

Actively Learning to Infer Social Ties

Honglei Zhuang · Jie Tang · Wenbin Tang · Tiancheng Lou · Alvin Chin · Xia Wang

Received: date / Accepted: date

Abstract We study the extent to which social ties between people can be inferred in large social network, in particular via active user interactions. In most online social networks, relationships are lack of meaning labels (e.g., “colleague” and “intimate friends”) due to various reasons. Understanding the formation of different types of social relationships can provide us insights into the micro-level dynamics of the social network. In this work, we precisely define the problem of inferring social ties and propose a Partially-Labeled Pairwise Factor Graph Model (PLP-FGM) for learning to infer the type of social relationships. The model formalizes the problem of inferring social ties into a flexible semi-supervised framework. We test the model on three different genres of data sets and demonstrate that it is possible to accurately infer 85% of social relationships. We further study how to maximally leverage user interactions to help improve the learning accuracy. Two active learning based algorithms are proposed to actively select relationships to query users for their labels. Experimental results show that with only a few user corrections, the accuracy of inferring social ties can be significantly improved. Finally, to scale the model to handle real large networks, a distributed learning algorithm has been developed.

The work is supported by the Natural Science Foundation of China (No. 61073073, No. 60973102, No. 61170061), Chinese National Key Foundation Research (No. 60933013, No.61035004), and National Basic Research Program of China (No. 2011CB302302). It is also supported by a research funding from Nokia Research Center.

Honglei Zhuang · Jie Tang · Wenbin Tang
Department of Computer Science and Technology, Tsinghua University, Beijing, China
Tel.: +86-10-62794365
Fax: +86-10-62794365
E-mail: zh109@mails.tsinghua.edu.cn

Tiancheng Lou
Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China

Alvin Chin · Xia Wang
Nokia Research Center Beijing, China.

1 Introduction

Online social networks, such as Facebook, MySpace, Twitter, and FourSquare, have already become a bridge to connect our real daily life and the virtual web space. Facebook, one of the largest social networks, has more than 750 million active users in July 2011; Foursquare, a location-based mobile social network, has attracted 10 million registered users in the middle of 2011. Just to mention a few, there is little doubt that most of our friends are online now. Considerable research has been conducted on social network analysis [1, 10, 28, 34], dynamic evolution analysis [19], social influence analysis [8, 17, 37], social behavior analysis [30, 35], and social tie analysis [6, 15, 23, 36]. However, most of these works ignore one important fact that makes the online social networks very different from the physical social networks, i.e., our physical social networks are colorful (“family members”, “colleagues”, and “classmates”) but the online social networks are still black-and-white: the users merely do not take the time to label the relationships. Indeed, statistics show that only 16% of mobile phone users in Europe have created custom contact groups [30, 13] and less than 23% connections on LinkedIn have been labeled. Understanding the formation of different types of social relationships can provide us insights into the micro-level dynamics of the social network. For example, awareness of the types of social relationships can offer abundantly additional information for many mining application such as friend recommendation and product advertisement.

In this work, we investigate to what extent social relationships can be inferred from the online social networks. Given users’ behavior history and interactions between users, can we estimate how likely they are to be family members? The first challenge is how to design a unified model so that it can be easily applied to different domains? There exist a few related studies. For example, Diehl et al. [7] try to identify the relationships by learning a ranking function. Wang et al. [41] propose an unsupervised algorithm for mining the advisor-advisee relationships from the publication network. However, [7] only considers the communication archive; [41] is a domain-specific unsupervised algorithm. Both algorithms are not easy to extend to other domains.

Another challenge is that online social networks are becoming more and more complex and dynamic. Even the best performance achieved by the state-of-the-art algorithms is still under 90%. The result is unsatisfactory and invariably contains a number of errors. A promising solution is to design an interactive interface to allow users to provide feedbacks on the inferring results. However, we should be aware that the interactive process might be tedious, error-prone, and time-consuming. For example, for inferring advisor-advisee relationships from the coauthor network, an author may have hundreds of coauthors. The user may soon become tired, if she/he is asked to carefully go through all her/his relationships to validate the inferring results. Ideally, an algorithm should be able to actively select only a few potentially wrong relationships to query the user, instead of passively waiting for user feedbacks. The problem is referred to as *actively learning to infer social ties*.

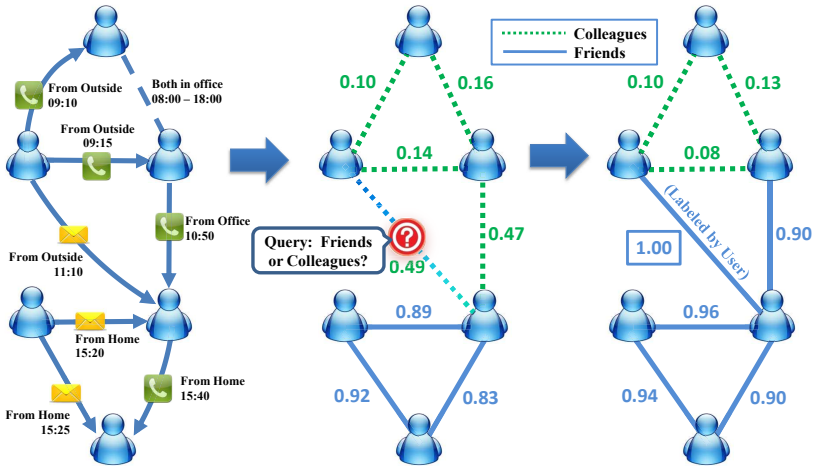


Fig. 1. An example of actively learning to infer social relationships in a mobile communication network. The left figure is the input of our problem, the middle figure shows the inferred relationships by the proposed learning model and the relationship selected by an active learning algorithm to query the user. The right figure is the improved result with the user's feedback.

Motivating Examples To illustrate the problem, Figure 1 gives an example of actively inferring friendships in a mobile communication network. The left figure gives the input of our problem: the mobile social network, which consists of users, calls made and messages sent between users, and users' attribute information such as location. The objective is to classify the type of social relationships in the network. The middle figure shows the result of the proposed PLP-FGM model, a semi-supervised learning model. The blue solid lines stand for friend relationship between users and the green dash lines indicate colleagues. The probability associated with each relationship represents how confident the learning model is in the inferred type of the relationship. Further, an active learning algorithm selects a uncertain relationship (highlighted with the question mark) to query the user. Once the user gives the answer, the learning model propagates the correction in the social network and further corrects the other relationships (Cf. the right figure).

Therefore, the fundamental problem is how to design a flexible model for effectively and efficiently learning to infer social ties in different networks. This problem is non-trivial and poses a set of unique challenges. First, what are the underlying factors that can determine a specific type of social relationship. Second, the input social network is partially labeled. We may have some labeled relationships, but most of the relationships are unknown. To learn a high-quality predictive model, we should not only consider the knowledge provided by the labeled relationships, but also leverage the unlabeled network information. Third, how to make optimal use of user interaction. The selection should consider both the uncertainty and the network structure information. Finally, real social networks are getting bigger with thousands even millions

of nodes. It is important to develop a method that can scale well to real large networks.

Results In this paper, we try to conduct a systematic investigation of the problem of actively learning to infer social ties in large networks. We precisely define the problem and propose a Partially-Labeled Pairwise Factor Graph Model (PLP-FGM) for solving this problem. To make optimal use of user interactions, two strategies, an influence maximization based strategy and a belief maximization based strategy, have been devised to actively select potentially wrong but most useful relationships to query the user. To scale to large networks, we have developed a distributed implementation of the learning algorithm based on MPI (Message-Passing Interface).

We evaluate the proposed model on three different data sets: Publication, Email, Mobile network. Experimental results demonstrate that the proposed PLP-FGM model can accurately infer 92.7% of advisor-advisee relationships from the coauthor network (Publication), 88.0% of manager-subordinate relationships from the email network (Email), and 83.1% of the friendships from the mobile network (Mobile). Our study also reveals several interesting phenomena:

- *Unlabeled data indeed helps.* With the unlabeled relationships, the accuracy of inferring social relationships can be significantly improved (+2.2-+11.8%).
- *Strong correlation between relationships.* For example, on the Email data set, if we consider three types of correlations (Co-recipient, Co-manager, Co-subordinate), we can obtain an 7.6% improvement on the inferring accuracy.
- *Network information helps.* For active selection of the relationship, both of the proposed methods which consider the network information outperform (+0.3-6.1%) the alternative baseline methods which consider only the uncertain information.

Organization The rest of paper is organized as follows. Section 2 formally formulates the problem and presents the basic idea of our approaches; Section 3 explains the PLP-FGM model; Section 4 describes several active learning strategies for PLP-FGM and gives a theoretical analysis; Section 5 gives experimental results; Finally, Section 6 discusses related work and Section 7 concludes.

2 Overview

In this section, we present the problem formulation and the basic ideas of our approach.

