This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

A General Framework for Implicit and Explicit Social Recommendation

Chin-Chi Hsu, Mi-Yen Yeh, Shou-De Lin

Abstract—Research of social recommendation aims at exploiting social information to improve the quality of a recommender system. It can be further divided into two classes. Explicit social recommendation assumes the existence of not only the users' ratings on items, but also the explicit social connections between users. Implicit social recommendation assumes the availability of only the ratings but not the social connections between users, and attempts to infer implicit social connections between users with the goal to boost recommendation accuracy. This paper proposes a unified framework that is applicable to both explicit and implicit social recommendation. We propose an optimization framework to learn the degree of social correlation and rating prediction jointly, so these two tasks can mutually boost the performance of each other. Furthermore, a well-known challenge for implicit social recommendation is that it takes quadratic time to learn the strength of pairwise connections. This paper further proposes several practical tricks to reduce the complexity of our model to be linear to the observed ratings. The experiments show that the proposed model, with only two parameters, can significantly outperform the state-of-the-art solutions for both explicit and implicit social recommender systems.

Index Terms—Recommender Systems, Social Networks

1 Introduction

SOCIAL recommendation, a study aiming at incorporating social information of users into a recommender system, has attracted decent attention in recent years. It can further be divided into two tracks: *explicit social recommendation* and *implicit social recommendation*. In explicit social recommendation, a variety of models have been developed to exploit the existing social network information to enhance the performance of a recommender system [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13]. A common and arguably most successful strategy is to integrate the social information, such as *trust* or *friendship*, into a collaborative filtering model in a certain way. Figure 1 describes an explicit social recommendation system given edge strength information is available, while Figure 2 shows another kind of explicit social recommender system where only binary relationship information (e.g., whether two people are friends) is available. Suppose there is a rating dataset including some ratings of four users $\{U_1, U_2, U_3, U_4\}$ to four items $\{V_1, V_2, V_3, V_4\}$. Such data can be denoted by a matrix where the "?" entries represent unknown ratings. A social-based recommender system reads the matrix together with a given or inferred user social network as the training examples, and then predicts the unknown ratings.

Note that "*social recommendation*" in this paper does not refer to recommending links in social networks; instead social networks serve as auxiliary information to improve the quality of a recommender systems.

The information of the strength of social relationship can be very useful to a recommender system, as it is reasonable to assume people trust the ratings from their closer friends comparing to those from their acquaintances. Given the rating data along with a binary social network, several works [1], [2], [9], [14] of explicit social recommendation have proposed methods to determine the social connection strength between users to enhance the quality of recommendation (see Figure 2).

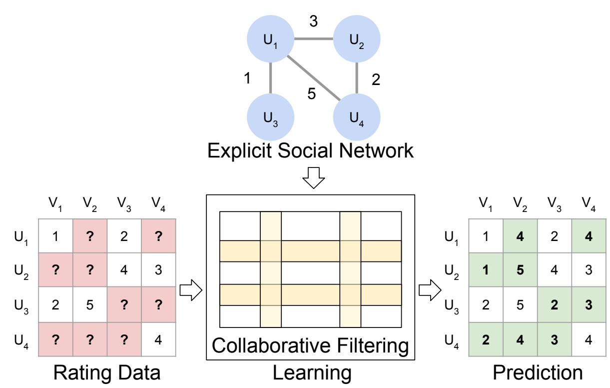

Fig. 1. Explicit social recommendation based on collaborative filtering techniques, assuming the social strength information(i.e. edge weight, scale of trust or friendship) is available.

Unfortunately, such trust or friendship data may not necessarily be available for every recommendation scenario due to the budget or privacy concerns. To address such concern, there emerges another research direction named *implicit social recommendation*, which aims at mining implicit user social relationship from historical rating data for better recommendation. Without any explicit social data, certain methods [15], [16], [17], [18], [19] have been proposed to generate an implicit social network from given ratings (see Figure 3). The pseudo links and/or their strengths can then play as a surrogate of the explicit social network to

- C.-C. Hsu and M.-Y. Yeh are with the Institute of Information Science, Academia Sinica, Taipei, Taiwan
E-mail: chinchi, miyen@iis.sinica.edu.tw
- S.-D. Lin are with the Department of Computer Science and Information Engineering National Taiwan University Taipei, Taiwan
E-mail: sdlin@csie.ntu.edu.tw

Manuscript received April 19, 2005; revised August 26, 2015.

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

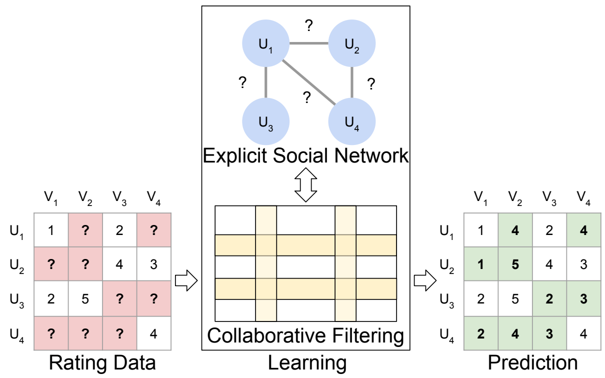

Fig. 2. Explicit social recommendation without social strengths between two connected users. Several existing solutions attempt to guess the social strengths using historical ratings to improve the overall performance.

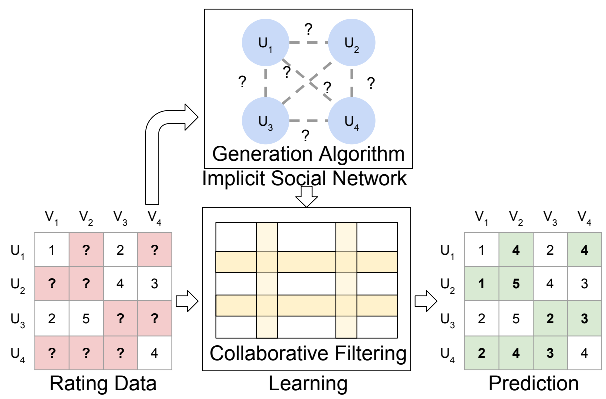

Fig. 3. Implicit social recommendation. A pseudo user social network is generated according to the rating data. It plays as the surrogate of the true social network for existing explicit social recommendation.

be incorporated into any explicit social recommendation model.

The goal of this paper is to propose a unified framework to accommodate both scenarios described above. Furthermore, it aims to address the following concerns in the existing social recommender systems.

Concerns for explicit social recommender systems: the quality of the given social information is sometimes questionable. Since most of the social data are collected from the web or social network services, inevitably they contain noises. For example, past empirical studies [4], [20] have shown that the auxiliary of *friendship* links is less useful than *trust* links in boosting the recommendation performance. Furthermore, although it is generally believed trust or friendship are positively correlated with the level of common-taste of people, this study [21] has shown that two users may not have similar rating tastes even they strongly trust each other. Thus, directly utilizing any given social connection may harm the recommendation performance.

Some concerns have been raised on using explicit social networks for recommendation. A past work [19] has conducted analysis on three social recommendation datasets to report three observations as follows: (1) Explicit social networks as adjacency matrices are typically sparse, which limits the amount of information that can be revealed. (2) For a user, the number of ratings is positively correlated to the number friends. It implies that explicit social networks could not help inferring preferences of cold-start users, who rate few items in rating datasets and likely do not possess too much social information. (3) Active users users rates many items) do not really possess similar rating patterns with their friends since active users does not necessary make friends based on interest sharing. It violates a common assumption in existing social recommendation works that *friends share similar preferences*. Consequently, this work concludes that explicit social networks are not always beneficial for recommendation.

Concern for implicit social recommender systems: most related works attempt to empirically *define* ad-hoc metrics to generate an implicit social network, and then feed it into an explicit social recommendation model. Such strategy requires careful tuning of thresholds and parameters in the metrics for different domains and different rating scales, thus can hardly be generalized. As will be shown in our experiments later on, efforts spent to determine a common parameter that can be effective in inferring implicit social networks across different datasets are usually futile, thus hinders the effectiveness of such models.

To address the above concerns, we propose a general social recommender model applicable to the scenarios with or without an explicit social network, as shown in Figure 4. Given an explicit binary social network, our model learns the strength of links from rating data to boost the quality of rating prediction. When the explicit social network is missing, our model learns jointly the existence and strength of social relationships from ratings. Different from most of the previous solutions for implicit social network that treats the learning of the social network and recommendation as two sequential but independent tasks, we propose a Variational Expectation Maximization (VEM) based solution that conducts the learning of social structure and rating prediction together. By jointly learning the social relationship and ratings from data, our data driven solution can not only absorb the potential damage brought up by noisy explicit social network, but also alleviates the concern of certain ad-hoc, less-general rules for deriving implicit social networks. To improve the usability and scalability, we further describe some implementation tricks to avoid quadratic time and space complexity. Experiments show that the proposed solution outperforms the state-of-the-art models in both explicit and implicit scenarios.

