Recommendation with Social Dimensions

Jiliang Tang[†], Suhang Wang[‡], Xia Hu[*],
Dawei Yin[†], Yingzhou Bi[⋆], Yi Chang[†] and Huan Liu[‡]
[*]Texas A&M University, hu@cse.tamu.edu
[†]Yahoo Labs, {jlt,daweiy,yichang}@yahoo-inc.com
[‡] Arizona State University, {suhang.wang, huan.liu}@asu.edu
[⋆] Guangxi Teachers Education University, yingzhou.bi@gmail.com

Abstract

The pervasive presence of social media greatly enriches online users' social activities, resulting in abundant social relations. Social relations provide an independent source for recommendation, bringing about new opportunities for recommender systems. Exploiting social relations to improve recommendation performance attracts a great amount of attention in recent years. Most existing social recommender systems treat social relations homogeneously and make use of direct connections (or strong dependency connections). However, connections in online social networks are intrinsically heterogeneous and are a composite of various relations. While connected users in online social networks form groups, and users in a group share similar interests, weak dependency connections are established among these users when they are not directly connected. In this paper, we investigate how to exploit the heterogeneity of social relations and weak dependency connections for recommendation. In particular, we employ social dimensions to simultaneously capture heterogeneity of social relations and weak dependency connections, and provide principled ways to model social dimensions, and propose a recommendation framework SoDimRec which incorporates heterogeneity of social relations and weak dependency connections based on social dimensions. Experimental results on real-world data sets demonstrate the effectiveness of the proposed framework. We conduct further experiments to understand the important role of social dimensions in the proposed framework.

Introduction

Recommender systems have been proven to be effective in mitigating the information overload problem (Herlocker et al. 1999; Sarwar et al. 2001; Su and Khoshgoftaar 2009). The increasing popularity of social media allows online users to participate in online activities which produce rich social relations, such as friendships in Facebook and trust relations in Epinions. Social relations provide an independent source for recommendation and bring about new opportunities for recommender systems. Exploiting social relations for recommendation attracts more and more attention in recent years (Golbeck 2006; Massa and Avesani 2007; Jamali and Ester 2009; Jiang et al. 2012). Their common assumption is that a user's preference is similar to or influenced by their directly connected friends, which can be explained by social correlation theories such as homophily (Miller McPherson and Cook 2001) and social influence (Marsden and Friedkin 1993).

Copyright © 2016, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

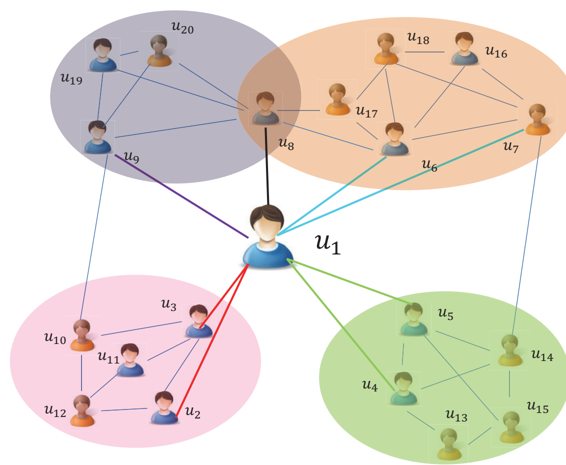

Figure 1: An illustration of u_1's social connections and the groups that she participates in.