2.1 Problem Formulation

First, we give necessary definitions of related concepts. A social network can be represented as $G = (V, E)$, where V is the set of $|V| = N$ users and $E \subset V \times V$ is the set of $|E| = M$ relationships between users. The objective of our work is to learn a model that can effectively infer the type of social relationships between two users. More precisely, we first formally define the output of our problem, namely *relationship semantics*.

Definition 1 Relationship semantics: Relationship semantics is a triple (e_{ij}, r_{ij}, p_{ij}) , where $e_{ij} \in E$ is a social relationship; $r_{ij} \in \mathcal{Y}$ is a label associated with the relationship; \mathcal{Y} is the set of all the labels; p_{ij} is the probability (confidence) obtained by an algorithm for inferring relationship type.

Social relationships might be undirected in some networks (e.g., the friendship discovered from the mobile calling network) or directed in other networks (e.g., the advisor-advisee relationship in the publication network). To be consistent, we define all social relationships as directed relationships. In addition, relationships may be static (e.g., the family-member relationship) or dynamic over time (e.g., colleague relationship). In this work, we focus on static relationships, and leave the dynamic case to our future work.

To infer relationship semantics, we could consider different factors such as user-specific information, link-specific information, and global constraints (Cf. § 5.1 and § 7 for examples). For example, to discover advisor-advisee relationships from a publication network, we can consider how many papers were coauthored by two authors; how many papers in total an author has published; when the first paper was published by each author. Besides, there may already exist some labeled relationships. Formally, we can define the input of our problem, a partially labeled network.

Definition 2 Partially labeled network: A partially labeled network is an augmented social network denoted as $G = (V, E^L, E^U, R^L, \mathbf{W})$, where E^L is a set of labeled relationships and E^U is a set of unlabeled relationships with $E^L \cup E^U = E$; R^L is a set of labels corresponding to the relationships in E^L ; \mathbf{W} is an attribute matrix associated with users in V where each row corresponds to a user, each column an attribute, and an element w_{ij} denotes the value of the j^{th} attribute of user v_i .

Based on the above concepts, we can define the problem of inferring social relationships. Given a partially labeled network, the goal is to detect the types (labels) of all unknown relationships in the network. More precisely,

Problem 1 Social relationship mining. Given a partially labeled network $G = (V, E^L, E^U, R^L, \mathbf{W})$, the objective is to learn a predictive function

$$f : G = (V, E^L, E^U, R^L, \mathbf{W}) \rightarrow R$$

Another important question is how we can learn the mapping function f effectively. In many situations, labeled data is limited and expensive. The

problem is, can we design a strategy to *actively* learn the model with minimal labeling cost? Formally,

Problem 2 Active relationship mining. Given a partially labeled network $G = (V, E^L, E^U, R^L, \mathbf{W})$, and a labeling budget b (number of user interactions). Our objective is to select a subset of unknown relationships $A \subset E^U$ within the constraint of b to label, so that the performance of predictive function f can be maximally improved.

Our formulation of inferring social relationships is very different from existing works on relation mining [4], which focuses on detecting the relationships from the content information, while we focus on mining relationship semantics in social networks. Diehl et al.[7] and Wang et al.[41] investigate the problem of relationship identification. However, they study the problem in specific domains (Email network or Publication network). Backstrom et al.[2] propose an algorithm based on supervised random walks for link prediction. Crandall et al.[6] incorporate geographic coincidences to infer social ties, while Wang et al.[42] consider user mobility and network proximity. Different from these works which aim at link prediction, our goal is to infer the types of relationships. There are also works on inferring the types of relationships. Hopcroft et al.[15] explore the problem of reciprocal relationship prediction and Tang et al. [36] have developed a framework for classifying the type of social relationships by learning across heterogeneous networks. Yang et al.[43] study the retweeting behavior. Leskovec et al.[23] focus on the prediction of edge signs (positive or negative). However, they do not consider how to make optimal use of user interaction.

2.2 Our Approach

For inferring the type of social relationships, we have three basic intuitions. First, the user-specific or link-specific attributes will contain implicit information about the relationships. For example, two users who make a number of calls in working hours might be colleagues; while two users who frequently contact with each other in the evening are more likely to be family members or intimate friends. Second, relationships among different users may have a correlation. For example, in the mobile network, if user v_i makes a call to user v_j immediately after calling user v_k , then user v_i may have a similar relationship (family member or colleague) with user v_j and user v_k . Third, we also need to consider some global constraints such as common knowledge or user-specific constraints.

Based on the intuitions above, we propose a Partially-Labeled Pairwise Factor Graph Model (PLP-FGM). It allows us to take all the factors mentioned above into account to better infer the social relationships. We will describe the model in details in Section 3.

For actively selecting helpful relationships to query the user, we define a quality function $Q(A)$, which measures the expected improvement of the

prediction performance by labeling relationships in set A . The problem can be then defined as an optimization problem of $Q(A)$, i.e.,

$$A^* = \arg \max_{A \subset Y^U} Q(A), |A| = b, b > 0$$

To quantify $Q(A)$, we could consider how a selected node can influence the others. For example, correction of a centered relationship may trigger a spread of the correction, thus help infer correlated relationships.

Based on the above intuitions, we develop an Influence-Maximization Selection (IMS) model and a Belief-Maximization Selection (BMS) model for actively inferring the types of social relationships. The IMS model selects the most influential nodes, by leveraging the network structure and the uncertainty obtained from PLP-FGM. The BMS model further incorporates the active selection process into the learning process of PLP-FGM.

3 Partially-Labeled Pairwise Factor Graph Model (PLP-FGM)

Typically, there are two ways to model the relationship inferring problem. The first way is to model each user as a node and for each node we estimate probability distributions of different relationships from the user to her/his neighbors in the social network. The graphical model consists of N variable nodes. Each node contains a $d \times |\mathcal{Y}|$ matrix to represent the probability distributions of different relationships between the user and her/his neighbors, where d is the number of neighbors of the node. This model is intuitive, but it suffers from some limitations. For example, it is difficult to model the correlations between two relationships, and its computational complexity is high. An alternative way is to model each relationship as a node in the graphical model and the relationship mining task becomes how to predict the semantic label for each relationship node in the model. This model contains M nodes ($2M$ when the input social network is undirected). This model is able to incorporate different correlations between relationships such as the above intuitions.

We propose a Partially-Labeled Pairwise Factor Graph Model (PLP-FGM). Figure 2 shows the graphical representation of the PLP-FGM. Each relationship (v_{i_1}, v_{i_2}) or $e_{i_1 i_2}$ in partially labeled network G is mapped to a *relationship node* r_i in PLP-FGM. We denote the set of relationship nodes as $Y = \{y_1, y_2, \dots, y_M\}$. The relationships in G are partially labeled, thus all nodes in PLP-FGM can be divided into two subsets Y^L and Y^U , corresponding to the labeled and unlabeled relationships respectively. For each relationship node $y_i = (v_{i_1}, v_{i_2}, r_{i_1 i_2})$, we combine the attributes $\{\mathbf{w}_{i_1}, \mathbf{w}_{i_2}\}$ into a *relationship attribute vector* \mathbf{x}_i .

Now we explain the PLP-FGM in details. The relationships in the input are modeled by relationship nodes in PLP-FGM. Corresponding to the three intuitions, we define the following three factors.

- *Attribute factor*: $f(y_i, \mathbf{x}_i)$ represents the posterior probability of the relationship y_i given the attribute vector \mathbf{x}_i ;

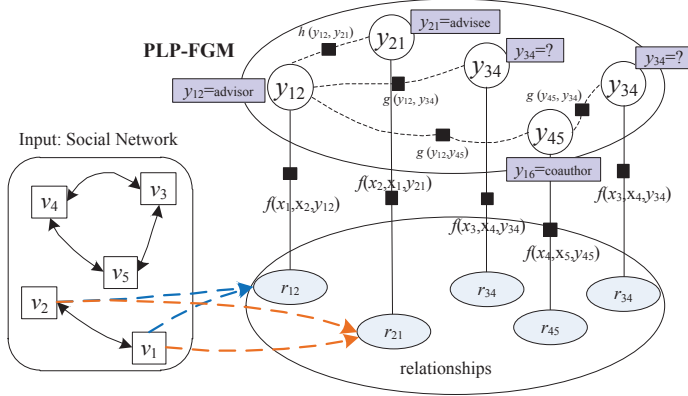


Fig. 2. Graphical representation of the PLP-FGM model.

- *Correlation factor*: $g(y_i, G(y_i))$ denotes the correlation between the relationships, where $G(y_i)$ is the set of correlated relationships to y_i .
- *Constraint factor*: $h(y_i, H(y_i))$ reflects the constraints between relationships, where $H(y_i)$ is the set of relationships constrained on y_i .

