

Friend or Frenemy?

Predicting Signed Ties in Social Networks

Shuang Hong Yang
College of Computing
Georgia Tech
shy@gatech.edu

Alex Smola
Yahoo! Research
smola@yahoo-inc.com

Bo Long
Yahoo! Labs
bolong@yahoo-inc.com

Hongyuan Zha
College of Computing
Georgia Tech
zha@cc.gatech.edu

Yi Chang
Yahoo! Labs
yichang@yahoo-inc.com

ABSTRACT

We study the problem of *labeling the edges* of a social network graph (e.g., acquaintance connections in Facebook) as either *positive* (i.e., trust, true friendship) or *negative* (i.e., distrust, possible frenemy) relations. Such signed relations provide much stronger signal in tying the behavior of online users than the unipolar Homophily effect, yet are largely unavailable as most social graphs only contain unsigned edges.

We show the surprising fact that it is possible to infer signed social ties with good accuracy *solely* based on users' behavior of decision making (or using only a small fraction of supervision information) via unsupervised and semi-supervised algorithms. This work hereby makes it possible to turn an unsigned *acquaintance network* (e.g. Facebook, Myspace) into a signed *trust-distrust network* (e.g. Epinion, Slashdot). Our results are based on a mixed effects framework that simultaneously captures users' behavior, social interactions as well as the interplay between the two. The framework includes a series of latent factor models and it also encodes the principles of balance and status from Social psychology. Experiments on Epinion and Yahoo! Pulse networks illustrate that (1) signed social ties can be predicted with high-accuracy even in fully unsupervised settings, and (2) the predicted signed ties are significantly more useful for social behavior prediction than simple Homophily.

Categories and Subject Descriptors

H.5.3 [Information Systems]: *Web-based Interaction*;

H.3.3 [Information Search and Retrieval]: *Information filtering* H.2.8 [Database Management]: *Data Mining*

General Terms: Algorithms, Performance

Keywords: Social ties, edge labeling, signed network, social behavior prediction, Homophily, sociological principles

1. INTRODUCTION

Social networks are playing an increasingly important role in shaping the business models of today's internet industry.

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

SIGIR'12, August 12–16, 2012, Portland, Oregon, USA.

Copyright 2012 ACM 978-1-4503-1472-5/12/08 ...\$10.00.

Not only are most emerging online services providing functionality to facilitate social interactions in their systems (e.g. last.fm, slashdot, Flickr, Yelp), but even many traditional internet systems are actively exploiting social networks to enhance their services, marketing and revenue, e.g., search engines (e.g., Bing), online retailers (e.g., Amazon), Web portals (e.g., Yahoo!). This is partly due to the widespread faith in the social effect of *Homophily* [20] that people socially acquainted tend to behave similarly. However, the extent to which this belief holds true is debatable — after all, it is unreasonable to expect two users who know each other (e.g., who are connected in Facebook) to behave alike in every aspect of their online decision making (e.g., purchase, ad click, movie rental, search intent, political view, sentimental moods). Instead, a much stronger signal of social ties is the network of trust relationships, since in many decisions we make, the opinions from the people we *trust* are what we really care about and what our decisions usually get influenced by. Unfortunately, this information is largely unavailable given the fact that the vast majority of social networks only contain acquaintance relationships, without distinguishing whom you trust from whom you don't.

Quite surprisingly, we show it is possible to infer a signed label for an acquaintance relationship — positive as trust (true friendship), negative as distrust (possible frenemy). In this way we make possible to turn an unsigned *acquaintance network* (e.g. Facebook, Myspace) to a signed *trust-distrust network* (e.g. Slashdot, Epinion). Moreover, by allowing us to tie the behavior of an online user to *whom he trusts* rather than *whom he knows*, the predicted signed relations are significantly more informative than Homophily for the purpose of behavior prediction.

1.1 Signed Social Ties

The main focus of this paper is (i) to predict the *signed social ties* (i.e., trust/distrust) out of the unsigned (i.e., acquaintance) relationships in social networks; and (ii) to examine how the predicted signed ties help behavior prediction. In particular, given that two users are socially connected (i.e., know each other), we would like to infer whether the connection is positive (e.g., trust, true friendship) or negative (e.g., distrust, possible frenemy); Furthermore, by tying the behavioral activities of online users with the inferred signed relationships, we examine how social targeting can be improved to promote online services, for example, to

provide better matched search results, News articles, games, advertisements, or products [21, 18, 17, 26].

By predicting signed social ties, we enable the social connections in a traditional social network to capture much richer relationships reflecting either *positive* (e.g. trust, approval, true friendship) or *negative* (e.g. distrust, opposition, frenemy) interactions. This will substantially extend the acquaintance relationship in standard social networks (e.g. Facebook), and also the unipolar tie-strength analysis of traditional link prediction [19, 22, 5], to a meaningful new dimension — we are able to *en-sign*¹ an unsigned network by turning it to a signed network. Moreover, models designed for signed social ties also have the potential to improve our understanding of social relations and the structure of our society, shedding light to sociological principles in general [6, 22, 5, 15].

1.2 Behavior-Relation Interplay

The interplay between social relations and user’s behavior of decision making is a key driving force underlying social activities [21, 18, 17]. On the one hand, social interactions can influence individual’s behavior of decision making [10, 1] — for instance, users can learn from each other, following their trusted friends’ good decisions and avoiding their mistakes. On the other hand, a user’s behavior could in turn impact his relationships with others — for instance, a user with malicious behavior turns to get bad reputations from others. In this paper, we present a framework for predicting the signed social ties by capturing such interplay. Specifically, we explore the following “mixture of effects” assumptions about social relations and users’ behavior of decision making:

1. The decision of a user is not decided solely by his own taste but also influenced by the opinions of other people whom he trusts.
2. The social relationship between two users depends not only on their intrinsic kinship but is also evolving according to the agreement between their opinions.

It is desirable to distinguish these two groups of factors when modeling behavior and social relationships. We present the *Behavior Relation Interplay* (BRI) model, a latent factor model that leverages behavioral evidence to infer social interactions and at the same time exploits the learned relations to tie users’ behavior. The key idea in BRI is to associate latent factors with both users and items and to define coupled models to simultaneously characterize both the social interactions and behavioral evidences.