Most existing social recommender systems treat a user's connections homogeneously and make use of her direct connections. However, connections in online social networks are intrinsically heterogeneous and are a composite of various relations (Tang and Liu 2009a; Sun and Han 2012; Tang, Gao, and Liu 2012). Users in online social networks form groups, and they have similar interests with other users in the same group although they may not directly connect. Figure 1 illustrates an example about u_1's social relations with $\{u_2, u_3, \ldots, u_9\}$, and the groups that she participates in. The user u_1 may treat her social relations differently in different domains (Tang, Gao, and Liu 2012). For example, u_1 may seek suggestions about "Sports" from $\{u_2, u_3\}$, but ask for recommendation about "Electronics" from $\{u_4, u_5\}$. Connected users in online social networks form groups where there are more connections among users within groups than among those between groups (Newman 2005). For example, in Figure 1, $\{u_1, u_2, u_3, u_{10}, u_{11}, u_{12}\}$ form a group while $\{u_1, u_3, u_5, u_{13}, u_{14}, u_{15}\}$ form another group. According to social correlation theories, similar users interact at higher rates than dissimilar ones; thus, users in the same group are likely to share similar preferences, establishing weak dependency connections when they are not directly connected (Tang and Liu 2009a). Note that in this paper we name the direct connections, e.g., relations between u_1 and $\{u_2, \ldots, u_9\}$ as strong dependency connections, and the connections within groups excluding the direct connections as weak dependency connections like relations between u_1 and $\{u_{10}, u_{11}, \ldots, u_{20}\}$. Weak dependency connections can provide important context information about users' interests (Tang, Chang, and Liu 2014), which have proven to be useful in job hunting (Granovetter 1973), the diffusion of ideas (Granovetter 1983), knowledge transfer (Levin and Cross 2004) and relational learning (Tang and Liu 2009a), while rarely investigated in recommendation.

The heterogeneity of social relations and weak dependency connections provide a new perspective for social recommender systems, and also present new challenges. In this paper, we investigate: (1) how to exploit the heterogeneity of social relations and weak dependency connections simultaneously, and (2) how to model them mathematically and then take advantage of them for recommendation. In an attempt to address these two challenges, we propose a novel recommendation framework SoDimRec which models the heterogeneity of social relations and weak dependency connections simultaneously based on social dimensions. Our contributions are summarized below:

- Employing social dimensions to simultaneously capture heterogeneity of social relations and weak dependency connections;
- Providing a principled way to model social dimensions for recommendation;
- Proposing a novel recommendation framework SoDimRec which incorporates heterogeneity of social relations and weak dependency connections based on social dimensions; and
- Evaluating the proposed framework extensively using real-world datasets to understand the working of the proposed framework.

Social Dimensions for Heterogeneity and Weak Dependency Connections

We first introduce notations used in this paper. We use $\mathbf{A}^i$ and $\mathbf{A}_j$ to represent the i-th column and the j-th row of $\mathbf{A}$, respectively. Let $\mathcal{U} = \{u_1, u_2, \ldots, u_n\}$ and $\mathcal{V} = \{v_1, v_2, \ldots, v_m\}$ be the sets of n users and m items respectively. We assume that $\mathbf{R} \in \mathbb{R}^{n \times m}$ is the user-item rating matrix. If u_i gives a rating to v_j, $\mathbf{R}_{ij}$ is the rating, otherwise we employ 0 to represent the unknown rating from u_i to v_j. Users can establish social relations to each other, and we use $\mathbf{T} \in \mathbb{R}^{n \times n}$ to denote user-user social relations where $\mathbf{T}_{ij}$ is the strength if u_j has a connection to u_i, and zero otherwise.

Social relations in social media are heterogeneous and different relations are mixed together. Recently, social dimensions are proposed to address this heterogeneity (Tang and Liu 2009b; 2009a) and reveal the potential affiliations of users in social networks based on network connectivity. A user can be involved in different social dimensions, corresponding to different types of relations (Tang and Liu 2009b). For example, in Figure 1, u_1 is in four social dimensions and she may have different types of relations with $\{u_2, u_3\}$, $\{u_4, u_5\}$ $\{u_6, u_7, u_8\}$ and $\{u_8, u_9\}$, respectively. According to social correlation theories, users with similar preferences interact at higher rates than dissimilar ones. Therefore extracting social dimensions is to identify a group of users who interact more frequently with each other, which boils down to a classical community detection problem (Tang and Liu 2009b; Leskovec, Huttenlocher, and Kleinberg 2010). Users in the same social dimension share similar preferences and weak dependency connections are established within a social dimension. Hence social dimensions can be used to simultaneously cover the heterogeneity of social relations and weak dependency connections, which correspondingly answers the first question in the introduction section.