Given a partially labeled network $G = (V, E^L, E^U, R^L, \mathbf{W})$, we can define the joint distribution over Y as

$$p(Y|G) = \prod_i f(y_i, \mathbf{x}_i) g(y_i, G(y_i)) h(y_i, H(y_i)) \quad (1)$$

The three factors can be instantiated in different ways. In this paper, we use exponential-linear functions. In particular, we define the attribute factor as

$$f(y_i, \mathbf{x}_i) = \frac{1}{Z_\lambda} \exp\{\lambda^T \Phi(y_i, \mathbf{x}_i)\} \quad (2)$$

where λ is a weighting vector and Φ is a vector of feature functions. Similarly, we define the correlation factor and constraint factor as

$$g(y_i, G(y_i)) = \frac{1}{Z_\alpha} \exp\left\{ \sum_{y_j \in G(y_i)} \alpha^T \mathbf{g}(y_i, y_j) \right\} \quad (3)$$

$$h(y_i, H(y_i)) = \frac{1}{Z_\beta} \exp\left\{ \sum_{y_j \in H(y_i)} \beta^T \mathbf{h}(y_i, y_j) \right\} \quad (4)$$

where \mathbf{g} and \mathbf{h} can be defined as a vector of indicator functions. This feature definition is often used in graphical models such as Markov Random Fields[14] or Conditional Random Fields[22].

Model Learning Learning PLP-FGM is to estimate a parameter configuration $\theta = (\lambda, \alpha, \beta)$, so that the log-likelihood of observation information (labeled relationships) are maximized. For presentation simplicity, we concatenate all factor functions for a relationship node y_i as $\mathbf{s}(y_i) =$

$(\Phi(y_i, \mathbf{x}_i)^T, \sum_{y_j} \mathbf{g}(y_i, y_j)^T, \sum_{y_j} \mathbf{h}(y_i, y_j)^T)^T$. The joint probability defined in (Eq. 1) can be rewritten as

$$p(Y|G) = \frac{1}{Z} \prod_i \exp\{\theta^T \mathbf{s}(y_i)\} = \frac{1}{Z} \exp\{\theta^T \sum_i \mathbf{s}(y_i)\} = \frac{1}{Z} \exp\{\theta^T \mathbf{S}\} \quad (5)$$

where $Z = Z_\lambda Z_\alpha Z_\beta$ is a normalization factor (also called partition function), \mathbf{S} is the aggregation of factor functions over all relationship nodes, i.e., $\mathbf{S} = \sum_i \mathbf{s}(y_i)$.

One challenge for learning the PLP-FGM model is that the input data is partially labeled. To calculate the partition function Z , one needs to sum up the likelihood of possible states for all nodes including unlabeled nodes. To deal with this, we use the labeled data to infer the unknown labels. Here $Y|Y^L$ denotes a labeling configuration Y inferred from the known labels. Thus, we can define the following log-likelihood objective function $\mathcal{O}(\theta)$:

$$\begin{aligned} \mathcal{O}(\theta) &= \log p(Y^L|G) = \log \sum_{Y|Y^L} \frac{1}{Z} \exp\{\theta^T \mathbf{S}\} \\ &= \log \sum_{Y|Y^L} \exp\{\theta^T \mathbf{S}\} - \log Z \\ &= \log \sum_{Y|Y^L} \exp\{\theta^T \mathbf{S}\} - \log \sum_Y \exp\{\theta^T \mathbf{S}\} \end{aligned} \quad (6)$$

To solve the objective function, we can consider a gradient decent method (or a Newton-Raphson method). Specifically, we first calculate the gradient for each parameter θ :

$$\begin{aligned} \frac{\partial \mathcal{O}(\theta)}{\partial \theta} &= \frac{\partial (\log \sum_{Y|Y^L} \exp \theta^T \mathbf{S} - \log \sum_Y \exp \theta^T \mathbf{S})}{\partial \theta} \\ &= \frac{\sum_{Y|Y^L} \exp \theta^T \mathbf{S} \cdot \mathbf{S}}{\sum_{Y|Y^L} \exp \theta^T \mathbf{S}} - \frac{\sum_Y \exp \theta^T \mathbf{S} \cdot \mathbf{S}}{\sum_Y \exp \theta^T \mathbf{S}} \\ &= \mathbb{E}_{p_\theta(Y|Y^L, G)} \mathbf{S} - \mathbb{E}_{p_\theta(Y|G)} \mathbf{S} \end{aligned} \quad (7)$$

Another challenge here is that the graphical structure in PLP-FGM can be arbitrary and may contain cycles, which makes it intractable to directly calculate the expectation $\mathbb{E}_{p_\theta(Y|G)} \mathbf{S}$. A number of approximate algorithms have been proposed, such as Loopy Belief Propagation (LBP) [27]. In this paper, we utilize Loopy Belief Propagation. Specifically, we approximate marginal probabilities $p(y_i|\theta)$ and $p(y_i, y_j|\theta)$ using LBP. With the marginal probabilities, the gradient can be obtained by summing over all relationship nodes. It is worth noting that we need to perform the LBP process twice in each iteration, one time for estimating the marginal probability $p(y|G)$ and the other for $p(y|Y^L, G)$. Finally with the gradient, we update each parameter with a learning rate η . The learning algorithm is summarized in Algorithm 1.

Inferring Unknown Social Ties We now turn to describe how to infer the type of unknown social relationships. Based on learned parameters θ , we can

<p>Input: learning rate η Output: learned parameters θ</p> <p>Initialize θ;</p> <p>repeat</p> <div style="margin-left: 20px;"> <p>Calculate $\mathbb{E}_{p_\theta(Y Y^L, G)} \mathbf{S}$ using LBP ;</p> <p>Calculate $\mathbb{E}_{p_\theta(Y G)} \mathbf{S}$ using LBP ;</p> <p>Calculate the gradient of θ according to Eq. 7:</p> $\nabla_\theta = \mathbb{E}_{p_\theta(Y Y^L, G)} \mathbf{S} - \mathbb{E}_{p_\theta(Y G)} \mathbf{S}$ <p>Update parameter θ with the learning rate η:</p> $\theta_{\text{new}} = \theta_{\text{old}} - \eta \cdot \nabla_\theta$ </div> <p>until <i>Convergence</i>;</p>

Algorithm 1: Learning PLP-FGM.

predict the label of each relationship by finding a label configuration which maximizes the joint probability (Eq. 1), i.e.,

$$Y^* = \operatorname{argmax}_{Y|Y^L} p(Y|G) \tag{8}$$

Again, we utilize the Loopy Belief Propagation to compute the marginal probability of each relationship node $p(y_i|Y^L, G)$ and then predict the type of a relationship as the label with the largest marginal probability. The marginal probability is then taken as the prediction confidence.

Time Complexity Analysis We use ν_1, ν_2, ν_3 to denote the number of attribute factors, correlation factors and constraints factors in our PLP-FGM respectively. In each round of LBP, the time cost of propagation is $\mathcal{O}(\nu_1 \cdot \dim(\Phi) + \nu_2 \cdot \dim(\mathbf{g}) + \nu_3 \cdot \dim(\mathbf{h}))$, where $\dim(\cdot)$ is the dimension of a vector. We execute the learning algorithm for n iterations, and in each round we execute LBP for n_{LBP} iterations. Thus we can estimate the time complexity as $\mathcal{O}((\nu_1 \cdot \dim(\Phi) + \nu_2 \cdot \dim(\mathbf{g}) + \nu_3 \cdot \dim(\mathbf{h})) \times n \times n_{LBP})$.

4 Actively Learning PLP-FGM

The quality function $Q(A)$ can be defined in different forms. Without any constraints, optimizing the quality function $Q(A)$ needs to enumerate all possible subsets $A \subset Y^U$, which is obviously NP-hard. Let us first review two baseline greedy algorithms; and then then we will present two new algorithms: Influence-Maximization Selection model (IMS) and Belief-Maximization Selection model (BMS); a theoretical analysis will be given.

4.1 Baseline Methods

Maximum Uncertainty(MU) A most common selection strategy for active learning is to select the most uncertain relationships. The uncertainty of an

unlabeled relationship y_i is measured by the *entropy* $H(y_i) = -\sum_{y \in \mathcal{Y}} p(y_i = y) \log p(y_i = y)$. Based on this intuition, we can define the quality function as

$$Q_{MU}(A) = H(A) \quad (9)$$

where $H(A) = \sum_{y_i \in A} H(y_i)$.

Information Density (ID) A drawback of the Maximum Uncertainty strategy is its tendency to choose outliers. Thus we employ another strategy, Information Density, proposed in [32]. The idea is to choose the most *representative* nodes in Y^U , which are supposed to be the most informative ones. Based on this intuition, we measure the informativeness of a node by its cosine similarity to all other unlabeled nodes in the sense of the attributes attached to a node. Formally, we define the quality function as

$$Q_{ID}(A) = \sum_{i \in A} H(y_i) \times \left[\frac{1}{|Y^U|} \sum_{j \in Y^U} \text{sim}(\mathbf{x}_i, \mathbf{x}_j) \right] \quad (10)$$

where $\text{sim}(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{x}_i \cdot \mathbf{x}_j}{\|\mathbf{x}_i\| \times \|\mathbf{x}_j\|}$. Note that we again employ the entropy of a relationship node $H(y_i)$ to leverage the “base” informativeness.