A key question is whether it is possible at all to infer signed ties in the absence of any labels. One of the more surprising results of our work is that it is, indeed, possible. To achieve this we design both unsupervised and semi-supervised algorithms for signed social tie prediction. In the unsupervised setting, we only observe the acquaintance connections but none of the sign labels, as is the typical case in the vast majority of the acquaintance social network sites such as Facebook. We show that BRI achieves fairly accurate predictions. The key idea is a mechanism called “*automatic sign determination*” that discovers the signs of social ties automatically from user activity data. We also examine semi-supervised prediction, where a small fraction of the sign

¹With slight abuse of expression, here by “en-sign” we mean “to assign a sign to” something that is unsigned.

labels are given. We show that BRI model provides comparable performance to the leave-one-out prediction results (i.e. all but one sign are observable) of the topology-based algorithm [14].

To further improve the prediction accuracy and deepen our understanding of signed social ties, we extend the BRI model to encode general sociological principles, particularly the theories of *status* and *structure balance* [6, 15]. One nice property of our model is that it naturally provides a mapping from a user to his social status and an interesting visualization of the network. We show the sociological principles are useful for signed social tie prediction.

Finally, we show that by predicting signed social ties, BRI models achieve significantly better performance in behavior prediction than models equipped with Homophily. In particular, BRI gains 5.8% improvement of RMSE in rating prediction on Epinion data, and 6.9% improvement of nDCG in top- k ranking on Yahoo! Pulse data.

As the paper consists of a series of progressive models, experimental results are presented right after each model is developed so as to motivate further enhancements. The paper is structured as follows: we first define the task and briefly reviews existing works in §2; the plain BRI model and learning algorithms are presented in §3; then §4 extends the model to encode sociological principles; §5 tests the proposed models in behavior prediction and §6 concludes.

2. BACKGROUND

2.1 Problem Definition

We consider a typical scenario in social networks where users routinely make decisions while actively interacting with one another. This motivates us to analyze the interplay between an individual user’s behavior of decision making and the tie between users. Suppose we are given a set of users

$$\mathcal{U} = \{1, 2, \dots, N\}$$

and a set of items (e.g., News articles, advertisement, retailing products, movies)

$$\mathcal{I} = \{1, 2, \dots, M\}.$$

The users are connected by a social network represented by the graph $\mathcal{G}(\mathcal{U}, \mathcal{C})$. Here $\mathcal{C} = \{c_{uv}\}$ denotes the set of edges, i.e., $c_{uv} \in \{1, \text{missing}\}$ defines the acquaintance connection between a pair of users $(u, v) \in \mathcal{U}^2$. In the network, users actively make decisions, for example, by clicking links, purchasing products, rating movies. Formally, a decision is a mapping

$$y : \mathcal{U} \times \mathcal{I} \rightarrow \mathcal{Y}.$$

That is: user u makes a decision regarding item i with a response $y_{ui} \in \mathcal{Y}$ (e.g. u rates movie i with a score of y_{ui}). For convenience, hereafter we use u and v to index users, and i and j for items unless stated otherwise.

We are motivated by the following two tasks:

En-signing social relations. A signed social tie is a mapping $s : \mathcal{C} \subset \mathcal{U}^2 \rightarrow \{\pm 1\}$ which characterizes the overall impression (trust/distrust) of user u to user v . In essence, our aim is to *label the edges* of the graph \mathcal{G} (i.e. the acquaintance relationships in a standard social network) *with signed labels* $+1$ or -1 .




Figure 1: An illustration of decision making in social networks: users make decisions with respect to items (e.g. by rating movies); simultaneously users also interact with one another exerting positive or negative social ties.

Behavior prediction. We are also interested in predicting the users’ behavior of decision making. In particular, given a user-item pair (u, i) , we want to predict the response y_{ui} with high precision.

We focus on the more challenging task of sign labeling in §3 and §4, and defer the second goal to §5.

2.2 Related Work

To provide a context of our work, we briefly review existing research on behavior prediction and relation analysis.

Behavior Prediction. Behavior prediction aims at predicting future (unseen) decision responses y_{ui} of a user u w.r.t an item i . We consider collaborative filtering (CF) approaches [12, 27], which tackle the problem by learning from past behavior. A widely used approach to CF learns informative compact latent factors to uncover the dyadic interactions. The basic idea is to associate latent factors with each user and each item, and let the decision be explained by the interactions (e.g., multiplication) of these terms.

While CF approaches gain a lot success in recommender systems, they may be insufficient in the context of social networks where user interactions influence decision making dramatically [10]. For example, instead of making decisions purely by following his own taste as in recommender systems, in a social network, a user can achieve more productive decisions by learning from his friend’s experiences. This aspect has been rapidly increasing in importance as many online services are starting to add a social network aspect to their system (e.g. last.fm, Yelp, Flickr) or directly rent existing social networks (e.g. Amazon, Netflix, Bing). Our recent work [26] exploited Homophily for behavior prediction, where the behavior of two socially-connected users is tied and reinforced by each other. This work further investigates a more realistic case where acquainted users are not necessarily reinforcing each other in decision making. We do so by modeling the signed social ties, a much stronger signal that allows us to tie the behavior of a user to whom he trusts rather than whom he knows.

Unsigned Relation Analysis. Traditional relation analysis [16, 23, 19, 5] focuses on unsigned relations and estimates the tie strength between two users by propagating the observed links through the network topology. For a pair of users (u, v) the observation whether they are connected (i.e.,

know each other) is binary $s_{uv} \in \{1, \text{missing}\}$. Many existing methods employ random walk or spectral algorithms. A random walk on the graph \mathcal{G} is a reversible Markov chain on the vertices \mathcal{U} . Vertices are considered close whenever the hitting time is small or whenever the diffusion probability is large. In some cases spectral smoothness is used to obtain similarity via the associated graph Laplacian [24]. This yields the intrinsic kinship defined by the dominant eigenvectors of the graph [16].

Signed Relation Analysis. Social relations are complex and subtle [15, 6], reflecting interactions that could be either *positive* (e.g. trust, agreement, approval) or *negative* (e.g. distrust, disagreement, opposition). Yet until recently a number of authors began to investigate signed relations by examining the links in emerging signed social networks (e.g., Slashdot, Epinion). [6] applied random walk algorithms to propagate both positive and negative links on Epinion network. [13] conducted similar experiments on Slashdot data. [15] examined the applicability of social psychology theories at multiple signed networks. [14] further applied classification approaches to predict the sign of social relationships. While the goal of [14] is seemingly similar to the sign-labeling task we explore here, this work is still within the scope of link analysis since the focus there is solely on the network topology. Furthermore, they adopted a leave-one-out setting — a rather optimistic scenario where the sign labels for all but one edge are observable. Instead, our focus is not on network topology, but rather on discovering the interplay between relations and behavior. Also, we require no or fewer labels to obtain comparable prediction as [14].