Another major strength of our model is that it learns the importance of each latent feature to determine the similarity between users. Conventional social recommender systems tend to regularize the latent factors of users so that the users with social connections tend to have similar latent factors. However, in the traditional setup, each latent dimension is assumed to contribute equally to the determination of similarity. We argue that this assumption is problematic since the "area of similarity' between every pairs of friends can be very different. For instance, A and B may become friends because they both like watching romance movies (assuming romance as one latent dimension), but it does not mean that A and B shall like action movies equally (assuming action represents another latent dimension). Similarly, friends A and C may have similar taste for action movies, but not

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

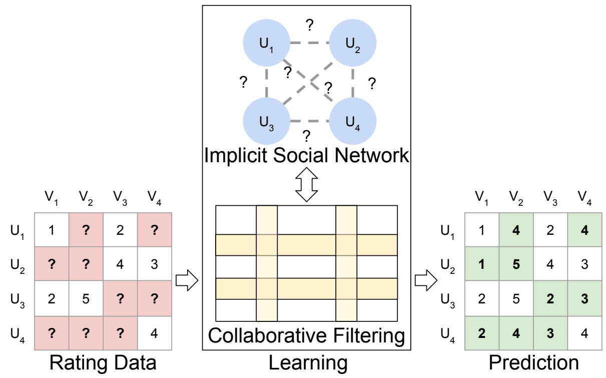

Fig. 4. The goal of our work in this paper. If there exists a binary social network from external source, then our proposed model works as Figure 2 illustrates. However, even without any given information of social structure, we would like to design a collaborative filtering model to automatically learn the underlying pairwise social connections from existing rating data. We hope that such social learning could boost the prediction accuracy of the recommender system.

necessary for other types of movies. The example shows that assuming two friends share similar latent vector (i.e. possess similar taste on all the latent dimensions) is an overly-strong assumption. In our model, such assumption is relaxed. Instead we let the model learn which latent dimension(s) best capture the similarity between two friends. In the previous example, our model can learn from data that A and B become friends because they have similar taste on romance movies, thus will only synchronize their preference on this dimension. We believe this is another major reason that our model can significantly outperform the competitors.

Another interesting finding we would like to share is that our solution can be applied to learn "implicit item relationship" as well. That is, we treat items as humans and learn their relationships for recommendation. Technically, our model learns which latent dimension(s) between two items should be aligned, and experiments show that considering such implicit item relationship in our model can further boost the recommendation quality significantly.

Below are the main contributions in this work:

- We propose a unified framework that can be adopted in either explicit and implicit social recommendation.
- To learn our joint learning framework, we designed a VEM-based optimization with two practical tricks to avoid quadratic complexity and cumbersome parameter tuning process.
- The experiments show that our model can not only outperform the state-of-the-art recommendation models and social recommendation models, but also produce superior results comparing to explicit social recommendation models.

Finally, we share the detailed mathematical derivation as well as the source code for reproducing the results. [1]

2 Related Work

As summarized in Section 1, prior works consist of either explicit social recommendation (using information from accessible external social networks) or implicit social recommendation (network topology and edge strengths inferred from the rating data). We review these two classes of works in this section.

Explicit social recommendation usually incorporates accessible social relations among users. In real-world applications, the relation of two users is usually defined by *trust* or *friendship*. Roughly speaking, trust indicates an one-way relation created by a user toward another. User A trusting user B does not necessarily imply B also trusting A. Conversely, friendship represents a mutual relationship between two users. User A making friends with user B implies A is also a friend of B. We refer readers to researches [20] and [4] that conduct experiments to compare the properties of these two types of relationships. Though some previous works rely on using the trust networks while others focus on friendship networks, we find that most of the proposed models are applicable for either case, regardless of whether directed trust networks or undirected friendship networks are exploited. Early solutions of explicit social recommendations [3], [5], [10] determine the recommended items by exploring the trust networks. Later on, researchers start to bring social information into the matrix factorization models with various assumptions social recommendation integration. SoRec [7] uses the shared user latent factors for both the rating matrix and the social matrix factorization. RSTE [8] predicts a user's ratings by the linear combination of the user and his or her trustees' latent factor vectors. SocialMF [6] defines that a user's latent factors should be close to the linear combination of his or her trusted friends' latent factors. Social Regularization [9] considers a pairwise assumption that two users trusting each other should have similar latent factors, and thus appends a regularization term to the classical matrix factorization model. There exist works modeling social influence based on collaborative filtering, such as conditional random field [12] and probabilistic Poisson factorization [1]. TrustSVD [4] incorporates the trust networks with SVD++, a variant of matrix factorization modeling the implicit influence from user latent factors through observed ratings. PTPMF [11] supposes that an observed social link shall be either weak or strong preference transfer between two connected users, and then declares two types of user latent factors to distinguish the social influence of weak-tied friendship and strong-tied friendship. However, often an explicit social network is available but the edge strength information is missing or less believable due to privacy or the budget constraints. Some solutions are proposed to extract helpful edge strength features from the existing given ratings. For example, combining with SocialMF, Fazeli et al. [14] survey and compare the performance of different trust strength metrics on an explicit social network. Fang et al. [2] fuse support vector regression with matrix factorization to learn both ratings and strengths from an explicit trust network. Social Regularization [9] considers cosine similarity or Pearson correlation coefficient between two users as candidates of social strength definition. Compared with the above approaches, the main contribution of our model is that it obtains the strength of each individual social connection through *learning the importance of each latent dimension*. Therefore, our model is flexible enough to determine the different types of social relationships between users, and

1. https://github.com/ntumslab/SocialCovariancePrior

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

thus learns the importance of each individual connection.

Implicit social recommendation attempts to extract latent social relations between two users from the given rating behaviors. The generated information serves as the surrogate of the explicit social networks in the explicit social recommender systems. Despite the quadratic time complexity of evaluating pairwise user or item social relations, Guo et al. [15] study user-based collaborative filtering that recommends items using a trust network generated from the predefined trust metrics. With the existence of extra features, Lin et al. [16] (rating time) and Guo et al. [22] (text review) present methods to obtain implicit social networks. There are some works that apply matrix factorization techniques on implicit social networks. RSTE [18]/Social Regularization [17] read implicit social networks generated by the evaluation of cosine similarity/Pearson correlation together with predefined thresholds to determine the social connections. Recently a novel generation method CUNE [19] has been proposed for implicit social recommendation by performing network embedding in a predefined user social network that comes from ratings. The method also requires a predefined threshold e.g. k to determine the top-k friends with the highest cosine similarity of network embeddings of two connected users. Comparing to the above models, the main strength of our approach is that *it models the rating prediction and the social strength learning as a joint optimization task.* These two tasks can mutually reinforce the quality of each other to achieve better results. Our approach is data-driven and thus does not need ad-hoc metrics or handcrafted thresholds to determine the existence/strength of social relationship.

3 MODEL DESIGN

In this section, we would like to present our unified social recommendation model, named *Social Covariance Prior (SCP)*. We first state the formal problem definition in Section 3.1, followed by the review of PMF in Section 3.2. Section 3.3 presents the overview of SCP.

3.1 Problem Definition

We are given *rating data* denoted by matrix $R \in \mathbb{R}^{N \times M}$, where N is the number of users and M is the number of items. An observed or non-missing entry R_{ij} records a numerical rating score that user i, $1 \leq i \leq N$, gives to item j, $1 \leq j \leq M$, as a training instance. Some datasets also provide an *explicit social network* to represent the social relations among users. An adjacent matrix $S^{(E)} \in \mathbb{R}^{N \times N}$ denotes the social network, where entry $S^{(E)}_{if}$ is binary (whether user i trusts or makes friends with user f). The goal of explicit social recommendation is to predict the unobserved values in R by incorporating both observed values in R and social information from $S^{(E)}$. For implicit social recommendation, the goal is, even without $S^{(E)}$, to mine an *implicit social network* $S^{(I)}$ from R to enhance the prediction accuracy.

3.2 Probabilistic Matrix Factorization (PMF)

Since the proposed SCP model is an extension of probabilistic matrix factorization, we first briefly introduce PMF. Mnih et al. [23] proposed a probabilistic view of matrix factorization. Matrix factorization assumes a low-rank $R \approx U^\top V$. That is, R can be approximated by the multiplication of user latent matrix $U \in \mathbb{R}^{K \times N}$ and item latent matrix $V \in \mathbb{R}^{K \times M}$ ($K \ll N, K \ll M$), where column vectors U_i, V_j reflect latent rating characteristics of user i and item j, respectively. Note that throughout this paper, we use K to denote the number of latent factors. The corresponding objective function is shown as follows:

$$\arg\min_{U,V} \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{M}\delta_{ij}\left(R_{ij} - U_i^\top V_j\right)^2 + \frac{\alpha_U}{2}\sum_{i=1}^{N}\|U_i\|_2^2 + \frac{\alpha_V}{2}\sum_{j=1}^{M}\|V_j\|_2^2 \tag{1}$$

$\delta_{ij} \in \{0,1\}$ indicates whether R_{ij} is observed in the training data. α_U, α_V are regularization parameters. Under probabilistic view, an observed rating R_{ij} is generated by a normal distribution of mean $U_i^\top V_j$ as the likelihood function. Corresponding to $l2$-norm regularization, zero-mean spherical Gaussian priors are used to generate U_i and V_j. Hence, the overall posterior to be optimized is

$$\begin{aligned}&\arg\max_{U,V} p(U, V|R, \theta)\\ =&\prod_{i=1}^{N}\prod_{j=1}^{M}\mathcal{N}(R_{ij}|U_i^\top V_j, \sigma_R^2)^{\delta_{ij}}\prod_{i=1}^{N}\mathcal{N}(U_i|0, \sigma_U^2 I)\prod_{j=1}^{M}\mathcal{N}(V_j|0, \sigma_V^2 I).\end{aligned} \tag{2}$$

$\mathcal{N}$ denotes a normal distribution, where the mean is the inner product $U_i^\top V_j$ and the variance is σ_R^2. σ_U^2, σ_V^2 are shared variances among users and items, and I is the identity matrix. $\theta = \{\sigma_R^2, \sigma_U^2, \sigma_V^2\}$ denotes all the hyperparameters. Note that $\mathcal{N}$ for vectors U_i and V_j becomes multivariate normal distribution. A common solution to (2) is to take $(-\log)$ (then we obtain (1)) and perform optimization using stochastic gradient descent (SGD). The probabilistic view enables us to combine matrix factorization with other probabilistic models like Latent Dirichlet Allocation (LDA) [24] more naturally. Moreover, the Bayesian treatment such as Markov Chain Monte Carlo (MCMC) [25] and variational inference [26], [27] can then be applied.

3.3 Social Covariance Prior (SCP)

3.3.1 SCP for Implicit User Social Recommendation

We first describe a general version of our model that assumes the non-existence of the social connections, where the implicit relationship has to be inferred using the ratings.