4.2 Proposed Methods

Influence-Maximization Selection (IMS) All the strategies mentioned above do not consider the network structure information. As relationship nodes in PLP-FGM are correlated, the most influential nodes are more likely to help improve the overall performance of the model. Existing work has studied several influence propagation model, including the Linear Threshold Model in [17]. The LTS set a threshold value ε_i for each node, and weights $b_{i,j}$ for its edges, satisfying $\sum_{j \in NB(i)} b_{i,j} \leq 1$. In each time stamp, if $\sum_{j \in NB(i) \wedge \text{activated}(j)} b_{i,j} \geq \varepsilon_i$, then the node i will be activated. We develop a variation of the LTS by incorporating a score for each node reflecting the strength of the influence passing in our model. The propagation process is described as:

- The graph is the same as the PLP-FGM model. In addition, we call a relationship node as “activated” when its label y_i is determined. The initial activated set of nodes is Y^L . We assign a threshold $\varepsilon_i = \sum_{y \in \mathcal{Y}} |p(y_i = y|G, Y^L) - \frac{1}{|\mathcal{Y}|}|$ for each node. Thus a node with higher uncertainty will be easier to be activated.
- When a node i is activated, it spreads its gained score increment ($g_i - \varepsilon_i$) to its neighbor nodes $j \in NB(i)$ with a weight $b_{i,j}$, i.e. $g_j \leftarrow g_j + b_{i,j}(g_i - \varepsilon_i)$. To simplify the problem, we set weight $b_{i,j} = 1/|NB(j)|$.

- If a node is labeled by the user, we set it as activated and assign its gained score as 1. The gained score for other nodes is set to 0 at the beginning. Once an inactivated node k gains a score which exceeds the threshold, i.e. $g_k > \varepsilon_i$, it will become activated and spread its gained score similarly. An activated node only spreads its gained score once and remains its status.

We define the quality function $Q_{IMS}(A)$ as the total number of activated nodes after the propagation process. Finding the set A that maximizes the quality function $Q_{IMS}(A)$ is NP-hard. Similar to [17], in this paper, we use a greedy strategy to approximate the solution. Note that unlike the LTS, we do not guarantee a lower bound of error for the greedy optimization method.

Belief-Maximization Selection (BMS) To quantify the influence of one node on the others, we employ the *belief* of each node obtained by Loopy Belief Propagation in our model. We define a heuristic by removing the effect of attributes from the belief score, denoted by $\mathcal{B}(y_i|G, Y^L)$. More precisely,

$$\mathcal{B}(y_i|G, Y^L) = \exp\{\theta^T \mathbf{s}(y_i) - \lambda^T \Phi(y_i, \mathbf{x}_i)\}$$

By normalizing the belief of one relationship node, we obtain the *belief marginal probability*.

$$p_{\mathcal{B}}(y_i|G, Y^L) = \frac{1}{Z_{\mathcal{B}}} \mathcal{B}(y_i|G, Y^L)$$

where $Z_{\mathcal{B}}$ is the normalization factor. It estimates the marginal probability distribution of a relationship node where the information of its attribute vector is absent.

A basic intuition is, the belief of a relationship node is *monotonically increasing* with respect to the number of relationship nodes of the same type, i.e., $\mathcal{B}(y_i = y|G, Y^L)$ is monotonically increasing with respect to the number of relationships with label y . Without loss of generality, we first consider the binary relationship mining problem, i.e., there are only two possible labels of relationships ($\mathcal{Y} = \{0, 1\}$). In the binary setting, we further consider the active selection for each type separately. This is because when mixing the different types of relationships together, it cannot be guaranteed to have a closed-form solution. Thus, when users provide only positive feedback, our objective is to find a set of positive nodes. Accordingly, we define the quality function of the *positive-oriented BMS strategy* as:

$$Q_{BMS^+}(A) = \sum_{y_i \in Y_{(1)}^U} p_{\mathcal{B}}(y_i = 1|G, Y^L \cup A) \quad (11)$$

where $Y_{(1)}^U = \{y_i | y_i \in Y^U \wedge \mathcal{B}(y_i = 1|G, Y^L) \geq \mathcal{B}(y_i = 0|G, Y^L)\}$.

Symmetrically, if the users provide only *negative* feedback, we can adopt a *negative-oriented BMS strategy*, with the following quality function:

Input: G, b
Output: a set of selected relationships A

Train PLP-FGM and get the parameter configuration θ ;
 $A^+, A^- \leftarrow \emptyset$;
for $b/2$ *times* **do**
 Use Loopy Belief Propagation(LBP) to obtain the probability distribution for each relationship;
 Find
 $y_{max}^+ = \operatorname{argmax}_{y_i \in Y^U} p(y_i = 1 | G, Y^L)(Q_{BMS^+}(A^+ \cup y_i) - Q_{BMS^+}(A^+))$;
 Move y_{max}^+ from Y^U to A^+ ;
 $Y^L \leftarrow Y^L \cup A^+$;
 Use Loopy Belief Propagation(LBP) to obtain the probability distribution for each relationship;
 Find
 $y_{max}^- = \operatorname{argmax}_{y_i \in Y^U} p(y_i = 0 | G, Y^L)(Q_{BMS^-}(A^- \cup y_i) - Q_{BMS^-}(A^-))$;
 Move y_{max}^- from Y^U to A^- ;
 $Y^L \leftarrow Y^L \cup A^-$;
end

Algorithm 2: Belief-Maximization Selection

$$Q_{BMS^-}(A) = \sum_{y_i \in Y_{(0)}^U} p_{\mathcal{B}}(y_i = 0 | G, Y^L \cup A) \quad (12)$$

The optimization of both quality functions $Q_{BMS^+}(A)$ and $Q_{BMS^-}(A)$ is NP-hard. However, as both quality functions are submodular (theoretical analysis is given in §4.3), a solution with an approximation ratio of $(1 - 1/e)$ can be obtained using a greedy algorithm: at each time, it selects the relationship which is expected to provide the maximum marginal increase of the quality function. Notice that we treat the examining relationship node y_i as if it is positive-labeled when optimizing $Q_{BMS^+}(A)$, or negative-labeled for $Q_{BMS^-}(A)$, since the active learning algorithm is label-unaware in the selection stage. In order to leverage the risk that a selected relationship is not labeled as expected, we employ a weighting factor $p(y_i | G, Y^L)$ to reflect how likely the relationship would be labeled as positive(negative).

To prevent making an imbalance selection, we intuitively use Q_{BMS^+} to choose $b/2$ nodes (where b is the number of relationships we expect to query the user each time), and then use Q_{BMS^-} for the rest. Algorithm 2 formally describes the selection process. This selection strategy is denoted by BMS. Note that BMS combines both BMS+ and BMS-. Thus it cannot guarantee a lower error bound of the approximation.

4.3 Theoretical Analysis

We give a theoretical analysis of proposed active learning models. The approximation ratio of the IMS model is given in [17]. Here we focus on the proof

of approximation guarantees of the BMS model. The proof is based on the submodular property, which indicates that the marginal gain from adding an element to a set S is at least as high as the marginal gain from adding the same element to a superset of S . The following is a formal definition of the submodular set function.

Definition 3 (Submodular) A set function F defined on set S is called submodular, if for all $A \subset B \subset S$ and $s \notin B$, it satisfies $F(A \cup \{s\}) - F(A) \geq F(B \cup \{s\}) - F(B)$.

Given a submodular function F , which is also monotone and non-negative, it is an NP-hard problem to find a k -element subset S to optimize F . But a greedy algorithm can result in an approximation ratio of $(1 - 1/e)$. It constructs the subset by selecting elements one at a time, each time choosing an element that provides the largest marginal increase in the function value. Thus, we have,

Theorem 1 For a non-negative, monotone submodular function F , let S be the k -element subset decided by the following algorithm: for k times, each time choose an element which gives the maximum marginal increase of F and move it to S . Let S^* denotes the optimal solution. Then we have $F(S) \geq (1 - \frac{1}{e})F(S^*)$.

Now we prove the submodularity of the quality function Q_{BMS^+} defined by Equation 11.

Theorem 2 The quality function Q_{BMS^+} satisfies the submodular property.