2.3 Data

Before providing a formal description of our approaches let us briefly discuss the data, which we will use to draw intuitions for model development, test proposed models and motivate further enhancements.

Epinion We use data from www.epinion.com for our analysis as it is the only publicly available dataset we are aware of that contains both signed social relationship *and* user behavioral data. Epinion is a well-known knowledge and opinion sharing web site, where users post reviews and assign ratings (on a scale from 1 to 5) to various of items such as retailing products, companies, or movies. More interestingly, the site maintains a signed social network that allows each user to indicate explicitly other users that he is connected to as either positive (trust) or negative (distrust). We apply our models to predict the trust/distrust relationships. The data set contains 132 thousand users, 1560 thousand items, 13.6 million ratings and 850 thousand signed relationships. Both the user-item matrix and the user-user network are very sparse, with densities of only 0.014% and 0.0048% respectively.

Yahoo! Pulse pulse.yahoo.com is an unsigned social network that allows users to communicate with friends and also express their preference for items with explicit indications of “like”. We examine data collected over one year, involving 10^8 users and a large collection of items, such as games, sports, News feeds, finance, travel, shopping, and local information services. Our evaluation focuses on a random subset consisting of about 400 items, 1.2 million users and 29 million “like” indications. Due to the unsigned characteristic of this network, we are not able to use this data set to evaluate




Figure 2: Histograms of behavioral correlations for positively and negatively connected users on epinion.com.

edge-sign prediction. Instead, we will use this data set for behavior prediction evaluation (Section 5).

2.4 Motivating Observations

This paper attempts to predict signed social ties by capturing the interplay between behavioral and social interactions. A key insight is that we consider the behavioral data (e.g. the ratings) to be useful for predicting signed social interactions. This is reasonable as in reality we are often willing to follow the opinions from a trusted friend and fight against those from a foe. As a result, decisions of positively-connected (i.e., trust) friends are more likely to agree, whereas for frenemies the chance of disagreement would be considerably higher. Here we empirically validate this intuition on the Epinion data, as can be seen in Figure 2.

In particular, we quantify behavioral agreement using the Pearson correlation score between the ratings of two users, and we compare this between users with positive and negative links. The Pearson score is defined as follows:

$$\gamma_{uv} := \frac{\sum_{i \in \mathcal{O}[uv]} y_{ui} y_{vi}}{\sqrt{\sum_{i \in \mathcal{O}[uv]} y_{ui}^2 \sum_{i \in \mathcal{O}[uv]} y_{vi}^2}}, \quad (1)$$

where $\mathcal{O}[uv]$ denote the set of items that both u and v rate. Figure 2 plots the distribution of γ_{uv} for positive-relationships (u trusts v) and negative-relationships (u distrusts v). We can see that the behavioral correlation is highly predictive for signed relations. Specifically, the majority (almost 90%) of the positive links indeed show positive behavioral correlations. The case for negative links is a bit more complicated — the distribution peaks at both -1 and +1, indicating that users with distrust relationships could agree or disagree with each other. But even in this case, we can still observe that for almost 40% of the negative edges, the signs of relations are consistent with the signs of behavioral correlations.

3. BEHAVIOR RELATION INTERPLAY

The nontrivial relationship between behavioral correlation and the sign of social relations motivates us to leverage both social and behavioral evidences when predicting signed social ties (and/or behavior vice versa). In this section, we present models to capture the interplay between users' behavior of decision making and social interaction in the context of social networks.




Figure 3: Illustration of behavior-relation interplay. (Left): user u makes a decision w.r.t. item i by either following his own interest in i or taking a friend v 's opinion on i ; (Right): social relation between u and v is a mixture of their prior impressions to each other and the calibration based on their behavioral agreement w.r.t. decision items.

3.1 The BRI Model

Similar to the latent factor CF models for behavior prediction, we devise a latent factor model for user behavior data $\{y_{ui}\}$. We associate latent factors ϕ_u with user u and ψ_i with item i . For notation convenience we require that each latent factor contains a constant component so as to absorb user/item-specific offset into latent factors. In particular, we assume the following model for behavior (The gist of our approach is captured by Figure 3):

Intrinsic taste We assume user u 's interest in item i is quantified by $\phi_u^\top \psi_i$.

Trust credit Similarly, we assume user u trusts an acquainted user v with a credit $\phi_u^\top \phi_v$.

Social decision making Furthermore, we assume that user u makes decision w.r.t. item i by taking a random walk, that is: with probability p he follows his own taste $\phi_u^\top \psi_i$, or with probability $(1 - p)$ he asks one of his friends, say v (with transition probability $\phi_u^\top \phi_v$) for advice, where v 's opinion is captured by $\phi_v^\top \psi_i$. As a results, aspects that are favored/disliked by his friends should impact an individual user's own choices. The decision y_{ui} therefore depends on $p\phi_u^\top \psi_i + (1 - p)\sum_v (\phi_u^\top \phi_v)\phi_v^\top \psi_i$, which is achieved equivalently by "coloring" the inner product between users by a normalized variant of the correlation

$$M_y := (\text{tr } C_y)^{-1} C_y \text{ where } C_y = \sum_v \phi_v \phi_v^\top \quad (2)$$

$$f_{ui} = \phi_u^\top [pI + (1 - p)M_y] \psi_i \quad (3)$$

Here $p \in [0, 1]$ encodes the the probability with which a user's behavior is decided by his own taste.

The same model can be applied to incorporate social relations among users. We assume a user u is fully characterized by its latent factor ϕ_u and devise the following model to characterize the relationships of user u to user v :

Prior impression We assume user u trusts an acquainted user v a priori (i.e., prior to any decision making behavior) with credit $\phi_u^\top \phi_v$.

Behavior agreement User u calibrates his impression regarding v based on the agreement² between his own and v 's decisions, e.g., according to the correlation $\sum_i (\phi_u^\top \psi_i)(\phi_v^\top \psi_i)$.

²Note that a dynamic model, i.e., $h_{uv}(t)$ over time t , is possible if we have access to the time stamp of each activity.