Since a user such as u_1 in Figure 1 can be involved in multiple social dimensions, we adopt an overlapping community detection algorithm based on NMF to extract social dimensions (Wang et al. 2011). Let $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_c\}$ be the set of c social dimensions. Let $\mathbf{G} \in \mathbb{R}^{n \times c}$ be the social dimension indicator matrix where $\mathbf{G}_{ik} = 1$ indicates that u_i belongs to $\mathcal{G}_k$ and zero otherwise. With social dimensions, we can define weak dependency connections between two users in the same social dimension who do not connect directly. For $u_i \in \mathcal{G}_k$, we use $\mathcal{H}_{ik}$ to denote the set of users with whom u_i establishes weak dependence connections in $\mathcal{G}_k$, which is formally defined as

$$\mathcal{H}_{ik} = \{u_j | u_j \in \mathcal{G}_k \text{ and } \mathbf{T}_{ij} = 0\}, \tag{1}$$

Social dimensions have been extensively studied with proven results. The significance of the introduction of social dimensions is two-fold. First, social dimensions can naturally incorporate the heterogeneity of social relations and weak dependency connections, and allow us to investigate them simultaneously. Second, modeling social dimensions paves a way to incorporate the heterogeneity of social relations and weak dependency connections into a coherent framework for recommendation. In the following section, we will investigate how to model heterogeneity and weak dependence connections via social dimensions, and introduce the proposed recommendation framework, correspondingly to answer the second question in the introduction section.

A Recommendation Framework Based on Social Dimensions

Before modeling, we first introduce a state-of-the-art recommender system based on matrix factorization as our basic model. Matrix factorization techniques have been widely employed for recommendation (Salakhutdinov and Mnih

2008; Koren 2008; Jamali and Ester 2010; Ma et al. 2011), and they assume that a few latent patterns influence user rating behaviors and perform a low-rank matrix factorization on the user-item rating matrix. Let $\mathbf{U}_i \in \mathbb{R}^{1\times K}$ and $\mathbf{V}_j \in \mathbb{R}^{1\times K}$ be the user preference vector for u_i and item characteristic vector for v_j respectively, where K is the number of latent factors. Matrix factorization based recommender systems solve the following problem:

$$\min_{\mathbf{U},\mathbf{V}} \sum_{i=1}^{n}\sum_{j=1}^{m} \mathbf{W}_{ij}(\mathbf{R}_{ij} - \mathbf{U}_i\mathbf{V}_j^\top)^2 + \alpha(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2),$$

where $\mathbf{U} = [\mathbf{U}_1^\top, \mathbf{U}_2^\top, \dots, \mathbf{U}_n^\top]^\top \in \mathbb{R}^{n\times K}$ and $\mathbf{V} = [\mathbf{V}_1^\top, \mathbf{V}_2^\top, \dots, \mathbf{V}_m^\top]^\top \in \mathbb{R}^{m\times K}$. The term $\alpha(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2)$ is introduced to avoid over-fitting, controlled by the parameter α. $\mathbf{W} \in \mathbb{R}^{n\times m}$ is a weight matrix where $\mathbf{W}_{ij}$ is the weight for the rating from u_i to v_j. In this work, we set $\mathbf{W}_{ij} = 1$ if $\mathbf{R}_{ij} \neq 0$ and 0 otherwise. Other schemes to define $\mathbf{W}$ can be found in (Pan et al. 2008; Li et al. 2010).