Proof First, we prove that for all $y_i \in Y^U$, the function $p_B(y_i = 1|S)$ is monotonic.

$$\begin{aligned} p_B(y_i = 1|S \cup \{x\}) &= \frac{\mathcal{B}(y_i = 1|S \cup \{x\})}{\mathcal{B}(y_i = 1|S \cup \{x\}) + \mathcal{B}(y_i = 0|S \cup \{x\})} \\ &= \frac{1}{1 + \frac{\mathcal{B}(y_i = 0|S \cup \{x\})}{\mathcal{B}(y_i = 1|S \cup \{x\})}} \end{aligned}$$

Let $k_1 = \frac{\mathcal{B}(y_i = 0|S \cup \{x\})}{\mathcal{B}(y_i = 1|S \cup \{x\})}$, then

$$p_B(y_i = 1|S \cup \{x\}) = \frac{1}{1 + k_1}$$

Similarly, let $k_2 = \frac{\mathcal{B}(y_i = 0|S)}{\mathcal{B}(y_i = 1|S)}$, then

$$p_B(y_i = 1|S) = \frac{1}{1 + k_2}$$

According to the assumption in 4.2, it is obvious that $k_1 \leq k_2$. Obviously,

$$p_{\mathcal{B}}(y_i = 1|S \cup \{x\}) \geq p_{\mathcal{B}}(y_i = 1|S)$$

And function $p(y_i = 1|S)$ is also submodular. Suppose $A \subset B \subset Y^U$, and there is another unlabeled relationship $x \notin B$.

Similarly, we define k_1, k_2, k_3, k_4 below:

$$\begin{aligned} k_1 &= \frac{\mathcal{B}(y_i = 0|A \cup \{x\})}{\mathcal{B}(y_i = 1|A \cup \{x\})}, & k_2 &= \frac{\mathcal{B}(y_i = 0|A)}{\mathcal{B}(y_i = 1|A)} \\ k_3 &= \frac{\mathcal{B}(y_i = 0|B \cup \{x\})}{\mathcal{B}(y_i = 1|B \cup \{x\})}, & k_4 &= \frac{\mathcal{B}(y_i = 0|B)}{\mathcal{B}(y_i = 1|B)} \end{aligned}$$

Note that $Q_{BMS^+}(A)$ sums over set $Y_{(1)}^U$, hence $k_1, k_2, k_3, k_4 \leq 1$. Specifically, we have

$$\alpha = \frac{k_1}{k_2} = \frac{k_3}{k_4} \leq 1, \beta = \frac{k_3}{k_1} = \frac{k_4}{k_2} \leq 1$$

Then we can obtain the following inequality,

$$\begin{aligned} \delta(A, x) &= p_{\mathcal{B}}(y_i = 1|A \cup \{x\}) - p_{\mathcal{B}}(y_i = 1|A) \\ &= \frac{1}{1+k_1} - \frac{1}{1+k_2} \\ &= \frac{(1-\alpha)k_2}{(1+\alpha k_2)(1+k_2)} \\ \delta(B, x) &= p_{\mathcal{B}}(y_i = 1|B \cup \{x\}) - p_{\mathcal{B}}(y_i = 1|B) \\ &= \frac{(1-\alpha)k_4}{(1+\alpha k_4)(1+k_4)} \\ &= \delta(A, x) \frac{(1+\alpha)\beta k_2 + \beta + \alpha\beta k_2^2}{(1+\alpha)\beta k_2 + 1 + \alpha\beta^2 k_2^2} \\ &\leq \delta(A, x) \end{aligned}$$

Now we can give the proof of the submodularity of quality function Q_{BMS^+} . Suppose $A \subset B \subset Y^U$, and there is another unlabeled relationship $y_x \notin B$.

$$\begin{aligned} \Delta(A, x) &= Q_{BMS^+}(A \cup \{y_x\}) - Q_{BMS^+}(A) \\ &= \sum_{y_i \in Y_{(1)}^U} p_{\mathcal{B}}(y_i = 1|Y^L \cup A \cup \{y_x\}) - \sum_{y_i \in Y_{(1)}^U} p_{\mathcal{B}}(y_i = 1|Y^L \cup A) \\ &= \sum_{y_i \in Y_{(1)}^U \setminus (A \cup \{y_x\})} [\delta(Y^L \cup A, y_x)] + 1 - p_{\mathcal{B}}(y_x = 1|A) \\ &\geq \sum_{y_i \in Y_{(1)}^U \setminus (B \cup \{y_x\})} [\delta(Y^L \cup B, y_x)] + 1 - p_{\mathcal{B}}(y_x = 1|B) \\ &= \Delta(B, x) \end{aligned}$$

Table 1. Statistics of three data sets.

Data set	Users	Unlabeled Relationships	Labeled Relationships
Publication	1,036,990	1,984,164	6,096
Email	151	3,424	148
Mobile	107	5,122	314

□

Therefore, we have proved that Q_{BMS^+} is submodular. The submodularity of Q_{BMS^-} can be proved in a similar way. According to Theorem 1, it guarantees a lower bound of the greedy algorithm employed for the BMS model.

5 Experimental Results

The proposed relationship mining approach is general and can be applied to many different scenarios. In this section, we present experiments on three different genres of data sets to evaluate the effectiveness and efficiency of our proposed approach. All data sets and codes are publicly available.¹

5.1 Experiment Setup

Data sets. We evaluate the proposed methods on three different data sets: Publication, Email, and Mobile. Statistics of the data sets are listed in Table 1.

- Publication. In the publication data set, we try to infer the advisor-advisee relationship from the coauthor network. The data set is provided by [41]. Specifically, we have collected 1,632,442 publications from Arnetminer [38] (from 1936 to 2010) with 1,036,990 authors involved. The ground truth is obtained in three ways: 1) manually crawled from researcher’s homepage; 2) extracted from Mathematics Genealogy project²; 3) extracted from AI Genealogy project³. In total, we have collected 2,164 advisor-advisee pairs as positive cases, and another 3,932 pairs of colleagues as negative cases. The mining results for advisor-advisee relationships are also available in the online system Arnetminer.org.
- Email. In the email data set, we aim to infer the manager-subordinate relationship from the email communication network. The data set consists of 136,329 emails between 151 Enron employees. The ground truth of manager-subordinate relationships is provided by [7].
- Mobile. In the mobile data set, we try to infer the friendship in mobile calling network. The data set is from Eagle et al. in [9]. It consists of call logs, bluetooth scanning logs and location logs collected by a software

¹ <http://arnetminer.org/socialtie/>

² <http://www.genealogy.math.ndsu.nodak.edu>

³ <http://aigp.eecs.umich.edu>

installed in mobile phones of 107 users during a ten-month period. In the data set, users provide labels for their friendships. In total, 314 pairs of users are labeled as friends.

Factor definition. In the Publication data set, relationships are established between authors v_i and v_j if they coauthored at least one paper. For each pair of coauthors (v_i, v_j) , our goal is to identify whether v_i is the advisor of author v_j . In this data set, we consider two types of correlations: 1) *co-advisee*. The assumption is based on the fact that one could have only a limited number of advisors in her/his research career. Based on this, we define a correlation factor h_1 between nodes r_{ij} and r_{kj} . 2) *co-advisor*. Another observation is that if v_i is the advisor of v_j (i.e., $r_{ij} = 1$), then v_i is very likely to be the advisor of some other student v_k who is similar to v_j . We define another factor function h_2 between nodes r_{ij} and r_{ik} .

In the Email data set, we try to discover the “manager-subordinate” relationship. A relationship (v_i, v_j) is established when two employees have at least one email communication. There are in total 3,572 relationships among which 148 are labeled as manager-subordinate relationships. We try to identify the relationship types from the email traffic network. For example, if most of an employee’s emails were sent to the same person, then the recipient is very likely to be her manager. A correlation named *co-recipient* is defined, that is, if a user v_i sent more than ϑ emails of which recipients including both v_j and v_k (ϑ is a threshold and is set as 10 in our experiment), then, the relationship r_{ij} and r_{ik} are very likely to be the same. Therefore, a correlation factor is added between the two relationships. Two constraints named *co-manager* and *co-subordinate* are also introduced in an analogous way as that for the publication data.

In the Mobile data set, we try to identify whether two users have a friendship if there were at least one voice call or one text message sent from one to the other. Two kinds of correlations are considered: 1) *co-location*: if more than three users arrived at the same location roughly the same time, we establish correlations between all the relationships in this groups. 2) *related-call*. When v_i makes a call to both v_k and v_j from the same location, or makes a call to v_k immediately after the call with v_j , we add a related-call correlation factor between r_{ij} and r_{ik} .

In addition, we also consider some other factors in the three data sets. A detailed description of the factor definition for each data set is given in Table 6 in Appendix.

Comparison methods. We compare our approach with the following methods for inferring relationship types:

SVM: It uses the relationship attribute vector \mathbf{x}_i to train a classification model, and predict the relationships by employing the classification model. We use the SVM-light package to implement SVM.

TPFG: It is an unsupervised method proposed in [41] for mining advisor-advisee relationships in publication network. This method is domain-specific and thus we only compare with it on the Publication data set.

Table 2. Performance of relationship mining with different methods on three data sets: Publication, Email and Mobile (%).