The fundamental idea is that we propose a series of social-based priors to regularize the underlying pairwise distances among user latent factors:

$$\prod_{(i,f)\in E_U}\left[\mathcal{N}(U_i|U_f, S_{Uif}^{-1})\mathcal{W}(S_{Uif}|K, \Lambda_U)\right]^{\frac{1}{T_{Ui}}}, \tag{3}$$

where E_U represents the set of directed edges in the user social network. Since an explicit social network is not available, here we assume E_U is a *fully connected graph*. T_{Ui} is the number of user i's neighbors (i.e., trustees or friends). Such prior will be incorporated into the original PMF equation to describe the user social network structure to be learned. Similar to previous works such as [9], SCP assumes that the social influence is implicitly transferred through the factorized matrix U. We model a pairwise social relation using a K-dimensional multivariate normal distribution $\mathcal{N}$

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

with the mean as neighbor f's latent factor vector, which regularizes the latent factors of social neighbors to be similar "to some extend" (i.e. rather than forcing the whole latent vector to be similar). The major difference from the previous works is that we introduce the personalized inverse covariance matrix S_{Uif} to fine-tune the similarity metrics between users. For each potential connection between persons i and f, we want to learn S_{Uif} to capture their similarity, or social edge strength (i.e. scale of trust or friendship), on the latent vectors U_i and U_f. The S_{Uif} is designed as a diagonal matrix to be learned, where the values of each diagonal element represent the variance of each dimension in the latent factor U_i. That says, S_{Uif} captures the strength of each latent dimension toward defining the edge strength. With such prior, we are now capable of modeling different "types" of social relationships, such as "A and B are friends because they are similar in latent dimension D1 and D2; while A and C are friends since they are similar in dimension D3, D4, and D5". Besides, to facilitate efficient learning, we choose Wishart distribution $\mathcal{W}$, the conjugate prior of precision matrix, with a fixed degree of freedom K and a scale matrix Λ_U shared by all S_{Uif} as regularization for S. As will be described later on, Λ_U will also be learned from data. Note that the zero-mean Gaussian spherical prior can be replaced with more general Gaussian priors such as those used in BPMF [25] or PPMF [26]. Nevertheless, here we apply the simple prior as the probabilistic view of $l2$-norm for fair comparison with previous social recommendation models widely using $l2$-norm regularization.

One major issue to address for implicit social recommender systems lies in the concern of efficiency. Because we need to learn S_{Uif} for a fully connected graph, it normally takes $O(N^2)$ time and space to evaluate all pairwise strengths between nodes. In Section 4.5.1, we will discuss how to address this issue.

3.3.2 SCP for Explicit User Social Recommendation

Explicit social recommendation can simply be treated as a special case in our model: when we are allowed to access explicit binary social networks together with rating data, we can simply define the edge sets E_U as the explicit social network. If the input network is undirected, each edge (i, f) is modelled as two directed edges $(i, f), (f, i) \in E_U$. Thus, instead of going through all pairs of nodes, if edge $(i, f) \notin E_U$ for users i, f, then we do not attempt to learn the similarity S_{Uif} and assume U_i is independent of U_f. Note that SCP does not utilize explicit edge strength even if it is provided. First, to our knowledge, there are very few publicly available explicit social network with edge strength information. Second, as mentioned previously, noisy edge strength could degrade the prediction performance, so we would rather learn it from data. Figure 5 marks the difference between incorporating an explicit social network and learning an implicit social network.

3.3.3 SCP for Item Social Networks

Since our model mines the implicit relationship between users, we can as well apply the same concept to mine the implicit relationship between items. Thus, the item social prior that is symmetric to (3) can be incorporated into our optimization procedure:

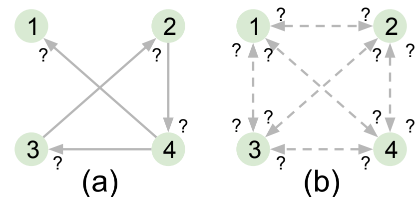

Fig. 5. Two scenarios of SCP learning a user social network. (a) If an explicit binary social network is accessible, then SCP learns the similarity (i.e. social strength) of each edge by rating data, regardless of external social strengths. (b) Without explicit social networks, SCP assumes that any pair of users have potential social relationships, and then learns every pairwise similarity. Note that one edge contains two directions to be learned different similarity values

$$\prod_{(j,g)\in E_V} \left[\mathcal{N}(U_j|U_g, S_{Vjg}^{-1})\mathcal{W}(S_{Vjg}|K, \Lambda_V)\right]^{\frac{1}{T_{Vj}}}. \tag{4}$$

Similar to the user social prior, the item social prior captures the similarity between two items by aligning their latent factors. Furthermore, every two items can be "similar" in different aspects captured by latent dimensions. For instance, movie A and B are similar because there is overlapping between the main actors. On the other hand, movie C and D are similar because they are talking about a similar topic.

3.3.4 Social Covariance Prior Overview

Integrating the components from previous sections, we now propose the general SCP model as follows:

$$\begin{aligned}
&\underset{U,V,S_U,S_V}{\arg\max} \prod_{i=1}^{N}\prod_{j=1}^{M} \mathcal{N}(R_{ij}|U_i^\top V_j, \sigma_R^2)^{\delta_{ij}} \\
&\prod_{i=1}^{N} \mathcal{N}(U_i|0,\sigma_U^2 I)^{1-b_U} \prod_{(i,f)\in E_U} \left[\mathcal{N}(U_i|U_f, S_{Uif}^{-1})\mathcal{W}(S_{Uif}|K,\Lambda_U)\right]^{\frac{b_U}{T_{Ui}}} \\
&\prod_{j=1}^{M} \mathcal{N}(V_j|0,\sigma_V^2 I)^{1-b_V} \prod_{(j,g)\in E_V} \left[\mathcal{N}(V_j|V_g, S_{Vjg}^{-1})\mathcal{W}(S_{Vjg}|K,\Lambda_V)\right]^{\frac{b_V}{T_{Vj}}},
\end{aligned} \tag{5}$$

where $b_U, b_V \in [0, 1]$ are balance parameters. They control the regularization ratio between the social prior and the zero-mean Gaussian spherical prior. Later on in our experiment, we will show that setting $b_U = b_V = 1$ yields better results, meaning that the original Gaussian spherical prior can be ignored. Figure 6 displays the graphical structure of SCP.

4 Model Learning

In this section, we will discuss how the variables in our model can be derived through the *mean-field Variational Expectation Maximization (VEM)* [28] method, one of the common Bayesian learning algorithms. The learning process and the prediction are described in Section 4.1 to 4.4. Section 4.5 describes several implementation tricks to improve the time and space efficiency. Finally we analyze the complexity of both time and space in Section 4.6.

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

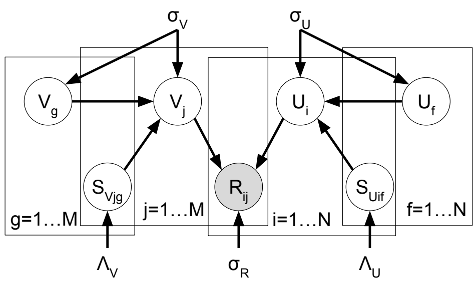

Fig. 6. Graphical model for Social Covariance Prior (SCP). In Section 4, we learn SCP using Variational Expectation Maximization (VEM) where E-step learns variational parameters θ' of random variables $\{U_i, V_j, S_{Ui}, S_{Vj}\}$ and M-step learns model parameters $\theta = \{\Lambda_U, \Lambda_V, \sigma_U^2, \sigma_V^2, \sigma_R^2\}$.

Stochastic Gradient Descent (SGD) is considered as a popular method to learn parameters in social recommendation. Its main advantage over VEM is to maximize the exact objective rather than the lower bound under VEM. We choose VEM because of two considerations. (1) Efficient computation. VEM saves significant efforts for parameter tuning. As regularization scales in SGD, overall four to five model parameters in SCP need to be carefully tuned, which is a time consuming process. VEM learns variational and model parameters to alleviate potential overfitting to validation data. (2) Efficient space usage. One of our major contributions is that based on VEM we propose space-saving tricks that reduce space complexity from $O(NK^2 + |E_U|K^2)$ to $O(NK)$ where $|E_U|$ is the number of edges, in order to eliminate the number of parameters that must be kept in memory during the learning process. Recall that when learning implicit social networks (Section 3.3.1), we have to learn a fully connected graph, and hence $O(NK^2 + |E_U|K^2) = O(N^2K^2)$. The trick is to exploit the fact of repeatedly taking V and E-steps in VEM learning. The old values of $O(N^2)$ social strength variational parameters x in the E-step can be instantly cast to the computation of M-step model variables, and hence we can release the $O(N^2K^2)$ space of x. Considering learning only diagonal elements for each $K \times K$ matrix parameter, the overall space complexity is from $O(N^2K^2)$ to $O(NK)$. The detail of two space-saving tricks is presented in Section 4.5.2. However the first trick cannot work under SGD, since all the old values of parameters x must be kept to to enable gradient descents $x' = x - \eta\frac{\partial O}{\partial x}$ with objective function O and learning rate η. Without the help of the space-saving trick, the SGD optimization needs to keep $O(N^2K)$ parameters in memory, which is infeasible to accept a large size of N users.