Social tie We assume the social interaction is a mixture of both prior impression and behavior agreement with mixture parameter q , i.e.: the social tie depends on $q\phi_u^\top\phi_v + (1-q)\sum_i(\phi_u^\top\psi_i)(\phi_v^\top\psi_i)$, or equivalently:

$$M_s := (\text{tr } C_s)^{-1} C_s \text{ where } C_s = \sum_i \psi_i \psi_i^\top \quad (4)$$

$$h_{uv} = \phi_u^\top [qI + (1-q)M_s] \phi_v \quad (5)$$

Here $q \in [0, 1]$ defines the mixture probability of prior impression. Unlike in random walk models where proximity in a graph is simply used to smooth secondary estimators of parameters (e.g. reachability, hitting times), we make direct use of it to model the latent variables ϕ s.

Based on the above descriptions, we specify the probabilities with spherical Gaussian distributions (extensions to other distributions are straightforward) and we summarize the overall model in the table below.

The Behavior Relation Interplay (BRI) model.	
$\forall u \in \mathcal{U}$	$\phi_u \sim \mathcal{N}(\phi_u 0, \sigma_u^2 I)$
$\forall i \in \mathcal{I}$	$\psi_i \sim \mathcal{N}(\psi_i 0, \sigma_i^2 I)$
$\forall u \in \mathcal{U}, i \in \mathcal{I}$	$y_{ui} \sim \mathcal{N}(y_{ui} f_{ui}, \sigma_y^2 I)$
$\forall u, v \in \mathcal{U}$	$s_{uv} \sim \mathcal{N}(s_{uv} h_{uv}, \sigma_s^2 I)$

Here $\sigma_{\mathcal{I}}, \sigma_{\mathcal{U}}, \sigma_s, \sigma_y$ are scalars specifying the variance.

We develop inference algorithms for the proposed model, relying on three types of dyadic evidences:

1. The user behavior trace $\{y_{ui}\}$ in the form of user-item interactions such as a user's rating for a movie.
2. The unsigned user-user social connections (e.g., the acquaintance relations at Facebook) $\{c_{uv}\}$, where $c_{uv} = 1$ when u and v know each other or missing otherwise.
Or alternatively whenever available:
3. The signed connections $\{s_{uv}\}$, $s_{uv} = \pm 1$ or missing, e.g., the trust/distrust or friend/foe relations at Epinion and Slashdot.

We refer to c_{uv} as *unlabeled tie* and to s_{uv} as *labeled tie*. Our goal is to infer the labels of social ties using none (i.e. unsupervised) or only a fraction (i.e. semi-supervised) of the label observations by fitting the latent factor model BRI on the above three sources of evidences.

3.2 Unsupervised BRI

We first consider the case where no label is observed. This is a typical setting for the majority of social network sites such as Facebook. Our key idea is to use an absolute-value link function $|\xi|$, which bridges the *hidden* signed-ties we aim to infer, h_{uv} , to the unsigned acquaintance connections we actually observed, c_{uv} . In this way, we enable the mechanism of *automatic sign determination* in the same spirit as the maximum margin clustering of [25] — we assume the objective function achieves optimality at either $h_{uv} = +1$ or $h_{uv} = -1$; through optimization, the model will automatically determine the optimal sign for h_{uv} because different signs correspond to different objective values.

Formulation. Given the observations of user-item behavioral interactions $(u, i) \in \mathcal{O}_y$ and unsigned user-user connections

$(u, v) \in \mathcal{O}_c$, we have the following optimization:

$$\begin{aligned} \underset{\phi, \psi}{\text{minimize}} \quad & \lambda_y \sum_{(u, i) \in \mathcal{O}_y} (y_{ui} - f_{ui})^2 + \lambda_c \sum_{(u, v) \in \mathcal{O}_c} (c_{uv} - |h_{uv}|)^2 \\ & + \lambda_{\mathcal{U}} \sum_{u \in \mathcal{U}} \|\phi_u\|^2 + \lambda_{\mathcal{I}} \sum_{i \in \mathcal{I}} \|\psi_i\|^2 \end{aligned} \quad (6)$$

where $\lambda_{\mathcal{U}}, \lambda_{\mathcal{I}}, \lambda_c, \lambda_y$ are trade-off parameters, f_{ui} and h_{uv} are defined according to Eq(3) and Eq(5).

Prediction. Given an incoming pair of users, (u, v) , we predict the sign of the social tie (i.e., trust or distrust) as follows:

$$\hat{s}_{uv} = \text{sign}(h_{uv} - h_0) \quad (7)$$

Here h_0 is a threshold which suitably determines the fraction of positive and negative ties. Whenever we have access to some labels on a validation set, it is straightforward to determine h_0 via line search: simply sort h_{uv} by magnitude and compute the balanced error for each distinct value of h_{uv} . This is in complete analogy to the line search procedure of [8] for F_β scores.

Optimization. Eq(6) is minimized by stochastic gradient descent. Strictly speaking the objective does not decouple entirely in terms of ϕ and ψ due to the matrices M_y and M_s . Nonetheless we found that an approximation which performs stochastic gradient descent on ϕ and ψ while keeping M_y, M_s fixed produces good results. Also, it can be decomposed into sessions and parallelized using the MapReduce framework.

The stochastic gradient descent procedure is entirely standard: the algorithm processes all observations and updates the parameters in the direction of the negative gradient. For instance, for \mathcal{O}_y the updates are as follows:

$$\phi_u \leftarrow (1 - \eta \bar{\lambda}_{\mathcal{U}}) \phi_u - \eta (f_{ui} - y_{ui})(p \psi_i + (1-p) M_y \psi_i) \quad (8)$$

$$\psi_i \leftarrow (1 - \eta \bar{\lambda}_{\mathcal{I}}) \psi_i - \eta (f_{ui} - y_{ui})(p \phi_u + (1-p) M_y \phi_u) \quad (9)$$

Here $\bar{\lambda}_{\mathcal{U}}$ and $\bar{\lambda}_{\mathcal{I}}$ are regularization constants suitably adjusted with respect to the prediction accuracy (they are obtained by normalizing Eq(6) in terms of size). The learning rate η is annealed with a discount factor after each iteration, as suggested by [12].

Algorithm 1 Stochastic gradient descent learning

initialize $M_y = M_s = I$, random ϕ_u and ψ_i
for $t = 1$ **to** T **do**
 Stochastic gradient descent w.r.t. ϕ_u with M_s fixed.
 Update M_s
 Stochastic gradient descent w.r.t. ψ_i with M_y fixed.
 Update M_y
end for

The absolute-value link function $|\xi|$ is approximated by a smooth variant $(\alpha^{-1} \log 2 \cosh \alpha \xi)$ with $\alpha = 100$. This approximation avoids non-differentiability in the objective function and aids convergence of the optimization.

