaller population of all users (Fig. 4).

Connected component remains stable as n varies, but decreases as k increases: The graphs corresponding to E_k^r and E_k^u have giant components for relatively low values of k . However, the size of the largest component shrinks rapidly once k passes a particular range (roughly between 50 and 100). (Fig. 5). This hints at a kind of qualitative transition in the structure of the network as a function of message volume, with the edges representing more than 100 communications each unable to sustain a very large component on their own.

VI. CONCLUSION

We have formulated several variants of the problem of reciprocity prediction in an on-line social network and have




Fig. 3: Proportion of nodes or edges (varying k)




Fig. 6: Scatter plot of users' interaction types




Fig. 4: Clustering coefficient




Fig. 5: Proportion in largest connected component

demonstrated good predictive ability using properties of nodes and their local network neighborhoods. In particular, features that approximate the relative status of two nodes v and w appear to be the most effective at predicting whether a link between v and w is reciprocated. We have also studied the subgraphs of reciprocated and unreciprocated links, finding that they differ in important ways with respect to component sizes and clustering.

There are a number of interesting further directions that could be pursued on this problem. First, our approach has used structural features of the system, but it would be interesting to combine these with other data such as textual content to obtain stronger estimates of reciprocation. Second, temporal information is another important source of data in analyzing

reciprocation, and understanding how reciprocated and unreciprocated interactions develop over time could provide further insight into the problem. Finally, we should remember that link reciprocity is one aspect of the broader issue of heterogeneity in the types of interactions encountered on a site such as Twitter; finding other dimensions along which to express this heterogeneity is an important question as well.

ACKNOWLEDGMENT

This work was supported in part by a Google Research Grant, a Yahoo! Research Alliance Grant, and NSF grants IIS-0910664, CCF-0910940, and IIS-1016099.

REFERENCES

- [1] H. Kwak, C. Lee, H. Park, and S. B. Moon, "What is Twitter, a social network or a news media?" in *Proc. 19th Intl. WWW Conf.*, 2010.
- [2] D. M. Romero and J. M. Kleinberg, "The directed closure process in hybrid social-information networks, with an analysis of link formation on Twitter," in *Proc. 4th Intl. Conf. on Weblogs and Social Media*, 2010.
- [3] J. R. Tyler and J. C. Tang, "When Can I Expect an Email Response? A Study of Rhythms in Email Usage," in *ECSCW'03: Proceedings of the eighth conference on European Conference on Computer Supported Cooperative Work*. Kluwer Academic Publishers, Sep. 2003.
- [4] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *Journal of the American Society for Information Science and Technology*, vol. 58, no. 7, pp. 1019–1031, 2007.
- [5] E. Gilbert and K. Karahalios, "Predicting tie strength with social media," in *Proc. 27th ACM Conference on Human Factors in Computing Systems*, 2009, pp. 211–220.
- [6] J. Leskovec, D. Huttenlocher, and J. Kleinberg, "Signed networks in social media," in *Proc. 28th ACM Conference on Human Factors in Computing Systems*, 2010, pp. 1361–1370.
- [7] D. M. Romero, B. Meeder, and J. Kleinberg, "Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on Twitter," in *Proc. 20th Intl. WWW Conf.*, 2011.
- [8] M. E. J. Newman, "Clustering and preferential attachment in growing networks," *Physical Review E*, no. 2, pp. 025 102+, Jul. 2001.
- [9] G. Salton and M. J. McGill, *Introduction to Modern Information Retrieval*. New York, NY, USA: McGraw-Hill, Inc., 1986.
- [10] L. Adamic, "Friends and Neighbors on the Web," *Social Networks*, 2003.
- [11] A. Barabasi, H. Jeong, Z. Neda, E. Ravasz, A. Schubert, and T. Vicsek, "Evolution of the social network of scientific collaborations," *Physica A: Statistical Mechanics and its Applications*, vol. 311, 2002.
- [12] L. Katz, "A New Status Index Derived From Sociometric Analysis," *Psychometrika*, 1953.