4.1 Overview

According to Bayesian treatment, we need to maximize the likelihood of of the observed ratings, given the parameters $\theta = \{\Lambda_U, \Lambda_V, \sigma_R^2, \sigma_U^2, \sigma_V^2\}$ and averaging over all possible values of the hidden variables $Z = \{U, V, S_U, S_V\}$:

$$p(R|\theta) = \int_Z p(R, Z|\theta)dZ,$$

where $p(R, Z|\theta)$ is represented in (5). Unfortunately, the integration incurs the intractability of the optimization. Instead, we can apply a tractable auxiliary probability $q(Z|\theta')$ *to maximize the lower bound of the log likelihood*. The lower bound can be derived from Jensen's inequality as follows:

$$\begin{aligned}\log p(R|\theta) &= \log \int_Z q(Z|\theta')\frac{p(R,Z|\theta)}{q(Z|\theta')}dZ = \log \mathbb{E}_{q(Z|\theta')}\left[\frac{p(R,Z|\theta)}{q(Z|\theta')}\right] \\ &\geq \mathbb{E}_{q(Z|\theta')}\left[\log\frac{p(R,Z|\theta)}{q(Z|\theta')}\right] \\ &= \mathbb{E}_{q(Z|\theta')}\Big[\log p(R|\theta)\Big] - \mathbb{E}_{q(Z|\theta')}\left[\log\frac{q(Z|\theta')}{p(Z|R,\theta)}\right] \\ &= \log p(R|\theta) - \mathbb{KL}\Big[q(Z|\theta')\|p(Z|R,\theta)\Big]. \end{aligned} \tag{6}$$

$\mathbb{E}_{q(Z|\theta')}(X)$ is the expected value of X over probability distribution $q(Z|\theta')$, and $\mathbb{KL}(p\|q)$ is the non-negative Kullback-Leibler divergence of two distributions p and q. To tighten the lower bound, we should have the divergence be 0, which implies $q(Z|\theta') = p(Z|R,\theta)$. However, the intractability of $p(Z|R,\theta)$ demands a tractable replacement. An approach to variational tractable $q(Z|\theta')$ is assuming all the involved hidden variables being independent of each other:

$$\begin{aligned} q(Z|\theta') = &\prod_{i=1}^{N}\mathcal{N}(U_i|\lambda_{Ui},\gamma_{Ui})\prod_{(i,f)\in E_U}\mathcal{W}(S_{Uif}|K+1,\Lambda_{Uif})^{\frac{b_U}{T_{Ui}}} \\ &\prod_{j=1}^{M}\mathcal{N}(V_j|\lambda_{Vj},\gamma_{Vj})\prod_{(j,g)\in E_V}\mathcal{W}(S_{Vjg}|K+1,\Lambda_{Vjg})^{\frac{b_V}{T_{Vj}}}, \end{aligned} \tag{7}$$

where $\theta' = \{\lambda_{Ui}, \lambda_{Vj}, \gamma_{Ui}, \gamma_{Vj}, \Lambda_{Uif}, \Lambda_{Vjg}\}$. One advantage is that there is no shared parameter among the hidden variables in $q(Z|\theta')$. The independence assumption is called *mean-field approximation*.

Now we have two sets of parameters: θ from the model distribution, and θ' from the variational distribution. VEM learns patterns by optimizing both sets of parameters that determine the presence of all probability distributions. Like classical expectation maximization (EM) algorithm, we iteratively execute variational *E-step* and *M-step*, updating θ' and θ respectively until a local optimum or predefined maximum number of iterations is reached.

Due to page limits and the symmetric forms of probability distributions of users and items, the following mathematical details focus on the derivation of user parameters. Readers can refer to our Supplemental Material for more derivation details of Section 4.2 and 4.3.

4.2 Variational E-step

In variational E-step, we have to optimize variational parameters θ' as fixing model parameters θ. Mean-field approach leads to a general principle [29]. It incurs convenience to derive θ' in $q(Z|\theta')$. For instance, for some user i, we have his or her variational distribution $q(U_i)$ as follows:

$$\log q(U_i|\theta') = \mathbb{E}_{-q(U_i|\theta')}\left[\log p(R, Z|\theta)\right] + C_0,$$

where $-q(U_i|\theta')$ represents the joint distribution of all the random variables Z except U_i, and C_0 absorbs the terms not relevant to U_i. We extend the expected value over independent variational random variables based on the mean-field assumption. The extension derivation refers to [30]. All the expected values as well as covariance matrices can be written as parameters in θ. In the end, we complete the square in order to write down a logarithmic form of multivariate

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

normal distribution with mean λ_{Ui} and covariance matrix γ_{Ui}. They are what we need to update at E-Step:

$$\gamma_{Ui} = \Bigg[\frac{1}{\sigma_R^2}\sum_{j=1}^{M}\delta_{ij}\Big(\lambda_{Vj}\lambda_{Vj}^\top + \gamma_{Vj}\Big) + (1-b_U)(\sigma_U^2)^{-1}I + \sum_{(i,f)\in E_U}\frac{b_U(K+1)}{T_{Ui}}\Lambda_{Uif} + \sum_{(f,i)\in E_U}\frac{b_U(K+1)}{T_{Uf}}\Lambda_{Ufi}\Bigg]^{-1}, \tag{8}$$

$$\lambda_{Ui} = \gamma_{Ui}\Bigg[\frac{1}{\sigma_R^2}\sum_{j=1}^{M}\delta_{ij}R_{ij}\lambda_{Vj} + \sum_{(i,f)\in E_U}\frac{b_U(K+1)}{T_{Ui}}\Lambda_{Uif}\lambda_{Uf} + \sum_{(f,i)\in E_U}\frac{b_U(K+1)}{T_{Uf}}\Lambda_{Ufi}\lambda_{Uf}\Bigg]. \tag{9}$$

We treat the similarity variables S_U the same ways. Here we write down the final equation:

$$\Lambda_{Uif} = \Big[(\lambda_{Ui}-\lambda_{Uf})(\lambda_{Ui}-\lambda_{Uf})^\top + \gamma_{Ui} + \gamma_{Uf} + \Lambda_U^{-1}\Big]^{-1}. \tag{10}$$

4.3 Variational M-step

This step takes care of the optimization of model parameters θ with fixed variational parameters θ'. We can perform the task by taking the derivative of the lower bound in (6), and find the closed-form solution of each parameter in θ.

We use σ_U^2 as example. The lower bound L in (6) is viewed as function of σ_U^2. Let $\frac{\partial L}{\partial \sigma_U^2} = 0$ and then obtain the closed-form solution:

$$\sigma_U^2 = \frac{1}{KN}\sum_{i=1}^{N}\Big(\|\lambda_{Ui}\|_2^2 + \mathrm{tr}(\gamma_{Ui})\Big). \tag{11}$$

At M-step, we update σ_U^2 using the above assignment. Following similar derivations, we have the update rules of other model parameters:

$$\Lambda_U = \frac{1}{KN}\sum_{(i,f)\in E_U}\frac{1}{T_{Ui}}(K+1)\Lambda_{Uif}, \tag{12}$$

$$\sigma_R^2 = \frac{1}{|R|}\sum_{i=1}^{N}\sum_{j=1}^{M}\delta_{ij}\Big[R_{ij}^2 - 2R_{ij}\lambda_{Ui}^\top\lambda_{Vj} + (\lambda_{Ui}^\top\lambda_{Vj})^2 + \mathrm{tr}(\gamma_{Ui}\gamma_{Vj}) + \lambda_{Ui}^\top\gamma_{Vj}\lambda_{Ui} + \lambda_{Vj}^\top\gamma_{Ui}\lambda_{Vj}\Big], \tag{13}$$

where $|R| = \sum_{i=1}^{N}\sum_{j=1}^{M}\delta_{ij}$ is the number of ratings in the training data.

4.4 Prediction

After the convergence of the VEM, we are able to predict an unobserved rating R_{ij} using the learned parameters. Without knowing the true rating likelihood, we apply the same strategy as M-step: finding R_{ij} to maximize the lower bound $L(R_{ij})$ of likelihood. This strategy can be done by letting $\frac{\partial L}{\partial R_{ij}} = 0$. The closed-form solution is used to predict unobserved ratings:

$$\hat{R}_{ij} = \lambda_{Ui}^\top\lambda_{Vj}. \tag{14}$$

4.5 Improving Efficiency

In this section we mention some practical tricks to improve the efficiency of SCP.

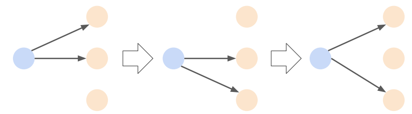

Fig. 7. Interpretation of randomized implicit social network. For each iteration (running E-step and M-step once), every user or item samples a small number of friends uniformly at random.

4.5.1 Time: Randomized Implicit Social Network

Generating implicit social connections is computationally expensive since we need to evaluate all possible pairwise connections. Previous works thus define heuristics to eliminate edges that are less likely to reflect strong social relations. For example, Ma et al. [17] leave merely those edges connecting two users who rate at least 10 identical items. However, such manually-crafted heuristics might not be suitable for every scenario, and it still costs $O(N^2)$ complexity to run the algorithm to determine which edge to be removed.

To address the efficiency issue, our idea is to do *random sampling* from a user's neighbors for each VEM iteration, as illustrated in Figure 7. To explain the trick, let us have a deeper analysis on the updating rules of our VEM model. In fact, there are three updating rules that require the enumeration of all possible pairs E_U. That is, a set of updated similarity scale matrices Λ_{Uif} for all neighbors f of user i is computed mainly for the updates of variational mean vector $\lambda_{Ui}, \lambda_{Vj}$ (shown in (9)), variational covariance matrix γ_{Ui}, γ_{Vj} (shown in (8)), and model scale matrix Λ_U, Λ_V (shown in (12)). Here we focus on the pairwise update terms in (8),

$$\sum_{(i,f)\in E_U}\frac{b_U(K+1)}{T_{Ui}}\Lambda_{Uif} + \sum_{(f,i)\in E_U}\frac{b_U(K+1)}{T_{Uf}}\Lambda_{Ufi} = \frac{b_U(K+1)}{N-1}\sum_{f=1,f\neq i}^{N}\Big(\Lambda_{Uif} + \Lambda_{Ufi}\Big). \tag{15}$$

The right-hand side of (15) reveals the mean of similarities between user i and all other users. Note that it is the mean value that we need. Thus, the law of large number tells us that for each VEM iteration, as long as we can sample sufficient amount of users $F \subseteq \{f : 1 \leq f \leq N, f \neq i\}$ uniformly at random, then the mean of similarities of the sampled neighbors will approach the original mean as follows:

$$\frac{b_U(K+1)}{N-1}\sum_{f=1,f\neq i}^{N}\Big(\Lambda_{Uif} + \Lambda_{Ufi}\Big) \approx \frac{b_U(K+1)}{|F|}\sum_{f\in F}\Big(\Lambda_{Uif} + \Lambda_{Ufi}\Big). \tag{16}$$

The same idea is applicable for the terms in (9) (12) as well since they all involve the estimation of mean values. In summary, at E-step, we only need to sample a constant number of neighbors for each user, which allows us to avoid the curse of evaluating a fully-connected graph.