Modeling Heterogeneity

When social relations are considered to be homogeneous, one of the most popular approaches to exploit social relations is to assume that the preference of a user should be close to the average preference of her social network (Jamali and Ester 2009; Ma et al. 2011). However, a user's social network is usually heterogeneous, and is a composite of various types of relations (Tang and Liu 2009a). For example, in figure 1, the connection $\langle u_1, u_3\rangle$ may be different from the connection $\langle u_1, u_4\rangle$ since they are in different social dimensions. Similarly, when the connection $\langle u_i, u_j\rangle$ in the k-th social dimension $\mathcal{G}_k$, the preference u_i should be close to the average preference of $\mathcal{G}_k$. With this intuition, we introduce the preferences of social dimensions to model heterogeneity of social relations.

Let $\hat{\mathbf{U}}_k \in \mathbb{R}^{1\times K}$ be the social dimension preference vector of $\mathcal{G}_k$ and $\hat{\mathbf{U}} = [\hat{\mathbf{U}}_1^\top, \hat{\mathbf{U}}_2^\top, \dots, \hat{\mathbf{U}}_c^\top]^\top$. A user can be involved in multiple social dimensions, and she might have different degrees of association with social dimensions. Next we will define the strengths of users associating with social dimensions.

Let $\mathcal{F}_k$ be the set of items frequently rated by users from the k-th social dimension $\mathcal{G}_k$, which provides items, users in $\mathcal{G}_k$ are interested in, and is formally defined as

$$\mathcal{F}_k = \{v_j | \sum_{u_i \in \mathcal{G}_k} sign(\mathbf{R}_{ij}) \geq \tau,\ \ j \in \{1, 2, \dots, m\}\ \}$$

where $sign(x) = 1$ if $x > 0$ and $sign(x) = 0$ otherwise. We empirically find that $\tau = 2$ works well in this work. Let $\hat{\mathbf{R}} \in \mathbb{R}^{c\times m}$ be the rating matrix for social dimensions where $\hat{\mathbf{R}}_k$ is the rating vector of $\mathcal{G}_k$, denoting the rating preference of $\mathcal{G}_k$ to its preferred item set $\mathcal{F}_k$. For $v_j \notin \mathcal{F}_k$, $\hat{\mathbf{R}}_{kj} = 0$, while for $v_j \in \mathcal{F}_k$, $\hat{\mathbf{R}}_{kj}$ is the average rating of users in $\mathcal{G}_k$ to v_j as $\frac{\sum_{u_i \in \mathcal{G}_k} \mathbf{R}_{ij}}{|\mathcal{G}_k|}$ where $|\mathcal{G}_k|$ is the number of users in $\mathcal{G}_k$.

Let $\mathbf{S} \in \mathbb{R}^{n\times c}$ where $\mathbf{S}_{ik}$ is the strength of u_i associating with $\mathcal{G}_k$. A user u_i is likely to strongly associate with $\mathcal{G}_k$ if her ratings $\mathbf{R}_i$ are similar to the ratings $\hat{\mathbf{R}}_k$ of $\mathcal{G}_k$ or she has many connections in $\mathcal{G}_k$. Therefore for $u_i \in \mathcal{G}_k$, we define the strength $\mathbf{S}_{ik}$ as ,

$$\mathbf{S}_{ik} = \beta \frac{\sum_{v_j \in \mathcal{F}_k} \mathbf{R}_{ij} \cdot \hat{\mathbf{R}}_{kj}}{\sqrt{\sum_{v_j \in \mathcal{F}_k} \mathbf{R}_{ij}^2}\sqrt{\sum_{v_j \in \mathcal{F}_k} \hat{\mathbf{R}}_{kj}^2}} + (1-\beta)\frac{|\mathcal{N}_i \cap \mathcal{G}_k|}{|\mathcal{N}_i|},$$

where $\mathcal{N}_i$ is the set of users in her social network. The strength $\mathbf{S}_{ik}$ is a linear combination of the rating similarity and the overlap between u_i's social network and $\mathcal{G}_k$. The parameter β is introduced to control their contributions, which is set to 0.7 in this work. For $u_i \notin \mathcal{G}_k$, $\mathbf{S}_{ik} = 0$.