Data set	Method	Accuracy	Precision	Recall	F1-score
Publication	SVM	76.6	72.5	54.9	62.1
	TPFG	81.2	82.8	89.4	86.0
	PLP-FGM-S	84.1	77.1	78.4	77.7
	PLP-FGM	92.7	91.4	87.7	89.5
Email	SVM	82.6	79.1	88.6	83.6
	PLP-FGM-S	85.6	85.8	85.6	85.7
	PLP-FGM	88.0	88.6	87.2	87.9
Mobile	SVM	80.0	92.7	64.9	76.4
	PLP-FGM-S	80.9	88.1	71.3	78.8
	PLP-FGM	83.1	89.4	75.2	81.6

PLP-FGM-S: The proposed PLP-FGM is based on the partially-labeled network. An alternative strategy is to train the model (parameters) with the labeled nodes only. We use this method to evaluate the necessity of the partial learning.

Evaluation measures. To quantitatively evaluate the proposed method, we consider two aspects: performance and scalability. For the relationship mining performance, we consider two-fold cross-validation (i.e., half training and half testing) and evaluate the approaches in terms of accuracy, precision, recall, and F1-score. For scalability, we examine the execution time of the model learning.

All the codes are implemented in C++, and all experiments are conducted on a server running Windows Server 2008 with Intel Xeon CPU E7520 1.87GHz (16 cores) and 128 GB memory. The distributed learning algorithm is implemented on MPI (Message Passing Interface).

5.2 Accuracy Performance

Table 2 lists the accuracy performance of inferring the type of social relationships by the different methods.

Performance comparison. Our method consistently outperforms other comparative methods on all the three data sets. In the Publication data set, PLP-FGM achieves a +27% (in terms of F1-score) improvement compared with SVM, and outperforms TPFG by 3.5% (F1-score) and 11.5% in terms of accuracy. We observe that TPFG achieves the best recall among all the four methods. This is because that TPFG tends to predict more positive cases (i.e., inferring more advisor-advisee relationships in the coauthor network), thus hurts the precision. As a result, TPFG underperforms our method by 8.6% in terms of precision. In Email and Mobile data set, PLP-FGM outperforms SVM by +4% and +5% respectively.

Unlabeled data indeed helps. From the result, it is clearly shown that by utilizing the unlabeled data, our model indeed obtains a significant improve-

Table 3. Factor contribution analysis on three data sets (%).

Data set	Factors used	Accuracy	Precision	Recall	F1-score
Publication	Attributes	77.1	71.1	59.8	64.9
	+ Co-advisor	83.5	80.9	69.8	75.0 (+10.1%)
	+ Co-advisee	83.1	79.7	70.2	74.7 (+9.8%)
	All	92.7	91.4	87.7	89.5(+24.6%)
Email	Attributes	80.1	79.5	81.2	80.3
	+ Co-recipient	80.8	81.5	79.7	80.6 (+0.3%)
	+ Co-manager	83.1	82.8	83.5	83.2 (+2.9%)
	+ Co-subordinate	85.0	84.4	85.7	85.0 (+4.7%)
	All	88.0	88.6	87.2	87.9 (+7.6%)
Mobile	Attributes	81.8	88.6	73.3	80.2
	+ Co-location	82.2	89.2	73.3	80.4 (+0.2%)
	+ Related-call	81.8	88.6	73.3	80.2 (+0.0%)
	All	83.1	89.4	75.2	81.6 (+1.4%)

ment. Without using the unlabeled data, our model (PLP-FGM-S) results in a large performance reduction (-11.8% in terms of F1-score) on the publication data set. On the other two data sets, we also observe a clear performance reduction.

Factor contribution analysis. We perform an analysis to evaluate the contribution of different factors defined in our model. We first remove all the correlation/constraint factors and only keep the attribute factor, and then add each of the factors into the model and evaluate the performance improvement by each factor. Table 3 shows the result of factor analysis. We see that almost all the factors are useful for inferring the social relationships, but the contribution is very different. For example, for inferring the manager-subordinate relationship, the co-subordinate factor is the most useful factor which achieves a 4.7% improvement by F1-score, and the co-manager factor achieves a 2.9% improvement, while the co-recipient factor only results in a 0.3% improvement. By combining all the factors together, we can further obtain a 2.9% improvement. An extreme phenomenon appears on the Mobile data set. With each of the two factors (co-location and related-call), we cannot obtain a clear improvement (0.2% and 0.0% by F1). However, when combining the two factors and the attribute factor together, we can achieve a 1.4% improvement, seven times higher than that obtained by the separated case. This is because our model not only considers different factors, but also considers the correlation between them.

5.3 Active Learning Performance

For active learning, in each data set, we first randomly select 10 relationships as the initial labeled set Y^L . And then we iteratively perform the active selection algorithm, each time selecting $b = 10$ relationships to query. After each round of selection, we learn the PLP-FGM model and evaluate the prediction

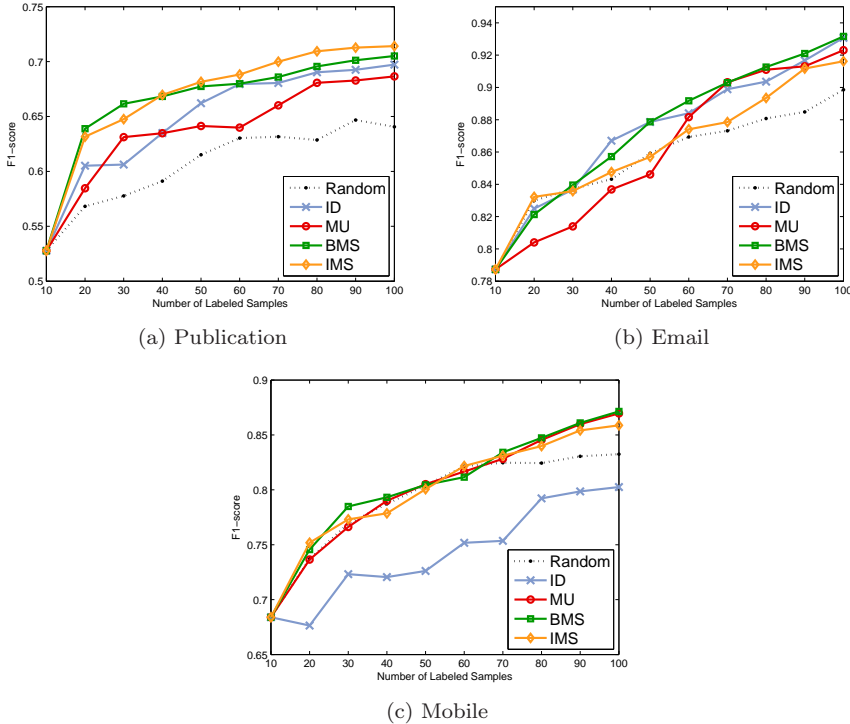


Fig. 3. Learning curves in terms of F1-score.

performance. We implement the experiment for 10 times on each data set and use the mean of F1-score for evaluation.

Comparison methods. We consider the following baseline methods:⁴

Random: It randomly selects b nodes in Y^U at each time.

Maximum Uncertainty(MU): It chooses the most b uncertain nodes among unlabeled relationships Y^U .

Information Density(ID): It chooses b nodes with the maximum average similarity to all other nodes in Y^U , proposed in [32].

Effect of active learning. We plot the learning curves on each data set in Figure 3, and list the average F1-score by all selection strategies in Table 4. The results clearly demonstrate the effectiveness of the active selection strategies. In the Publication data set, the overall F1-score of the IMS strategy with 100 samples labeled outperforms the Random algorithm by +7.4%. In Email and Mobile data set, the BMS strategy achieves the best performance, with an improvement of +3.3% and +3.9% respectively.

⁴ We did not consider the *co-advisee* correlation in the model when dealing with the Publication data set and the *co-subordinate* correlation for the Email data set, since they conflict with the assumption of monotonic belief 4.2.

Table 4. Average F1-score by all selection strategies (%). The results were obtained by randomly selecting 10 relationships as the initial labeled set Y^L , and then iteratively perform the active selection algorithms, each time with $b = 10$ relationships to query.

Data set	Random	MU	ID	BMS	IMS
Publication	60.6	63.7	64.8	66.4	66.8
Email	85.6	86.2	87.3	87.6	86.3
Mobile	79.2	80.0	74.3	80.4	79.9

Performance comparison. In the publication data set, both the proposed BMS and IMS strategy significantly performs better than all the baseline methods (paired t -tests with 95% significance). BMS also significantly outperforms Random and ID strategy in Mobile data set, while its performance is close to MU. The performance of IMS is shown better than ID in Mobile data set, but is close to other baseline methods. In Email data set, BMS significantly outperforms Random, while the performance of other methods seems close to each other. Generally, the proposed BMS strategy performs more consistently, and obtain better result in two of the three data sets. The performance of IMS strategy is the best in Publication data set, but seems close to baseline methods in the other two data sets.

Network information helps. According to factor contribution analysis mentioned before, *co-advisor* factor in Publication data set contributes the most. This explains why the proposed methods achieve better performance than the alternative baseline methods in Publication data set. The average F1-score of BMS and IMS reaches 65% with less than 30 labeled samples, while ID uses more than 40, and MU uses more than 60. In Email and Mobile data set, BMS still takes advantage of the network information, but the improvement shrinks due to the considerable decrease of factor contribution.