4.5.2 Space: Saving Matrix Parameter

SCP is composed of three kinds of $K \times K$ matrix parameters: variational covariance matrices $\gamma_{Ui}, \gamma_{Vj} \forall i, j$, variational scale matrices $\Lambda_{Uif}, \Lambda_{Vjg} \forall i, j, f, g$, and model scale matrices Λ_U, Λ_V. There are overall NK^2 values in $\gamma_{Ui} \forall i$ and $|E_U|K^2$ values in $\Lambda_{Uif} \forall i, f$ to be trained, which incurs

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

certain burden on the optimization process. Furthermore, overfitting is also a concern since we have to train a number of parameters. A simple trick we have applied here is to make all the matrix parameters *diagonal*. In other words, all the off-diagonal elements are fixed to zero. Some previous works [26], [31] have also utilized such trick. When performing VEM, we need not to update off-diagonal elements at all. The diagonal trick enables us to save and train NK values in $\gamma_{Ui} \forall i$ and $|E_U|K$ values in $\Lambda_{Uif} \forall i, f$. Details are shown in our Supplemental Material.

In Section 4.5.1 we have described a trick to reduce the time complexity to linear using sampling techniques. It still requires $|E_U|K$ non-zero elements to store each of the variational scale matrices Λ_{Uif} and Λ_{Vjg}, and up to $O(|F|NK)$ storage for implicit social networks. Inspired by [27], we observe that for our task there is no need to store those variational scale matrices throughout the process. First, those variables are used only in the training stage. They play no role during the prediction stage (see Section 4.4). Second, in the VEM training stage, we observe that an updated Λ_{Uif} is always used in summation in (8), (9) and (12) later on. Hence, after Λ_{Uif} updates its value, we can use it in (8), (9), and (12) for the summation term right away, and then its space can be released since it is no longer needed. In the end, $O(|F|NK)$, the most dominant factor to space complexity, can be eliminated. In terms of space, our model becomes more scalable for big datasets after applying the tricks mentioned. Details are put in the Supplemental Material.

4.6 Complexity Analysis

Section 4.5.2 has concluded that keeping only diagonal matrix parameters takes only $O(NK + MK)$ in terms of space. Other parameters do not demand space larger than such. Consequently, the overall space complexity is constrained to $O(|R| + |E| + NK + MK)$, where $|R|$ is the number of ratings in the training data, $|E|$ is the number of edges if an explicit social networks is provided, or zero for the case of implicit recommendation.

As for running time, since the number of iterations needed for convergence varies for different datasets, we analyze the time complexity of one VEM iteration. For variational variable sets $\lambda_U, \lambda_V, \gamma_U$ and γ_V, the update rules (8) and (9) take total time $O(|R|K^2 + NK^3 + MK^3)$ including $O(K^3)$ time to generate the inverse matrix. In practice, K is usually small as the size of the latent factor (e.g., 10 in our experiments). With the trick mentioned in Section 4.5.2, the similarity-related variables Λ_U, Λ_V can be computed in $O(|E|K)$ time for explicit social networks. However for implicit social networks, the complexity becomes $O(|F|NK + |F|MK)$ where $|F|$ is the fixed number of sampled users or items, which we have discussed in Section 4.5.1. To update σ_R^2, it takes time $O(|R|K)$. Note that in practice $|F|$ and K are far smaller than $|R|, |E|, N$ and M. Therefore, our model is scalable since the time complexity is linear to $|R|, |E|, N$ or M.

5 Experiments

5.1 Experiment Settings

5.1.1 Datasets

To evaluate the proposed approach, we conduct experiments using four datasets (Epinion [2], Ciao [3], Flixster [4] and FilmTrust [5]) containing explicit user binary social networks and two datasets (MovieLens 1M [6] and Amazon [7]) as the representative of a large rating dataset without explicit social network. For each dataset, we discard the repeated ratings such that if a user rated an item more than once, then we keep only the latest rating. The statistics of the datasets are shown in Table 1. For each rating in each dataset, we subtract it by a constant mean of all ratings in that dataset. Note that for the largest dataset (Amazon), it takes about 8.5 hours for SCP to finish the computation with RMSE 1.178, while all competitors failed to finish the computation within 72 hours. Therefore we cannot really compare the results on Amazon dataset in the following section.

TABLE 1
Statistics of datasets in our experiments. "N/A" means no explicit user social network. Cold-start users have less than 20 ratings in a dataset.

Dataset	Ratings	Users	Items	Edges	Cold-start users	Relation
FilmTrust [32]	35494	1508	2071	1632	854	Trust
Epinions [33]	912441	22164	296277	354897	11362	Trust
Ciao [33]	282650	7375	105114	57544	3889	Trust
Flixster [34]	8196077	147612	48794	2538746	111120	Friendship
MovieLens 1M [35]	1000209	6040	3706	N/A	0	N/A
Amazon [36]	82677131	21176522	9874211	N/A	20594442	N/A

5.1.2 Competitors

Next, we introduce the following baseline and the state-of-the-art social recommendation models for comparison.

Probabilistic Matrix Factorization (PMF) [23]. As have described in Section 3.2, PMF is a popular model-based collaborative filtering approach without using social relationship. It serves as the baseline in our experiments.

NeuMF [37]. To our knowledge, NeuMF is the state-of-the-art recommender system that extends vanilla matrix factorization. It does not import explicit social networks or assume implicit social networks in its neural network structure. We compare SCP with NeuMF to show the performance benefit of learning implicit social networks.

Explicit Social Recommendation Models. We implemented several well-known explicit social recommender systems. The corresponding objective function for each is shown below:

Social Regularization (SocReg) [9]

$$\arg\min_{U,V} \frac{1}{2} \sum_{i=1}^{N} \sum_{j=1}^{M} \delta_{ij} \left(R_{ij} - U_i^\top V_j \right)^2$$

2. www.public.asu.edu/˜jtang20/datasetcode/truststudy.htm
3. www.public.asu.edu/˜jtang20/datasetcode/truststudy.htm
4. www.cs.ubc.ca/˜jamalim/datasets/
5. www.librec.net/datasets.html
6. grouplens.org/datasets/movielens/
7. jmcauley.ucsd.edu/data/amazon/links.html

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

$$+ \frac{\alpha_U}{2}\|U\|_2^2 + \frac{\alpha_V}{2}\|V\|_2^2 + \frac{\beta}{2}\sum_{(i,f)\in E_U}\omega_{if}\|U_i - U_f\|_2^2, \quad (17)$$

where ω_{if} denotes the edge strength between user i and f.

SocialMF [6]

$$\arg\min_{U,V} \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{M}\delta_{ij}\left(R_{ij} - U_i^\top V_j\right)^2$$
$$+ \frac{\alpha_U}{2}\|U\|_2^2 + \frac{\alpha_V}{2}\|V\|_2^2 + \frac{\beta}{2}\sum_{i=1}^{N}\|U_i - \sum_{f\in F_i}\omega_{if}U_f\|_2^2, \quad (18)$$

where F_i denotes the set of neighbors of user i.

TrustSVD [4]

$$\arg\min_{U,V,Y,W,b} \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{M}\delta_{ij}\left(R_{ij} - \hat{R}_{ij}(U,V,Y,W,b)\right)^2 + \frac{\beta}{2}C_1(U)$$
$$+ \frac{\beta}{2}\sum_{(i,f)\in E_U}(\omega_{if} - \hat{\omega}_{if}(U,W))^2 + \frac{\alpha}{2}C_0(U,V,Y,W,b), \quad (19)$$

where Y represents latent factors of items that a user rates in the training data, W denotes latent factors of users that a user trusts, b is the bias of users and items. Functions $\hat{R}, \hat{\omega}$ are designed to predict unobserved entries in R, ω respectively, and functions C_0, C_1 contain regularization terms.

PTPMF [11]

$$\arg\min_{U,V} \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{M}\delta_{ij}\left(R_{ij} - \hat{R}_{ij}(U^w, U^s, V, P)\right)^2$$
$$+ \frac{\alpha_U}{2}\|U\|_2^2 + \frac{\alpha_V}{2}\|V\|_2^2 + \frac{\beta}{2}\sum_{i=1}^{N}\|U_i^w - \sum_{f\in F_i}\omega_{if}U_f\|_2^2$$
$$+ \frac{\beta}{2}\sum_{i=1}^{N}\|U_i^s - \sum_{f\in F_i}\omega_{if}U_f\|_2^2, + \frac{\gamma}{2}C(P) \quad (20)$$

where U^w, U^s respectively imply latent factors of users interacting with friends of weak or strong social strengths. $\hat{R}$ denotes the rating estimate. Parameters P learn the user preferences between weak and strong social strengths. C is the regularization term.

Both SocReg and SocialMF are well-known explicit social recommendation works, while TrustSVD and PTPMF are the state-of-the-art solutions in this domain. Note that all the four competitors aim to minimize the squared error of rating prediction, which corresponds to maximizing likelihood of normal distributions over ratings in SCP. Some of the previous works in Section 2 are not considered since their objectives do not target at the minimization of squared error. We believe that it is fairer to compare all the experimented models with the same objective form.

Implicit social recommendation. For each of the five datasets, we follow a classical proposal in [17] to generate implicit user social networks. That is, the existence of an edge is confirmed if two users commonly rate at least 10 items and the Pearson Correlation Coefficient (PCC) between their ratings is larger than 0.5. Then the generated networks are fed into the above models (i.e. SocReg, SocialMF, TrustSVD and PTPMF) as the input social network to generate the recommendation results. Note that although [17] proposes the concept of dissimilar edges, it is not applied in our experiments since it is specifically designed for SocReg. We also implement the state-of-the-art proposal CUNE [19] to generate another implicit social network for each rating dataset. The approaches determine the friend set of a user with the top-50 highest cosine similarities between the network embeddings of the user and another. All the hyper-parameters of CUNE follow the default setting in [19]. Recall that SCP learns an implicit user social network optimization itself, therefore it does not use the generated implicit social network. We show an overview of our experiments using Table 2.