3.3 Empirical Results for Unsupervised BRI

We conducted experiments on the Epinion data set. The social relationships in Epinion are dominated by positive ones (over 85% of the edges are positive). Following the settings of [14, 6], we derive a balanced dataset containing equal numbers of positive and negative relations by randomly down-sampling the positive edges. Moreover, because




Figure 4: Unsupervised sign prediction performance on Epinion. From left to right: random guess, behavioral correlation based method, the Friendship-Interest Propagation (FIP) model [26], and the proposed BRI model. Results are reported with latent dimensionality $k = 5, 10$ and 50 respectively.

the behavioral data in this dataset is extremely skewed (the mean value of the ratings is around 4.6), we standardize all the ratings to the range $[-1, 1]$ by applying a transformation: $(y - 3.0)^+ - 1.0$.

We evaluate the BRI model in terms of its unsupervised sign prediction accuracy. For this purpose, we hide the signs of all the relationships in the training set, i.e., all the edges are treated as “acquaintance” regardless whether they were originally trust or distrust. For comparison, we also evaluate the following three unsupervised methods:

- **RANDOM.** A random guessing predictor yields 50% accuracy because the dataset is balanced.
- **CORRELATION.** The 1-dimensional classifier:

$$\hat{s}_{uv} = \text{sign}(\gamma_{uv} - \gamma_0)$$

based on the behavioral correlations γ , where γ_0 is the optimal threshold found by line-search.

- **FIP.** The Friendship-Interest Propagation model [26] has been proved effective for link (tie-strength) prediction using both social and behavior data. FIP can be viewed as a special case of BRI with (1) both $p = 1$ and $q = 1$; and (2) no automatic sign determination.

Throughout this paper, all the models are evaluated by using the 1-dimensional prediction scheme Eq(7). In particular, for the output score ρ_{uv} of a model, we predict the sign as $\hat{s}_{uv} = \text{sign}(\rho_{uv} - \rho_0)$, where ρ_0 is an optimal threshold scalar found by line search on validation data.

The prediction accuracy is depicted in Figure 4, where the results of BRI and FIP are reported at latent dimensionality $k = 5, 10$ and 50 respectively. As can be seen from the figure, BRI achieves up to 71% accuracy in this unsupervised setting and it significantly outperform all the three baselines. In particular, BRI gains 42% improvement over RANDOM, 21% over FIP and 11% over CORRELATION. The result is quite promising especially considering that all the signs are hidden in training. It also proves the BRI model is potentially usable in practice for en-signing unsigned social networks, given that we can predict trust/distrust from acquaintance relationships with over 70% accuracy.

One key notion that motivates our work is that we assume the behavioral data should be useful for tie-sign prediction.




Figure 5: Top: Unsupervised sign prediction accuracy as a function of the coloring proportion $p = q$; Bottom: the same statistic for the weight for behavioral data, λ_y .

As can be seen from Figure 4, using the behavioral data alone, the simple correlation-based predictor achieves the second best performance among all the four models. This observation validates our hypothesis and indicates that behavioral information is indeed highly predictive for signed social ties. In Figure 5 (b), we further report the accuracy of BRI as a function of the weight for behavioral data, i.e., λ_y in Eq(6). As we can see, while the best result is achieved at a moderate value of λ_y , interestingly, the performance does not degrade badly if we further increase λ_y . In contrast, if λ_y is decreasing toward 0, the performance degrades quickly towards that of random prediction — in an extreme, BRI trained on unsigned social data alone does not perform substantially better than RANDOM. These observations reveal that, in the unsupervised setting, the behavioral information is even more predictive than the topology of the acquaintance graph.

Note that although the FIP model also leverages the behavioral data, it is ineffective in sign prediction. FIP estimates the unipolar (i.e., in the interval $[0, 1]$) tie-strength indicating whether two users *know each other*. Nonetheless, it is incapable to predict whether they *trust or distrust each other*, i.e. FIP cannot discern the sign of a connection. From another perspective, this comparison between FIP and BRI confirms that the proposed mechanism of *automatic sign determination* is effective for sign prediction.

The performance of the BRI model is affected by the parameters of mixture probabilities, i.e., p and q in Eq(3-5).

Table 1: Semi-supervised sign prediction performance on Epinion data. We compare five methods: random guess (RAND), Friendship-Interest propagation (FIP) [26], behavioral correlation based classifier with optimal threshold (CORR), Leave-one-out Logistic Regression based on local topology features [14] (LOO-LR), and the proposed BRI model trained with different percentages of supervision information (i.e., tie labels).

Model	RAND	FIP	CORR	LOO-LR	BRI									
Labels	0	0	0	All - 1	0	100	1%	5%	10%	20%	30%	50%	70%	90%
Accuracy	0.500	0.591	0.644	0.934	0.709	0.709	0.731	0.747	0.776	0.818	0.836	0.869	0.894	0.912

In Figure 5, we illustrate how the prediction accuracy of BRI changes with these two parameters, where for simplicity p and q are set to the same values. As can be seen from the figure, the accuracy curve is typically in a inverted U-shape with the optimal performance achieved at around 0.2. This observation suggests that both the two groups of factors (i.e. intrinsic taste vs influence, prior impression vs behavior agreement) are important and should be taken into account simultaneously as what we do in the BRI model.

3.4 Semisupervised BRI

We now consider semi-supervised formulation, where part of the labels for social ties are observed — in practice, these labels could be obtained directly from users through an on-line survey or simply by manual annotations. In particular, assume that in addition to the behavioral observations \mathcal{O}_y and the acquaintance relations \mathcal{O}_c , we also have access to a small set of labeled relations (i.e., trust-distrust) on pairs $(u, v) \in \mathcal{O}_s$. In this case we add

$$\lambda_s \sum_{(u,v) \in \mathcal{O}_s} (s_{uv} - h_{uv})^2 \quad (10)$$

to the objective function of Eq(6). This formulation has a close connection with the semi-supervised classification model of [2]. It maximizes the fidelity to the *target output* (e.g. s_{uv}) of the labeled data, while approximating the target output for unlabeled data up to an arbitrary sign. A notable distinction is that we are learning latent factor models (e.g., the BRI model) rather than a global feature mapping (i.e., the classifier of the form $w^\top x + b$ in [2]).

The learning algorithm and prediction formula are similar to the unsupervised case and therefore are omitted.