To model heterogeneity of social relations, we force the preference vector of a user close to social dimension preference vectors, controlled by their association strengths as,

$$\min \quad \sum_{i=1}^{n} \sum_{k \in \mathcal{I}(G)} \mathbf{S}_{ik} \|\mathbf{U}_i - \hat{\mathbf{U}}_k\|_2^2 \tag{2}$$

where $\mathcal{I}(G)$ is the set of social dimensions that u_i involves in, which is formally defined as

$$\mathcal{I}(G) = \{k | \mathbf{G}_{ik} = 1,\ k \in \{1, 2, \dots, c\}\ \} \tag{3}$$

In Eq.(2), a large value of $\mathbf{S}_{ik}$ indicates that u_i is strongly associated with $\mathcal{G}_k$ hence the distance between u_i's preference $\mathbf{U}_i$ and the k-th social dimension preference $\hat{\mathbf{U}}_k$ should be small, while a small value of $\mathbf{S}_{ik}$ suggests that their distance could be large.

Modeling Weak Dependency Connections

Two users u_i and u_j, who do no connect directly, are in the same social dimension $\mathcal{G}_k$ and then establish a weak dependency connection. Before modeling, we first investigate the existence of weak dependency connections by answering the question - are two users with a weak dependency connection more likely to share similar preferences than two randomly chosen users?

To answer the question, we calculate two similarities for each weak dependency connection, i.e., weak dependency similarity w and random similarity r. For example, for the weak dependency connection between u_i and u_j, w is the rating cosine similarity between u_i and u_j, while r is the similarity between u_i and a randomly chosen user without weak dependency connections. Finally we obtain two similarity vectors, $\mathbf{s}_w$ and $\mathbf{s}_r$. $\mathbf{s}_w$ is the set of all weak dependency similarities w while $\mathbf{s}_r$ is the set of r. We conduct a two-sample t-test on $\mathbf{s}_w$ and $\mathbf{s}_r$. The null hypothesis and the alternative hypothesis are defined as

$$H_0 : \mathbf{s}_w \leq \mathbf{s}_r, \quad H_1 : \mathbf{s}_w > \mathbf{s}_r. \tag{4}$$

In our two studied datasets, the null hypothesis is rejected at significance level $\alpha = 0.01$ with p-values of 3.14e-28 and 7.21e-15, respectively. Evidence from t-test suggests a positive answer to the question: *users with weak dependency connections are more likely to share similar preferences.* With the verification of the existence of weak dependency connections, next we will introduce details about how to model them.

Weak dependency connections may have different strengths. Let $\mathcal{S}_{ijk}$ be the strength of the weak dependence connection between u_i and u_j in $\mathcal{G}_k$. $\mathcal{S}_{ijk}$ indicates the interest similarity of u_i and u_j in $\mathcal{G}_k$, hence we use cosine rating similarity to calculate $\mathcal{S}_{ijk}$ in terms of $\mathcal{F}_k$ of $\mathcal{G}_k$, which is formally defined as

$$\mathcal{S}_{ijk} = \frac{\sum_{v_l \in \mathcal{F}_k} \mathbf{R}_{il} \cdot \mathbf{R}_{jl}}{\sqrt{\sum_{v_l \in \mathcal{F}_k} \mathbf{R}_{il}^2} \sqrt{\sum_{v_l \in \mathcal{F}_k} \mathbf{R}_{jl}^2}} \tag{5}$$

Two users with a weak dependence connection are likely to share similar user interests. To model a weak dependence connection between u_i and u_j, we force their user preferences $\mathbf{U}_i$ and $\mathbf{U}_j$ close to each other as

$$\min \sum_i \sum_{k \in \mathcal{I}(G)} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} \|\mathbf{U}_i - \mathbf{U}_j\|_2^2 \tag{6}$$

where a large value of $\mathcal{S}_{ijk}$ indicates that their user preferences $\mathbf{U}_i$ and $\mathbf{U}_j$ should be very close, while a small value of $\mathcal{S}_{ijk}$ indicates that the distance of their user preferences $\mathbf{U}_i$ and $\mathbf{U}_j$ could be large.