TABLE 2
Overview of our experiments on implicit social recommendation

Model	Two stages of implicit social recommendation: 1: Generating implicit social networks	2: Learning recommendation with implicit social networks
SocReg [9] SocialMF [6] TrustSVD [4] PTPMF [11]	Algorithms: (1) PCC [17] (2) CUNE [19]	SocReg SocialMF TrustSVD PTPMF
SCP	One-stage approach Learning implicit social networks and recommendation together	

All the competitor approaches are trained using stochastic gradient descent (SGD), as mentioned in the corresponding paper. In the original papers of most competitors, vanilla SGD is used with a fixed learning rate. To save tuning time, the learning rate is dynamically adapted by ADADELTA [38] where the insensitive parameters is set to $\rho = 0.95, \epsilon = 10^{-6}$, according to the original paper. NeuMF using Adam [39] is the exception for better performance, while it is not the case for the other competitors. We tune regularization parameters α, β by searching values in $\{10^p \mid p \in \mathbb{Z}, -4 \le p \le 1\}$, and let $\alpha_U = \alpha_V$. $\gamma = 10^{-5}$ following the experiments in PTPMF. For SCP parameters b_U, b_V, we tune them in the value set $\{0.1k \mid k \in \mathbb{Z}, 0 \le k \le 10\}$. We always fix the number of latent factors $K = 10$, which is also assumed by all the competitor models in their original papers. Finally, the number of sampled neighbors $|F| = 100$ for SCP.

5.1.3 Evaluation

As SCP and all the competitor models minimize the squared error as their objective, we report model performance using a common evaluation metric *Root Mean Squared Error (RMSE)*, which is monotonically increasing in the squared error. It is widely used for evaluating rating prediction in a recommender system. Smaller RMSE implies better prediction performance.

On the other hand, we also evaluate the ranking of the items for each user using Hit Rate and AUC, the Area under ROC curve (since some experiments have binary outcomes). For each user, we examine top 50 recommended items with the highest predicted ratings, following the setting in [11].

We use 5-fold cross validation, and report average results of the 5 folds. Moreover, we choose 10% of the training data as the validation data for each of the model. After each epoch of the VEM or SGD optimization, we evaluate RMSE on the validation data. If we observe RMSE decreasing by less than 10^{-4}, the model stops training as convergence is considered to be reached. Since SCP and all the baseline

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

models optimize MSE for rating prediction, it is proper to use RMSE to determine the timing of early stopping.

5.2 Experiment Results

Here we would like to evaluate five hypotheses (H1 to H7) to justify the effectiveness of our model. Furthermore, we conduct sensitivity analysis of parameters b_U, b_V in SCP.

5.2.1 Performance Comparisons

TABLE 3
RMSE and one-fold running time (seconds) of PMF using SGD and VEM methods; (*) denotes the RMSE being significantly lower than the other under 95% confidence level.

		FilmTrust	Epinions	Ciao	Flixster	MovieLens 1M
SGD	RMSE	0.8187	1.1506	1.0740	0.8744	0.8797
	Time	0.34	11.53	3.05	102.32	14.75
VEM	RMSE	0.8063	1.1078*	1.0610*	0.8474*	0.8582*
	Time	0.43	12.67	1.18	184.87	7.80

H1: VEM-based learning model produces competitive results in learning PMF comparing to SGD-based learning methods. One key idea of our model is to resort to VEM instead of SGD in optimization, since VEM allows us to embed a complex social prior and learn both latent user/item factors as well as implicit social co-variance together. As commented at the head of in Section 4, SCP cannot be simply optimized all the parameters of its social priors using SGD. Thus we think it is meaningful to compare whether VEM learns as good and fast as SGD in a traditional PMF setting, which is equivalent to SCP without social priors, to confirm VEM being a competitive alternative. Table 3 lists RMSE on PMF using both techniques. The results show that the VEM-based learning strategy is at least as good as SGD in terms of accuracy and efficiency, which strengthen our motivation of replacing SGD with VEM.

TABLE 4
Performance comparisons between recommender systems with explicit user social networks. (*) implies the statistically significant difference between top-1 and top-2 values under 95% confidence.

		SocReg	SocialMF	TrustSVD	PTPMF	SCP
FilmTrust	RMSE	0.818	0.819	0.809	0.821	**0.805**
	AUC	53.315%	53.272%	53.619%	53.024%	**53.784%**
	Hit Rate	5.385%	5.252%	3.899%	7.255%	**37.016%***
Epinions	RMSE	1.139	1.144	1.101	1.148	**1.091***
	AUC	56.757%	54.912%	**58.832%***	55.187%	57.594%
	Hit Rate	0.134%	0.265%	0.173%	0.218%	**2.687%***
Ciao	RMSE	1.046	1.046	1.025	1.071	**1.013***
	AUC	50.623%	50.673%	**54.079%***	51.503%	53.229%
	Hit Rate	1.432%	1.397%	1.494%	2.188%	**4.572%***
Flixster	RMSE	0.855	0.855	0.870	0.868	**0.848***
	AUC	54.492%	54.473%	54.584%	54.318%	**54.796%***
	Hit Rate	1.078%	1.305%	0.176%	1.983%	**6.143%***

H2: Given an explicit social network, our model can outperform existing explicit social recommender models. Since we claim our model as a general solution for both explicit and implicit social recommendations, we would like to first evaluate its effectiveness on explicit social recommendation. Table 4 demonstrates the performance differences of various social recommender systems and evaluation metrics. For all the four datasets with explicit user social networks, SCP is not significantly worse than any competitor in any metric. The results confirm that SCP can be regarded as an effective solution to the classical explicit social recommendation problem.

TABLE 5
Performance comparisons between recommender systems with explicit user social networks. Cold-start users have less than 20 ratings, while warm-start users rate at least 20 items in a dataset. * implies the statistically significant difference between top-1 and top-2 values under 95% confidence.

			SocReg	SocialMF	TrustSVD	PTPMF	SCP
FilmTrust	RMSE	Warm	0.809	0.809	**0.803**	0.815	0.803
		Cold	0.852	0.856	0.835	0.849	0.811
	AUC	Warm	54.957%	54.677%	55.176%	54.335%	**55.774%**
		Cold	51.549%	51.759%	51.945%	51.615%	51.639%
	Hit Rate	Warm	11.927%	11.621%	8.073%	15.841%	**52.141%***
		Cold	0.375%	0.375%	0.703%	0.679%	25.433%
Ciao	RMSE	Warm	1.040	1.040	1.023	1.058	**1.012**
		Cold	1.079	1.081	1.036	1.141	1.019
	AUC	Warm	50.882%	50.982%	**56.615%***	52.124%	55.605%
		Cold	50.361%	50.361%	51.507%	50.874%	50.820%
	Hit Rate	Warm	2.846%	2.760%	2.817%	4.320%	**8.044%***
		Cold	0.165%	0.175%	0.309%	0.278%	1.461%

In Table 4, SCP significantly outperforms other models in item ranking metrics, despite the fact that it optimizes RMSE. Also note that by No-Free-Lunch Theorem [40], an RMSE-optimized model is expected to have high performance in RMSE, but does not guarantee the same performance in other metrics. Therefore as observed in Table 4, for some datasets all models do not perform exceptionally in AUC and Hit Rate. Designing SPC for ranking-based objection will be the future work of us.

For further investigation, we separate all the users into two disjoint sets. The cold-start set C consists of users that have less than c ratings. Other users are classified into the warm-start set W. In our experiments, $c = 20$ following the setting in [37]. Due to page limit, we show cold/ware start comparion only on two datasets FilmTrust and Ciao in Table 5. The results show that SCP does extremely well on cold start users, implying that learning hidden social relationship can provide more information for sparse users and to boost the prediction outcomes. By Table 1, there are 57% and 53% of users classified into the cold-start set in FilmTrust and Ciao. The overall results in Table 4 thus show large value differences between SCP and other models.

H3: Without a given social network, our model can still outperform the other implicit social recommender solutions. Now we assume that the explicit user social networks are not available for training a model. For competitor models, we generate implicit user social networks as presented in 5.1.2. Note that SCP does not need the generated network since the approach assumes a fully-connected implicit social network and learns the edge strength automatically. Table 6 shows the performance comparisons. The results imply that SCP significantly outperforms the competitors via either PCC or CUNE in almost all metrics.

Table 6 also shows the running time of each model applied to implicit social recommendation [8]. While both PCC and CUNE requires $O(N^2)$ user pair comparisons to determine artificial social edges, SCP uses a random

8. Experiment machine: Eight Intel(R) Xeon(R) CPU X5570, 2.93GHz.

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

TABLE 6
Performance and one-fold running time (seconds) comparisons between recommender systems with implicit user social networks. * implies the statistically significant difference between top-1 and top-2 values under 95% confidence.