3.5 Empirical Results for Semisupervised BRI

Table 1 gives the overall results of semi-supervised sign prediction on Epinion data, where we report the performance of the BRI model that are trained using progressive proportions (1% to 90%) of the labeled data. As a reference, we also report the results obtained by RANDOM guessing, CORRELATION-based 1-dimensional classifier, the Friendship-Interest Propagation (FIP) model of [26] and the leave-one-out logistic regression model (denoted LOO-LR) of [14].

Our first observation from the table is that the performance of the BRI model is approaching that of LOO-LR. As we increase the labeled data used for training, we see that the prediction accuracy of the BRI model is steadily improving, indicating that the labeled data can indeed be used to improve the prediction ability of BRI, somewhat as we expected. In particular, with 1% of labeled data, we gain over 3% improvement of accuracy; and if we use over 90% of the labeled data, BRI can achieve comparable per-

formance as the leave-one-out prediction of [14] but using far less supervision information compared to the latter.

Interestingly, it seems that the labeled data are not always helpful for BRI. For example, the model trained using 100 (i.e. less than 0.05%) labeled edges does not perform significantly better (in some cases, even slightly worse) than the unsupervised model. This observation suggests that the amount of the labeled data should exceed a minimum threshold in order for it to take effect. This is quite different from the traditional semi-supervised learning in which labeled data, whenever available, seems to always help prediction. Nonetheless, it is reasonable in our context because we are learning a latent factor model over extremely large user/item spaces — the model is almost inevitably overparameterized and with a lot of local optima due to the non-convexity; the labeled data cannot guarantee to drive the model toward a right direction if there is too little of it.

4. ENCODING SOCIAL PSYCHOLOGY

Theories from social psychology [6] provide profound perspectives for understanding the formation of signed relations in social networks. According to [15, 14], satisfactory sign prediction accuracy could be obtained by applying social psychology principles. For example, [14] constructed local topology-based features according to status and balance theories and employed a logistic regression classifier based on these features to predict the signed relationships. In a leave-one-out setting, their model achieves up to 93% sign prediction accuracy on Epinion data. In this section, we extend the basic BRI models to encode social psychological principles, especially the theories of status and structural balance.

4.1 Encoding Structural Balance Theory

The first theory we try to encode is that of *structural balance* [7, 15]. Roughly speaking it implies the intuition that “a friend’s friend is a friend” and that “an enemy’s enemy is a friend”. Basically, it considers the balance of signs on a triad involving three users. It states that a triad is balanced if and only if it contains an odd number of positive edges.

BALANCE THEORY [7, 15]: *A triad (u, v, w) is balanced in either of the two cases: (1) if $s_{uv} = 1$ and $s_{uw} = 1$, then $s_{vw} = 1$; (2) if $s_{uv} = -1$ and $s_{uw} = -1$, then $s_{vw} = 1$.*

To encode the balance theory, we introduce an additional (1-dimensional) latent factor β_u for each user u . Unlike the latent profiles ϕ , β is only used to capture the structural balance of the network topology. In particular, we assume an additional component:

$$g_{uv} = \beta_u \beta_v, \quad (11)$$

The social tie is modeled using the following mixture-of-

effect model:

$$h_{uv} = q_1 \phi_u^\top \phi_v + q_2 \phi_u^\top M_s \phi_v + q_3 \beta_u \beta_v, \quad (12)$$

where $q_1, q_2, q_3 \in [0, 1]$ and $q_1 + q_2 + q_3 = 1$.

The new mixture component encodes the status theory because it implies: if $g_{uv} > 0$ and $g_{uw} > 0$, then $g_{vw} > 0$ (i.e. a friend’s friend is a friend); if $g_{uv} < 0$ and $g_{uw} < 0$, then $g_{vw} > 0$ (i.e. an enemy’s enemy is a friend). This rank-1 parametrization of structural balance matches the models used in collaborative filtering, albeit with reduced dimensionality. Effectively we are setting aside a subspace of the latent factors exclusively for social interactions.

4.2 Encoding Status Theory

In the previous sections we implicitly viewed the social network as an *undirected* graph. Here, we extend the BRI model to encode status theory [6, 15], which was developed for directed networks. That is, from now on, the social ties are viewed as asymmetric relationships. Basically, status theory assumes that there exists a partial order over the user space such that positive edges only goes from low-status nodes to nodes with higher status; therefore, the relationships are transitive, i.e., if $s_{uv} = +1$ and $s_{vw} = +1$ then $s_{uw} = +1$, and vice versa.

STATUS THEORY[6, 15]: *A positive directed link $s_{uv} = +1$ indicates that the head node u has a higher status than the tail node v ; a negative link $s_{uv} = -1$ indicates that u has a lower status than v .*

To encode status theory we introduce a global user-independent parameter θ to capture the partial ordering of users. θ maps the latent user profile ϕ_u to a scalar quantity $\ell_u = \theta^\top \phi_u$, which reflects the corresponding user u ’s *social status*. According to status theory, we characterize social ties from u to v by modeling the relative status difference between user u and user v :

$$\ell_{uv} = \theta^\top (\phi_u - \phi_v) = \underbrace{\theta^\top \phi_u}_{=: \ell_u} - \underbrace{\theta^\top \phi_v}_{=: \ell_v}. \quad (13)$$

Note that status theory implies that the social tie is typically an antisymmetric relationship, i.e. if $s_{uv} = +1$, then $s_{vu} = -1$. This implication is, however, too strong. For example, in the Epinion case, only 2% of the bi-directional relationships are truly antisymmetric. To relax this assumption, we therefore use a mixture of effects model as in the previous section. This yields:

$$h_{uv} = q_1 \phi_u^\top \phi_v + q_2 \phi_u^\top M_s \phi_v + q_3 (\ell_u - \ell_v). \quad (14)$$

Since ℓ_u, ℓ_v induce a total order on users it satisfies the requirements of social status theory: by transitivity $\ell_{uv} > 0$ and $\ell_{vw} > 0$ imply $\ell_{uw} > 0$.

4.3 Empirical Results

We evaluate the two extended BRI models in the semi-supervised setting and report the results in Tabel 2. As a reference, we also report the performance of the balance model and the status model on its own. As can be seen from the table, by encoding the status theory, we achieve results that are comparable with the leave-one-out prediction of [14], but we use far less supervision, i.e. observed tie labels. Also note that encoding the balance theory only leads to very marginal improvements. The latter is to be

Table 2: Semi-supervised sign prediction performance on Epinion. We compare the following models: Structural balance (BAL), Status (STA), basic BRI model and combinations of these models, denoted by BRI+STA and BRI+BAL respectively.