In-depth analysis of BMS. There are also some variations of BMS and we conduct a comparison between them. BMS+ selects all b nodes optimizing $Q_{BMS+}(A)$, while BMS- employs $Q_{BMS-}(A)$. Figure 4 shows the average F1-score of the different versions. In Publication and Email data set, the difference between BMS and BMS+ is subtle, while the performance of BMS- drops. It might be resulted from different criteria of these three strategies. BMS+ tends to obtain true-positive samples, whereas BMS- is more likely to acquire true-negative samples. F1-score excludes the impact of true-negative samples, and therefore undermines the performance of BMS-. The gap disappears in Mobile data set, probably due to the weak contribution of its correlation and constraints factors.

5.4 Distributed Learning

As real social networks may contain millions of users and relationships, it is important for the learning algorithm to scale well with large networks. To

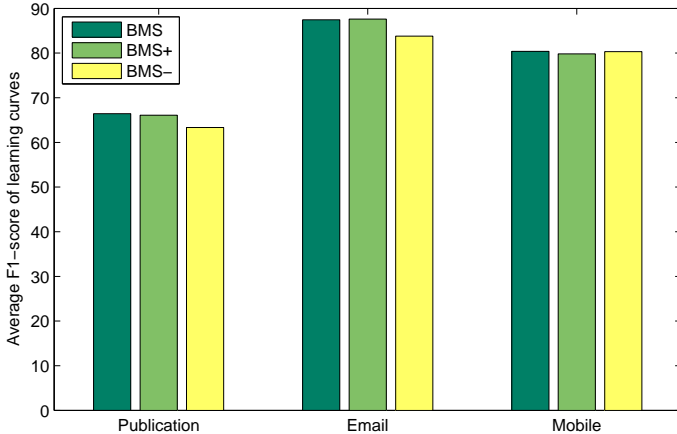


Fig. 4. Performance comparison between variations of BMS.

Table 5. Data transferred in distributed learning algorithm.

Phase	From	To	Data Description
Initialization	Master	Slave i	i -th subgraph
Iteration Beginning	Master	Slave i	Current parameters θ
Iteration Ending	Slave i	Master	Gradient in i -th subgraph

address this issue, we develop a distributed learning method based on MPI (Message Passing Interface). The learning algorithm can be viewed as two steps: 1) compute the gradient for each parameter via loopy belief propagation; 2) optimize all parameters with the gradient descents. The most expensive part is the step of calculating the gradient. Therefore we develop a distributed algorithm to speed up the process.

We adopt a *master-slave* architecture, i.e., one master node is responsible for optimizing parameters, and the other slave nodes are responsible for calculating gradients. At the beginning of the algorithm, the graphical model of PLP-FGM is partitioned into P roughly equal parts, where P is the number of slave processors. This process is accomplished by graph segmentation software METIS[16]. The subgraphs are then distributed over slave nodes. Note that in our implementation, the edges (factors) between different subgraphs are eliminated, which results in an approximate solution. In each iteration, the master node sends the newest parameters θ to all slaves. Slave nodes then start to perform Loopy Belief Propagation on the corresponding subgraph to calculate the marginal probabilities, then further compute the parameter gradient and send it back to the master. Finally, the master node collects and sums up all gradients obtained from different subgraphs, and updates parameters by the gradient descent method. The data transferred between the master and slave nodes are summarized in Table 5.

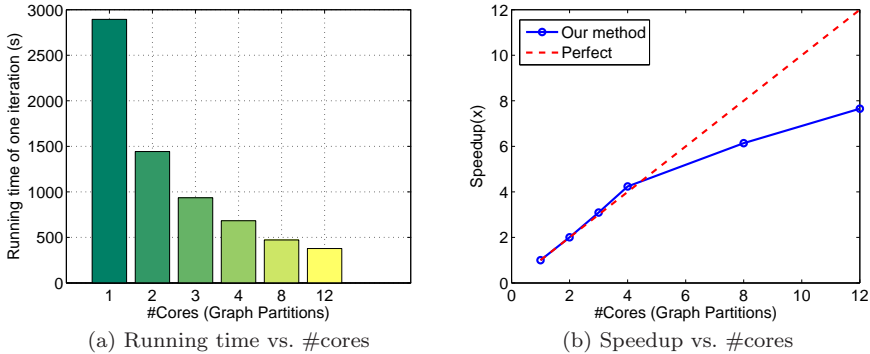


Fig. 5. Scalability performance.




Fig. 6. Approximation of graph partition.

We conduct a series of experiments to evaluate the scalability performance of our distributed learning algorithm on the Publication data set. Figure 5 shows the running time and speedup of the distributed algorithm with different number of computer nodes (2,3,4,8,12 cores) used. The speedup curve is close to the perfect line at the beginning. Although the speedup inevitably decreases when the number of cores increases, it can achieve $\sim 8\times$ speedup with 12 cores. It is noticeable that the speedup curve is beyond the perfect line when using 4 cores, it is not strange since our distributed strategy is approximated. In our distributed implementation, graphs are partitioned into subgraphs, and the factors across different parts are discarded. Thus, the graph processed in distributed version contains less edges, making the computational cost less than the amount in the original algorithm. The effect of subgraph partition is illustrated in Figure 6. By using good graph partition algorithm such as METIS, the performance only decreases slightly (1.4% in accuracy and 1.6% in F1-score). A theoretical study of the approximate ratio for the distributed learning algorithm would be an interesting issue and is also one of our ongoing work.

6 Related work

Relationship Mining Relationship mining is an important problem in social network analysis. One research branch is to predict and recommend unknown links in social networks. Liben-Nowell et al.[24] study the unsupervised methods for link prediction. Menon et al.[26] propose a log-linear matrix model for dyadic prediction. Backstrom et al. [2] propose a supervised random walk algorithm to estimate the strength of social links. Leskovec et al. [23] employ a logistic regression model to predict positive and negative links in online social networks, where the positive links indicates the relationships such as friendship, while negative indicating opposition. However, these works consider only the black-white social networks, and do not consider the types of the relationships. There are also several works on mining the relationship semantics. Diehl et al. [7] try to identify the manager-subordinate relationships by learning a ranking function. Wang et al. [41] propose an unsupervised probabilistic model for mining the advisor-advisee relationships from the publication network. Eagle et al. [9] present several patterns discovered in mobile phone data, and try to use these pattern to infer the friendship network. Tang et al.[39] develop a classification framework of social media based on differentiating different types of social connections. However, these algorithms mainly focus on a specific domain, while our model is general and can be applied to different domains. Moreover, these methods do not explicitly consider the correlation information between different relationships. Hopcroft et al.[15] explore the problem of reciprocal relationship prediction and Tang et al. [36] have developed a framework for classifying the type of social relationships by learning across heterogeneous networks. However, they do not consider how to make optimal use of user interaction.

Another related research topic is relational learning[4,11]. However, the problem presented in this paper is very different. Relational learning focuses on the classification problems when objects or entities are presented in relations, while this paper explores the relationship types in social network. A number of supervised methods for link prediction in relational data have also been developed [40,29].

Active Learning for Networked Data The exploding data size in social network, bioinformatics, and many other fields nourishes the development of studies on active learning of networked data. Settles et al.[32] provide a survey for active learning strategies in sequence labeling tasks. There have been several works focus on some special graphical structures. Krause et al.[20] develop an algorithm to nonmyopically optimize the active learning task in chain graphical models. Cesa-Bianchi et al.[5] studies active learning on trees. Also, active learning on several specific model has been explored. Martinez et al.[25] combine active learning method with a CRF-based classifier and propose an automated online learning framework. Golovin et al.[12] develop a greedy algorithm for Bayesian active learning with noise. Shi et al. [33] study the problem of batch mode active learning for networked data. They propose

a batch mode active learning method by combining three criteria (i.e., minimum redundancy, maximum uncertainty and maximum impact). Similar to our intuition, Kuwadekar et al.[21] also attempt to leverage the network structure to benefit the active learning performance. But they study the problem on across-network task, while we focus on within-network setting. Kimura et al.[18] propose an effective algorithm to extract influential nodes in a social network. In addition, Bilgic et al.[3] propose a method for active inference on networked data, based on uncertainty sampling, committee-based sampling and clustering; Roy et al.[31] directly optimize expected future error.

7 Conclusion

In this paper, we study the problem of inferring social ties in large networks. We define the problem in a semi-supervised framework, and propose a Partially-Labeled Pairwise Factor Graph Model (PLP-FGM) for learning to infer the relationship semantics. In PLP-FGM, relationships in social network are modeled as nodes, and the type of social relationships are modeled as hidden variables. An efficient algorithm is presented to learn model parameters and to predict unknown relationships. Experimental results on three different types of data sets validate the effectiveness of the proposed model. To further scale up to large networks, a distributed learning algorithm has been developed. Experiments demonstrate good parallel efficiency of the distributed learning algorithm. In order to effectively learn the mapping function, we propose two active learning strategies: Influence-Maximization Selection and Belief-Maximization Selection, both aiming to capture the inter-relationship influence. Experimental results show that BMS and IMS often achieve significant better performance than baseline methods.