			SocReg	SocialMF	TrustSVD	PTPMF	SCP
FilmTrust	RMSE	PCC	0.817	0.815	0.805	0.823	**0.798**
		CUNE	0.816	0.817	0.814	0.823	
	AUC	PCC	53.295%	53.417%	54.005%	52.975%	**54.169%**
		CUNE	53.382%	53.206%	53.732%	53.059%	
	Hit Rate	PCC	5.040%	4.191%	3.806%	3.992%	**41.260%***
		CUNE	3.024%	5.053%	3.899%	2.215%	
	Time	PCC	2.99	2.57	34.97	14.32	**1.19**
		CUNE	98.61	98.69	160.15	275.74	
Epinions	RMSE	PCC	1.144	1.099	1.102	1.123	**1.063***
		CUNE	1.098	1.123	1.104	1.117	
	AUC	PCC	54.909%	57.054%	58.661%	56.463%	**60.101%***
		CUNE	57.982%	55.869%	59.397%	56.162%	
	Hit Rate	PCC	0.278%	0.164%	0.172%	2.133%	**4.135%***
		CUNE	0.079%	0.184%	0.519%	0.160%	
	Time	PCC	216.52	240.07	3040.72	301.50	**99.84**
		CUNE	2173.74	2086.75	6266.55	7948.12	
Ciao	RMSE	PCC	1.046	1.018	1.022	1.036	**0.986***
		CUNE	1.019	1.049	1.026	1.043	
	AUC	PCC	50.565%	52.501%	54.056%	51.417%	**54.492%***
		CUNE	52.617%	51.623%	54.195%	52.037%	
	Hit Rate	PCC	1.473%	0.773%	1.407%	1.334%	**7.064%***
		CUNE	1.730%	0.849%	1.538%	0.697%	
	Time	PCC	27.82	36.14	698.41	57.93	**21.69**
		CUNE	549.74	553.55	1653.94	2732.47	
Flixster	RMSE	PCC	0.862	0.855	0.869	0.871	**0.846***
		CUNE	0.855	0.854	0.892	0.861	
	AUC	PCC	54.493%	54.532%	54.560%	54.237%	**54.935%***
		CUNE	54.516%	54.556%	54.546%	54.495%	
	Hit Rate	PCC	0.075%	1.307%	0.045%	1.058%	**8.543%***
		CUNE	1.131%	1.141%	0.144%	0.297%	
	Time	PCC	11188.65	11504.58	86281.33	17693.54	**306.63**
		CUNE	75554.39	75660.52	168959.38	107468.44	
MovieLens 1M	RMSE	PCC	0.876	0.871	0.883	0.875	**0.855***
		CUNE	0.871	0.868	0.874	0.867	
	AUC	PCC	73.209%	74.305%	74.252%	73.525%	**75.489%***
		CUNE	73.879%	74.338%	74.153%	74.609%	
	Hit Rate	PCC	10.377%	8.735%	35.950%	21.026%	**74.073%***
		CUNE	12.520%	12.669%	64.487%	8.997%	
	Time	PCC	158.25	174.03	3219.03	1168.46	**11.44**
		CUNE	778.33	786.51	3832.90	6625.07	

sampling tricks (Section 4.5.1) to alleviate the computation burden. Consequently SCP is faster than all the competitors in implicit social recommendation. The experiment results reflect that SCP incorporates dense implicit social strength learning without sacrificing efficiency.

H4: Learning implicit item networks are usually useful. In Section 3.3.3, we propose extended SCP to consider implicit item social networks. Here we want to verify whether learning both the user and the item networks together can indeed boost the performance. The results are presented in Table 7. We realize that it does reach the lowest RMSE values in all datasets if learning of both user and item social networks, but not the case if the ranking metric is used. Concluded from the present experiments, learning item social networks is effective if we aim to boost rating prediction (RMSE), but there is the potential cost of lower item ranking performance.

H5: Learning social strengths can improve recommendation of SCP. To verify the effectiveness of multi-dimensional edge strength learning in SCP, we fix all the edge strengths 1. The evaluation of this variant is listed in Table 7. It shows the importance of learning social strengths for each edge.

TABLE 7
Performance comparisons between recommender systems without social networks (PMF and NeuMF) and SCP exploiting: (EU) explicit user social networks; (IU) implicit user social networks; (IUI) implicit user and item social networks; (IU_S1) implicit user social networks with all fixed edge strengths 1. * implies the statistically significant difference between top-1 and top-2 values under 95% confidence. N/A denotes no explicit social network in the dataset.

		PMF	NeuMF	SCP_EU	SCP_IU	SCP_IUI	SCP_IU_S1
FilmTrust	RMSE	0.820	0.804	0.805	0.798	**0.786***	0.801
	AUC	53.589%	53.664%	53.784%	54.169%	**54.497%**	54.478%
	Hit Rate	5.332%	19.934%	37.016%	**41.260%**	24.801%	25.862%
Epinions	RMSE	1.151	1.141	1.091	1.063	**1.037***	1.073
	AUC	53.648%	57.644%	57.594%	60.101%	**61.232%***	59.964%
	Hit Rate	0.424%	1.696%	2.687%	4.135%	**4.360%**	3.623%
Ciao	RMSE	1.046	1.018	1.013	0.986	**0.958***	0.992
	AUC	50.750%	54.518%	53.229%	54.492%	**55.160%***	54.524%
	Hit Rate	1.429%	3.186%	4.572%	**7.064%**	6.362%	6.986%
Flixster	RMSE	0.862	0.851	0.848	0.846	**0.837***	0.846
	AUC	54.363%	54.669%	54.796%	54.935%	**54.964%**	54.919%
	Hit Rate	1.399%	8.567%	6.143%	8.543%	**10.850%***	2.571%
MovieLens 1M	RMSE	0.883	0.879	N/A	0.855	**0.851***	0.855
	AUC	74.063%	73.839%	N/A	75.489%	**75.667%**	75.482%
	Hit Rate	29.281%	66.030%	N/A	74.073%	**77.245%***	74.106%

H6: Learning implicit social networks does benefit recommender systems. Also observing the results in Table 7, we discover that implicit social networks commonly lead to higher performance than explicit social networks for SCP. It is consistent with the claim in [19] that learning implicit social networks can even be a better choice than gathering information about explicit social networks. This table also shows that SCP outperforms the state-of-the-art model NeuMF that does not consider social information at all.

TABLE 8
Mean average precision of each user's top-k similar implicit neighbors of the Pearson correlation coefficient-defined network (PCC), CUNE-defined network (CUNE) and the SCP-learned network (SCP).

		$k = 10$	$k = 20$	$k = 50$	$k = 100$	$k = 200$
FilmTrust	PCC	0.350%	0.385%	0.401%	0.401%	0.401%
	CUNE	0.589%	0.684%	0.813%	0.813%	0.813%
	SCP	**1.174%**	**1.371%**	**1.525%**	**1.621%**	**1.713%**
Ciao	PCC	0.013%	0.013%	0.014%	0.014%	0.014%
	CUNE	0.065%	0.084%	0.115%	0.115%	0.115%
	SCP	**0.166%**	**0.209%**	**0.277%**	**0.340%**	**0.419%**

H7: SCP can discover hidden user social connection better than the human-defined implicit user social network. Although learning hidden social connections between people is not the goal of this paper, it is not hard to realize that using Equation 3 we can generate the probability that the latent factors of two users are similar as a heuristic to infer social connections. Here we want to verify whether the discovered network matches the true network better than the one identified by competitors. To verify our hypothesis, we utilize Mean Average Precision (MAP) to compare three uncovered implicit user social networks: one comes from SCP learning, the others come from the competitor's solutions using Pearson correlation coefficient and CUNE as described in Section 5.1.2. We report the comparison using two smallest datasets, FilmTrust and Ciao, with explicit user

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

social networks as ground truth. Table 8 lists the MAP comparisons of k most similar neighbors of each user in both user implicit social networks. Results reveal that in all cases, our model outperforms the competitor in revealing the connections between users.

5.2.2 Parameter Sensitivity

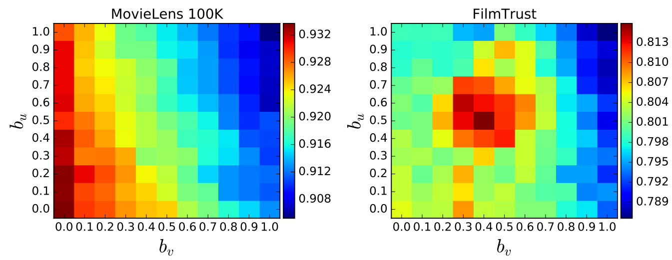

Fig. 8. RMSE in terms of different parameter combinations for datasets MovieLens 100K and FilmTrust.

As have been claimed in Section 3 to 4, our model automatically learns almost all the hidden variables from data. Nevertheless, to balance the contribution between Gaussian spherical prior and social prior, there remains two parameters b_U, b_V that still need to be determined through cross validation. Here we show how these two parameters influence the overall performance to suggest reasonable assignments for them.

To learn implicit social networks in SCP, from 0 to 1, we experiment different combinations of b_U, b_V settings, and then draw the RMSE distributions in Figure 8. Due to the page limit, we present the distributions for only implicit social recommendation with the incorporation of user and item social networks on MovieLens 100K, sharing the same reference as MovieLens 1M, and FilmTrust. The resulting distribution shown in the left figure (i.e., MovieLens 100K) is different from that in the right (i.e., FilmTrust), implying distinct spectrum of social influences between the two datasets. For example, for MovieLens 100K, the smooth distribution simply indicates that the overall performance does improve with more social influence factors considered, where the worse performance lies in the plane where only smaller social influence exists. For FilmTrust, it seems that the worst performance lies in the middle where these two factors are close to 0.5. We also discover that in general setting b_U close to 1 yields good results. In fact $b_U = b_V = 1$ generates the best results for all the datasets we have tested. Such finding supports the idea of eliminating Gaussian spherical priors because of the strong regularization capability of fully-connected implicit social networks. Without Gaussian priors, our model gains more freedom to train background social relations given rating information. That is, if the users do not have time to tune the parameters using validation dataset, we suggest the above values as default for implicit social recommendation.

5.3 Discussions

The most surprising result in our experiments is that SCP learning implicit social networks outperforms exploiting explicit social networks. It might seem counter intuition at the first glance, since explicit social networks bring additional user attributes other than ratings. Actually our conclusion is the same as [19], which also obtains higher performance of implicit social recommendation. In Section 1 we have summarized that [19] claims three potential drawbacks of explicit social recommendation. Following we review the three points to explain the strength of SPC.

Observation 1: Explicit social networks as adjacency matrices could be too sparse to boost recommendation performance. However, SCP assumes and learns a fully connected implicit social network, its adjacency matrix shall be dense and more information can be provided.

Observation 2: Cold-start users have fewer friends in explicit social networks. Since SCP implicit social networks are fully connected, there are $(N-1)$ friends for each user, no matter whether it belongs to cold start or not.