%Labels	0%	1%	5%	20%	50%	90%
BAL	0.539	0.557	0.582	0.601	0.662	0.735
STA	0.617	0.719	0.757	0.799	0.824	0.843
BRI	0.709	0.731	0.747	0.818	0.869	0.912
BRI+BAL	0.711	0.734	0.751	0.821	0.870	0.912
BRI+STA	0.714	0.739	0.763	0.833	0.884	0.925

expected since adding a small subspace to an already existing multiplicative latent factor model will not change the expressive richness of the model significantly.

Among all the methods, the rank-1 balance model performs the worst, e.g., in the unsupervised case, it performs even much worse than FIP and the correlation-based approach. Somewhat surprisingly, the plain status model perform quite satisfactorily. In most cases, the accuracy scores achieved by the status model are very close to those obtained by the basic BRI model. This observation is consistent with the empirical insight reported in [15] that status theory turns to explain the signed relationship better. However, these results do not entirely invalidate the balance theory because we use a rank-1 model that could be too simplistic for any practical network topology.

It is worth noting that the status model provides a *status mapping* $\ell_u = \theta^\top \phi_u$ that embeds the users into a 1-dimensional space with the coordinates corresponding to users’ *social status*. In Figure 6, we visualize the Epinion network by randomly sampling 0.1% edges as a bipartite graph based on this embedding. In particular, the upper and lower horizontal lines denote the embedding spaces for head nodes and tail nodes respectively; edges that match the status theory (i.e., $s_{uv} = +1$ while $\ell_u > \ell_v$ or $s_{uv} = -1$ while $\ell_u < \ell_v$) are plotted as green lines; in contrast, edges mismatch the theory are depicted as red lines; to give a clear illustration, negative edges are inverted, i.e., $s_{uv} = -1$ are replaced by $s_{vu} = +1$. This figure shows a clear picture about social status and signed relations. Over 84% of signed relations, i.e., trust/distrust, are indeed consistent with the theory, going from higher-status head nodes to lower-status tails. Noticeably, while there are confusions, they are mainly in the middle area between users with relatively small difference of social status. Also note that the confusions are becoming increasingly sparse for high-status head users, indicating that these users are more likely to be trusted by most other users, which is intuitively reasonable. An interesting investigation would be to examine how the social status discovered by this work relates to the well-known concept of *authority* [11] in social networks. We would like to leave such investigation for future research.

5. BEHAVIOR PREDICTION

A direct motivation of this work was to take advantage of signed relationships, when available, for behavior prediction purposes. The reasonable expectation in this context was that a stronger social signal would allow us to obtain better behavior prediction than what can be achieved via

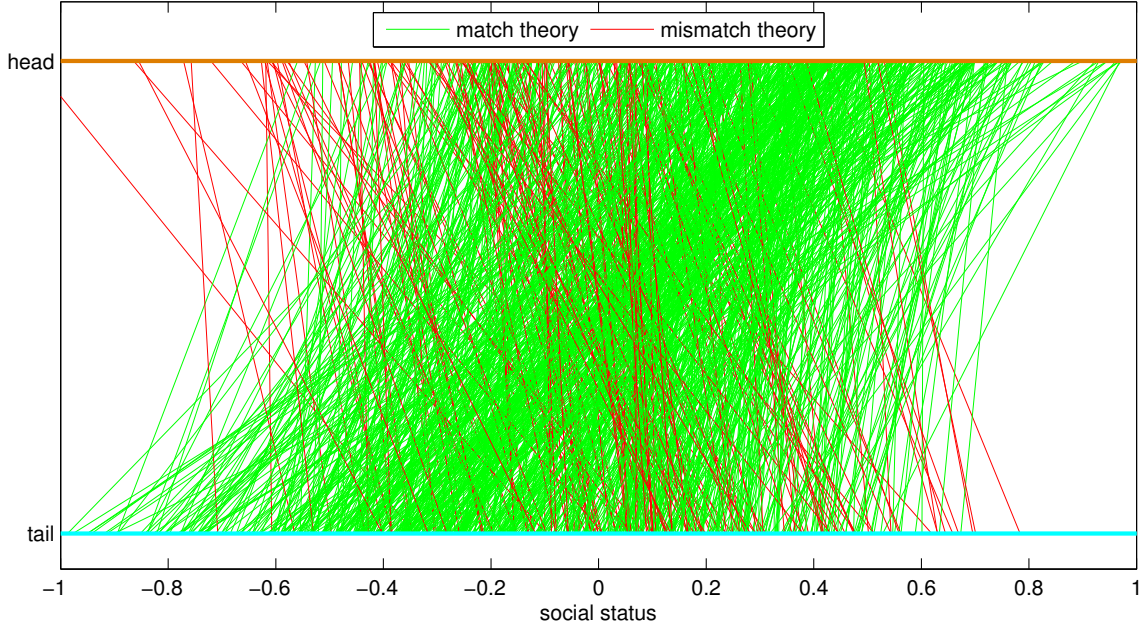


Figure 6: A visualization of social status and signed social ties for epinion.com. Users are embedded into a 1-dimensional space using the status scores ℓ_u as coordinates. Relationships that are consistent with the status theory are denoted as green lines, or otherwise as red lines.

Table 3: Behavior prediction performance on Epinion and Yahoo! Pulse: comparison among collaborative filtering (SVD), the FIP model [26] and our BRI model. Performance is evaluated in terms of rating prediction on Epinion, and top- k ranking on Pulse.

Model	Epinion		Yahoo! Pulse		
	RMSE	MAE	AP@5	AR@5	nDCG@5
SVD	0.556	0.387	0.726	0.208	0.734
FIP	0.553	0.386	0.761	0.225	0.767
BRI	0.524	0.350	0.778	0.232	0.785

unipolar Homophily. In this section, we test whether, even in the absence of explicitly available signed social ties, we are able to improve behavior prediction estimates. For this purpose, we compare three models: a plain collaborative filtering approach (SVD), the Homophily-based model (FIP) [26] and the proposed BRI model. We evaluate these models on both Epinion and Yahoo! Pulse data. For Epinion, we evaluate the rating prediction performance in terms of *root mean square error* (RMSE) and *mean absolute error* (MAE). For Pulse, because the behavior data is binary (either 1 or missing), we evaluate the top- k ranking performance, i.e. we compare the top suggestions of the model with the true actions taken by a user (i.e. whether he *liked* the item). We use the following three metrics commonly used in the IR community:

AP is the *average precision*. AP@ k averages the precision of the top- k ranked list of each query (e.g. user).