So far PLP-FGM is utilized to infer social ties in a static social network. However, most social network in our life is dynamic and it would be worthwhile to extend the model to the dynamic case. PLP-FGM works in a semi-supervised framework, requiring labeled data to learn the predictive function. Some unsupervised methods can be studied for the relationship mining problem for a broader utilization.

Detecting the relationship semantics makes online social networks colorful and closer to our real physical networks. It represents a new research direction in social network analysis. As future work, it is interesting to further study to which extent we should trust user feedbacks. In addition, it would be also interesting to investigate how the inferred relationship semantic information can help other applications such as community detection, influence analysis, and link recommendation.

References

1. R. Albert and A. L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1), 2002.

2. L. Backstrom and J. Leskovec. Supervised random walks: predicting and recommending links in social networks. In *WSDM*, pages 635–644, 2011.
3. M. Bilgic, L. Mihalkova, and L. Getoor. Active learning for networked data. In J. Fürnkranz and T. Joachims, editors, *ICML*, pages 79–86. Omnipress, 2010.
4. M. E. Califf and R. J. Mooney. Relational learning of pattern-match rules for information extraction. In *AAAI/IAAI*, pages 328–334, 1999.
5. N. Cesa-Bianchi, C. Gentile, F. Vitale, and G. Zappella. Active learning on trees and graphs. In *COLT*, pages 320–332, 2010.
6. D. Crandall, L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg. Inferring social ties from geographic coincidences. *PNAS*, 107(52):22436, 2010.
7. C. P. Diehl, G. Namata, and L. Getoor. Relationship identification for social network discovery. In *AAAI*, pages 546–552. AAAI Press, 2007.
8. P. Domingos and M. Richardson. Mining the network value of customers. In *KDD*, pages 57–66, 2001.
9. N. Eagle, A. S. Pentland, and D. Lazer. Mobile phone data for inferring social network structure. *Social Computing, Behavioral Modeling, and Prediction*, pages 79–88, 2008.
10. M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.
11. L. Getoor and B. Taskar. *Introduction to statistical relational learning*. The MIT Press, 2007.
12. D. Golovin, A. Krause, and D. Ray. Near-optimal bayesian active learning with noisy observations. *CoRR*, abs/1010.3091, 2010.
13. R. Grob, M. K. 0002, R. Wattenhofer, and M. Wirz. Cluestr: mobile social networking for enhanced group communication. In *GROUP*, pages 81–90, 2009.
14. J. M. Hammersley and P. Clifford. Markov field on finite graphs and lattices. *Unpublished manuscript*, 1971.
15. J. E. Hopcroft, T. Lou, and J. Tang. Who will follow you back? reciprocal relationship prediction. In *CIKM'11*, 2011.
16. G. Karypis and V. Kumar. *MeTis: Unstructured Graph Partitioning and Sparse Matrix Ordering System, Version 4.0*, Sept. 1998.
17. D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.
18. M. Kimura, K. Saito, R. Nakano, and H. Motoda. Extracting influential nodes on a social network for information diffusion. *Data Min. Knowl. Discov.*, 20(1):70–97, 2010.
19. J. Kleinberg. Temporal dynamics of on-line information streams. In *Data Stream Management: Processing High-speed Data*. Springer, 2005.
20. A. Krause and C. Guestrin. Optimal value of information in graphical models. *J. Artif. Intell. Res. (JAIR)*, 35:557–591, 2009.
21. A. Kuwadekar and J. Neville. Relational active learning for joint collective classification models. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, ICML '11, pages 385–392, New York, NY, USA, June 2011. ACM.
22. J. D. Lafferty, A. McCallum, and F. C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the 18th International Conference on Machine Learning (ICML'01)*, pages 282–289, 2001.
23. J. Leskovec, D. P. Huttenlocher, and J. M. Kleinberg. Predicting positive and negative links in online social networks. In *WWW*, pages 641–650, 2010.
24. D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *JASIST*, 58(7):1019–1031, 2007.
25. O. Martinez and G. Tsechpenakis. Integration of active learning in a collaborative crf. *Computer Vision and Pattern Recognition Workshop*, 0:1–8, 2008.
26. A. K. Menon and C. Elkan. A log-linear model with latent features for dyadic prediction. In *ICDM*, pages 364–373, 2010.
27. K. Murphy, Y. Weiss, and M. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *UAI*, volume 9, pages 467–475, 1999.
28. M. E. J. Newman. The structure and function of complex networks. *SIAM Reviews*, 45, 2003.

29. A. Popescul and L. Ungar. Statistical relational learning for link prediction. In *IJCAI03 Workshop on Learning Statistical Models from Relational Data*, volume 149, page 172, 2003.
30. M. Roth, A. Ben-David, D. Deutscher, G. Flysher, I. Horn, A. Leichtberg, N. Leiser, Y. Matias, and R. Merom. Suggesting friends using the implicit social graph. In *KDD*, pages 233–242, 2010.
31. N. Roy and A. McCallum. Toward optimal active learning through sampling estimation of error reduction. In *ICML*, pages 441–448, 2001.
32. B. Settles and M. Craven. An analysis of active learning strategies for sequence labeling tasks. In *EMNLP*, pages 1070–1079, 2008.
33. L. Shi, Y. Zhao, and J. Tang. Batch mode active learning for networked data. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2011.
34. S. H. Strogatz. Exploring complex networks. *Nature*, 410:268–276, 2003.
35. C. Tan, J. Tang, J. Sun, Q. Lin, and F. Wang. Social action tracking via noise tolerant time-varying factor graphs. In *KDD*, pages 1049–1058, 2010.
36. J. Tang, T. Lou, and J. Kleinberg. Inferring social ties across heterogenous networks. In *WSDM'12*, 2012.
37. J. Tang, J. Sun, C. Wang, and Z. Yang. Social influence analysis in large-scale networks. In *KDD*, pages 807–816, 2009.
38. J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. Arnetminer: Extraction and mining of academic social networks. In *KDD'08*, pages 990–998, 2008.
39. L. Tang and H. Liu. Leveraging social media networks for classification. *Data Min. Knowl. Discov.*, 23(3):447–478, 2011.
40. B. Taskar, M. F. Wong, P. Abbeel, and D. Koller. Link prediction in relational data. In *NIPS*. MIT Press, 2003.
41. C. Wang, J. Han, Y. Jia, J. Tang, D. Zhang, Y. Yu, and J. Guo. Mining advisor-advisee relationships from research publication networks. In *KDD*, pages 203–212, 2010.
42. D. Wang, D. Pedreschi, C. Song, F. Giannotti, and A.-L. Barabási. Human mobility, social ties, and link prediction. In *KDD*, pages 1100–1108, 2011.
43. Z. Yang, J. Guo, K. Cai, J. Tang, J. Li, L. Z. 0007, and Z. Su. Understanding retweeting behaviors in social networks. In *CIKM*, pages 1633–1636, 2010.

Appendix: Feature Definition

In this section, we introduce how we define the attribute factor functions. In the Publication data set, we define five categories of attribute factors: Paper count, Paper ratio, Coauthor ratio, Conference coverage, First-paper-year-diff. The definitions of the attributes are summarized in Table 6. In the Email data set, traffic-based features are extracted. For a relationship, we compute the number of emails for different communication types. In the Mobile data set, the attributes we extracted are #voice calls, #messages, Night-call ratio, Call duration, #proximity and In-role proximity ratio.

Table 6. Attributes used in the experiments. In the Publication data set, we use P_i and P_j to denote the set of papers published by author v_i and v_j respectively. For a given relationship (v_i, v_j) , five categories of attributes are extracted. In the Email data set, for relationship (v_i, v_j) , number of emails for different communication types are computed. In the Mobile data set, the attributes are from the voice call/message/proximity logs.

Data set	Factor	Description	
Publication	Paper count	$ P_i , P_j $	
	Paper ratio	$ P_i / P_j $	
	Coauthor ratio	$ P_i \cap P_j / P_i , P_i \cap P_j / P_j $	
	Conference coverage	The proportion of the conferences which both v_i and v_j attended among conferences v_j attended.	
	First-paper-year-diff	The difference in year of the earliest publication of v_i and v_j .	
Email	Traffics	Sender	Recipients Include
		v_i	v_j
		v_j	v_i
		v_i	v_k and not v_j
		v_j	v_k and not v_i
		v_k	v_i and not v_j
		v_k	v_j and not v_i
		v_k	v_i and v_j
Mobile	#voice calls	The total number of voice call logs between two users.	
	#messages	Number of messages between two users.	
	Night-call ratio	The proportion of calls at night (8pm to 8am).	
	Call duration	The total duration time of calls between two users.	
	#proximity	The total number of proximity logs between two users.	
	In-role proximity ratio	The proportion of proximity logs in "working place" and in working hours (8am to 8pm).	