Observation 3: There is large difference in rating distributions between active users and their friends. SCP learns an individual edge strength for each latent factor, and hence allows two friends to be 'similar' only in certain latent dimensions.

6 Conclusion

Probabilistic matrix factorization has been a successful machine learning model toward recommender systems. Based on the social correlation intuition, we propose a new approach to incorporate social network information into probabilistic matrix factorization. We list our contributions again:

Contribution 1: In terms of effectiveness, we successfully build a joint model simultaneously to learn factorized matrices and social network structures. Experiments support that our new approach outperforms previous works that either focus on explicit social recommendation or implement implicit social recommendation in two separate stages.

Contribution 2: Distinct from learning a shared social strength, our work allows learning an individual social strength for each latent factor. We believe that the multi-dimensional social strength learning can benefit the overall recommendation quality.

Contribution 3: In terms of efficiency, to address the scalability problem resulting from fully connected implicit social networks, we propose several practical tricks during the learning process so the complexity can be reduced from quadratic to linear.

Future works consist of two parts. First, we would like to investigate how the content information can be incorporated into the existing model for a hybrid recommender system. Second, we will also consider bringing into our model the temporal information from both social and ratings sides to further boost the performance.

References

[1] A. J. B. Chaney, D. M. Blei, and T. Eliassi-Rad, "A probabilistic model for using social networks in personalized item recommendation," in *Proc. of ACM RecSys*, 2015, pp. 43–50.

[2] H. Fang, Y. Bao, and J. Zhang, "Leveraging decomposed trust in probabilistic matrix factorization for effective recommendation," 2014.

[3] J. A. Golbeck, "Computing and applying trust in web-based social networks," Ph.D. dissertation, 2005.

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

[4] G. Guo, J. Zhang, and N. Yorke-Smith, "Trustsvd: Collaborative filtering with both the explicit and implicit influence of user trust and of item ratings," 2015.

[5] M. Jamali and M. Ester, "Trustwalker: A random walk model for combining trust-based and item-based recommendation," in *Proc. of ACM SIGKDD*, 2009, pp. 397–406.

[6] M. Jamali and M. Ester, "A matrix factorization technique with trust propagation for recommendation in social networks," in *Proc. of ACM RecSys*, 2010, pp. 135–142.

[7] H. Ma, H. Yang, M. R. Lyu, and I. King, "Sorec: Social recommendation using probabilistic matrix factorization," in *Proc. ACM CIKM*, 2008, pp. 931–940.

[8] H. Ma, I. King, and M. R. Lyu, "Learning to recommend with social trust ensemble," in *Proc. ACM SIGIR*, 2009, pp. 203–210.

[9] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King, "Recommender systems with social regularization," in *Proc. ACM WSDM*, 2011, pp. 287–296.

[10] P. Massa and P. Avesani, "Trust-aware recommender systems," in *Proc. of ACM RecSys*, 2007, pp. 17–24.

[11] X. Wang, S. C. Hoi, M. Ester, J. Bu, and C. Chen, "Learning personalized preference of strong and weak ties for social recommendation," ser. WWW '17, 2017.

[12] X. Xin, I. King, H. Deng, and M. R. Lyu, "A social recommendation framework based on multi-scale continuous conditional random fields," in *Proc, ACM CIKM*, 2009, pp. 1247–1256.

[13] B. Yang, Y. Lei, D. Liu, and J. Liu, "Social collaborative filtering by trust," in *Proc. IJCAI*, 2013, pp. 2747–2753.

[14] S. Fazeli, B. Loni, A. Bellogin, H. Drachsler, and P. Sloep, "Implicit vs. explicit trust in social matrix factorization," in *Proc. ACM RecSys*, 2014, pp. 317–320.

[15] G. Guo, J. Zhang, D. Thalmann, A. Basu, and N. Yorke-Smith, "From ratings to trust: An empirical study of implicit trust in recommender systems," in *Proc. ACM SAC*, 2014, pp. 248–253.

[16] C. Lin, R. Xie, X. Guan, L. Li, and T. Li, "Personalized news recommendation via implicit social experts," *Information Sciences*, vol. 254, pp. 1 – 18, 2014.

[17] H. Ma, "An experimental study on implicit social recommendation," in *Proc. of ACM SIGIR*, 2013, pp. 73–82.

[18] H. Ma, I. King, and M. R. Lyu, "Learning to recommend with explicit and implicit social relations," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, pp. 29:1–29:19, May 2011.

[19] C. Zhang, L. Yu, Y. Wang, C. Shah, and X. Zhang, "Collaborative user network embedding for social recommender systems," ser. SDM '17, 2017.

[20] H. Ma, "On measuring social friend interest similarities in recommender systems," in *Proc. of ACM SIGIR*, 2014, pp. 465–474.

[21] C.-m. Au Yeung and T. Iwata, "Strength of social influence in trust networks in product review sites," in *Proc. of ACM WSDM*, 2011, pp. 495–504.

[22] G. Guo, J. Zhang, D. Thalmann, and N. Yorke-Smith, "Etaf: An extended trust antecedents framework for trust prediction," in *Proc. of IEEE/ACM ASONAM*, 2014, pp. 540–547.

[23] A. Mnih and R. R. Salakhutdinov, "Probabilistic matrix factorization," in *Proc. of NIPS*, 2008, pp. 1257–1264.

[24] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, Mar. 2003.

[25] R. Salakhutdinov and A. Mnih, "Bayesian probabilistic matrix factorization using markov chain monte carlo," in *Proc. of ICML*, 2008, pp. 880–887.

[26] H. Shan and A. Banerjee, "Generalized probabilistic matrix factorizations for collaborative filtering," in *Proc. of IEEE ICDM*, 2010, pp. 1025–1030.

[27] Y. J. Lim and Y. W. Teh, "Variational bayesian approach to movie rating prediction," in *Proc. of KDD Cup and Workshop*, 2007.

[28] J. M. Bernardo, M. J. Bayarri, J. O. Berger, A. P. Dawid, D. Heckerman, A. F. M. Smith, M. W. (eds, M. J. Beal, and Z. Ghahramani, "The variational bayesian em algorithm for incomplete data: with application to scoring graphical model structures," 2003.

[29] E. P. Xing, M. I. Jordan, and S. Russell, "A generalized mean field algorithm for variational inference in exponential families," in *Proc. of UAI*, 2003, pp. 583–591.

[30] K. B. Petersen and M. S. Pedersen, "The matrix cookbook," nov 2012.

[31] H. Jing, A.-C. Liang, S.-D. Lin, and Y. Tsao, "A transfer probabilistic collective factorization model to handle sparse data in collaborative filtering," in *Proc. of IEEE ICDM*, 2014, pp. 250–259.

[32] G. Guo, J. Zhang, and N. Yorke-Smith, "A novel bayesian similarity measure for recommender systems," in *Proc. of IJCAI*, 2013, pp. 2619–2625.

[33] J. Tang, H. Gao, H. Liu, and A. D. Sarma, "eTrust: Understanding trust evolution in an online world," in *Proc. of ACM SIGKDD*, 2012.

[34] M. Jamali, "Flixster dataset," http://www.cs.ubc.ca/~jamalim/datasets/.

[35] GroupLens, "Movielens datasets," http://grouplens.org/datasets/movielens/.

[36] R. He and J. McAuley, "Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering," ser. WWW '16, 2016.

[37] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," ser. WWW '17, 2017.

[38] M. D. Zeiler, "ADADELTA: An Adaptive Learning Rate Method," *ArXiv e-prints*, Dec. 2012.

[39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR '15*, 2015.

[40] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Transactions on Evolutionary Computation*, pp. 67–82, 1997.

Chin-Chi Hsu received the master degree from the Department of Computer Science and Information Engineering, National Taiwan University in 2014. He is a research assistant in the Institute of Information Science, Academia Sinica, Taiwan. His research domains include machine learning, data mining and social network analysis. In particular, he researched time series mining, social recommendation and node importance ranking in networks.

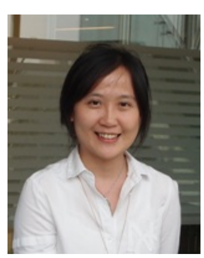

Mi-Yen Yeh is currently Associate Research Fellow of Institute of Information Science (and Research Center for IT Innovation under joint appointment) at Academia Sinica, Taiwan. She received her B.S. and Ph.D. degrees in Electrical Engineering from National Taiwan University, Taiwan, in 2002 and 2009, respectively. Her main research area is on data mining and databases, with a specific focus on mining ordered data, social network analysis, and data management on non-volatile memory. She received the best paper award (in system software and security) of the 28th annual ACM Symposium on Applied Computing (SAC 2013), Distinguished Postdoctoral Fellowship in Academia Sinica, and Research Exploration Award in Pan Wen Yuan Foundation.

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TKDE.2018.2821174, IEEE Transactions on Knowledge and Data Engineering

Shou-De Lin is currently a full professor in the CSIE department of National Taiwan University. He holds a BS in EE department from National Taiwan University, an MS-EE from the University of Michigan, and an MS in Computational Linguistics and PhD in Computer Science both from the University of Southern California. He leads the Machine Discovery and Social Network Mining Lab in NTU. Before joining NTU. Prof. Lin's research includes the areas of machine learning and data mining, social network analysis, and natural language processing. His international recognition includes the best paper award in IEEE Web Intelligent conference 2003, Google Research Award in 2007, Microsoft research award in 2008, merit paper award in TAAI 2010, best paper award in ASONAM 2011, US Aerospace AFOSR/AOARD research award winner for 5 years. He is the all-time winners in ACM KDD Cup, leading or co-leading the NTU team to win 5 championships. He also leads an NTU research team to win WSDM Cup 2016. He has served as the senior PC for SIGKDD and area chair for ACL. He is currently the associate editor for International Journal on Social Network Mining, Journal of Information Science and Engineering, and International Journal of Computational Linguistics and Chinese Language Processing.

1041-4347 (c) 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.