AR or *average recall* of the top- k rank list of each query.

nDCG or *normalized Discounted Cumulative Gain* is the normalized position-discounted precision score.

In all three metrics we use $k = 5$ since most social networks use a similar number of items for service recommendation. For the Pulse data, we also implement the bias correction procedure suggested by [26].

The overall results are reported in Table 3. For both data sets we see that the BRI model consistently outperform the two baselines in all the measures. In particular, it improves the RMSE by 5.8% on Epinion, and the nDCG@5 score by 6.9% on Yahoo! Pulse. Note that BRI also significantly outperforms the FIP model. In terms of methodology, both FIP and BRI exploit social relations to improve behavior prediction; the difference, however, lies in how they leverage the social relations: FIP models the *unsigned tie-strength* between users which is then used to *reinforce* the decision making behavior of connected users; in contrast, BRI models the *signed ties*, distinguishing those whom we truly trust from those we do not. In this way, although behavior of two connected users is tied to each other, but not necessarily reinforcing. BRI therefore provides a mechanism that allows social users to follow the opinions from people they trust while at the same time fighting against the viewpoints of their foes. Our results suggest that signed social ties, such as trust-distrust relations, are a substantially more predictive signal in tying and regularizing social behavior than Homophily.

6. SUMMARY

We examined the problem of predicting signed social ties, such as trust and distrust, based on the acquaintance relationships in social networks. This allows us to determine, with surprising accuracy, whether a link corresponds to a trustworthy friend or rather a frenemy. We present models that infer signed ties by capturing the interplay between

social relations and users' behavior of decision making, and extend the models to encode general principles from social psychology. We investigate sign prediction in both unsupervised and semi-supervised settings, conduct experiments in a variety of perspectives and report promising results. We demonstrate that the predicted signed ties are much stronger signals for relating social behavior than traditional Homophily.

Acknowledgements

We thank Su-Lin Wu for helpful discussions. Part of this work was supported by NSF Grants IIS-1049694 and IIS-1116886, and NSFC grant 61129001/F010403.

7. REFERENCES

- [1] B. Awerbuch and R. Kleinberg. Competitive collaborative learning. *Journal of Computer and System Sciences*, 74(8):1271–1288, 2008.
- [2] K. Bennet and A. Demiriz. Semi-supervised support vector machines. In *Advances in Neural Information Processing Systems 11*, pages 368–374, 1998.
- [3] M. J. Brzozowski, T. Hogg, and G. Szabo. Friends and foes: ideological social networking. In *Proceeding of the 26th ACM CHI Conference on Human Factors in Computing Systems*, pages 817–820, 2008. ACM.
- [4] Y. Chen, D. Pavlov, and J. Canny. Large-scale behavioral targeting. In *Proceedings of the 15th ACM Conference on Knowledge Discovery and Data Mining*, pages 209–218. ACM, 2009.
- [5] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Conference on Human factors in computing systems CHI*, pages 211–220, 2009. ACM.
- [6] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *World Wide Web WWW*, pages 403–412, 2004. ACM.
- [7] F. Heider. Attitudes of cognitive organization. *Journal of Psychology*, 21:107–112, 1946.
- [8] T. Joachims. A support vector method for multivariate performance measures. In *Proceedings of the 22nd international conference on Machine learning (ICML '05)*, pages 377–384, 2005. ACM.
- [9] O. Kallenberg. *Probabilistic symmetries and invariance principles*. Springer, 2005.
- [10] T. Kameda, Y. Ohtsubo, and M. Takezawa. Centrality in sociocognitive networks and social influence: An illustration in a group decision-making context. *Journal of Personality and Social Psychology*, 73(2):296–309, August 1997.
- [11] J. M. Kleinberg. Hubs, authorities, and communities. *ACM Computing Surveys*, 31, December 1999.
- [12] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *IEEE Computer*, 42(8):30–37, 2009.
- [13] J. Kunegis, A. Lommatzsch, and C. Bauckhage. The slashdot zoo: mining a social network with negative edges. In *World Wide Web 2009*, pages 741–750. ACM.
- [14] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Predicting positive and negative links in online social networks. In *Conference on World Wide Web WWW*, pages 641–650, New York, NY, USA, 2010. ACM.
- [15] J. Leskovec, D. Huttenlocher, and J. Kleinberg. Signed networks in social media. In *Human factors in computing systems CHI*, pages 1361–1370, New York, NY, USA, 2010. ACM.
- [16] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *Conference on Knowledge and Information Management CKIM*, pages 556–559, New York, NY, USA, 2003. ACM.
- [17] H. Ma, M. Lyu, and I. King. Learning to recommend with trust and distrust relationships. In *ACM conference on Recommender Systems RecSys*, pages 189–196, 2009. ACM.
- [18] P. Massa and P. Avesani. Trust-aware recommender systems. In *ACM conference on Recommender Systems RecSys*, pages 17–24, 2007. ACM.
- [19] M. Newman. Detecting community structure in networks. *The European Physical Journal B - Condensed Matter and Complex Systems*, 38(2):321–330, 2004.
- [20] M. McPherson, L. S. Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [21] J. O'Donovan and B. Smyth. Trust in recommender systems. In *Intelligent User Interfaces IUI*, pages 167–174, 2005. ACM.
- [22] J. P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A. L. Barabási. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences*, 104(18):7332–7336, 2007.
- [23] M. Rosvall and C. T. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123, 2008.
- [24] A. J. Smola and R. Kondor. Kernels and regularization on graphs. In *Proceeding of Annual Conference on Computational Learning Theory*, pages 144–158, 2003.
- [25] L. Xu, J. Neufeld, B. Larson, and D. Schuurmans. Maximum margin clustering. In *Advances in Neural Information Processing Systems 17*. MIT Press, 2005.
- [26] S.-H. Yang, B. Long, A. Smola, N. Sadagopan, Z. Zheng, and H. Zha. Like like alike – joint friendship and interest propagation in social networks. In *WWW '11: Proceedings of the 20th international conference on World wide web*, 2011. ACM.
- [27] S.-H. Yang, B. Long, A. Smola, H. Zha, and Z. Zheng. Collaborative-competitive filtering: learning recommender using context of user choice In *SIGIR '11: Proceedings of the 34th ACM conference on Research and development in Information Retrieval*, 2011. ACM.