Our Framework - SoDimRec

With solutions to model heterogeneity and weak dependency connections, our framework SoDimRec is to minimize the following problem

$$\begin{aligned}
\min_{\mathbf{U},\mathbf{V},\hat{\mathbf{U}}} & \sum_{i=1}^{n} \sum_{j=1}^{m} \mathbf{W}_{ij} (\mathbf{R}_{ij} - \mathbf{U}_i \mathbf{V}_j^\top)^2 \\
& + \lambda_1 \sum_{i=1}^{n} \sum_{k \in \mathcal{I}_G} \mathbf{S}_{ik} \|\mathbf{U}_i - \hat{\mathbf{U}}_k\|_2^2 \\
& + \lambda_2 \sum_{i=1}^{n} \sum_{k=1}^{c} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} \|\mathbf{U}_i - \mathbf{U}_j\|_2^2 \\
& + \alpha (\|\hat{\mathbf{U}}\|_F^2 + \|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2)
\end{aligned} \tag{7}$$

In Eq. (7), the second term models the heterogeneity and λ_1 is introduced to control its contribution; while the third term is used to model weak dependency connections, controlled by λ_2. In this work, we employ the alternating least square approach to solve the problem in Eq. (7), which is proven to be efficient for solving these low-rank approximation problems and is also easy to parallelize for large-scale data (Zhou et al. 2008). The problem in Eq. (7) is not a standard least squares problem due to the weight matrix $\mathbf{W}$ thus we need to extend this approach for our problem. We use $\mathcal{J}_c$ to denote the objective function in Eq. (7), and the derivative of $\mathcal{J}_c$ with respect to each entry of $\mathbf{U}$ $\mathbf{U}_{ir}$ is

$$\begin{aligned}
\frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_{ir}} = & \sum_j \mathbf{W}_{ij} (\mathbf{R}_{ij} - \mathbf{U}_i \mathbf{V}_j^\top) \mathbf{V}_{jr} + \alpha \mathbf{U}_{ir} \\
& + \lambda_1 (\sum_{k \in \mathcal{I}(G)} \mathbf{S}_{ik} (\mathbf{U}_{ir} - \hat{\mathbf{U}}_{kr})) \\
& + \lambda_2 (\sum_{k \in \mathcal{I}(G)} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} (\mathbf{U}_{ir} - \mathbf{U}_{jr}))
\end{aligned} \tag{8}$$

Then $\frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_i}$ is computed as

$$\begin{aligned}
\frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_i} & = (\frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_{i1}}, \frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_{i2}}, \ldots, \frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_{ik}}) = \\
& \mathbf{U}_i \Big(\mathbf{V}^\top \hat{\mathbf{D}}_i \mathbf{V} + \alpha \mathbf{I} + \lambda_1 \sum_{k \in \mathcal{I}(G)} \mathbf{S}_{ik} \mathbf{I} \\
& + \lambda_2 \sum_{k \in \mathcal{I}(G)} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} \mathbf{I} \Big) - (\mathbf{R}_i \hat{\mathbf{D}}_i \mathbf{V} \\
& + \lambda_1 \sum_{k \in \mathcal{I}(G)} \mathbf{S}_{ik} \hat{\mathbf{U}}_k + \lambda_2 \sum_{k \in \mathcal{I}(G)} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} \mathbf{U}_j).
\end{aligned} \tag{9}$$

where $\hat{\mathbf{D}}_i \in \mathbb{R}^{n \times n}$ is a diagonal matrix with entries of i-th row in $\mathbf{W}$ on the diagonal. Setting $\frac{\partial \mathcal{J}_c}{\partial \mathbf{U}_i}$ to 0, we get

$$\begin{aligned}
\mathbf{U}_i = & (\mathbf{R}_i \hat{\mathbf{D}}_i \mathbf{V} + \lambda_1 \sum_{k \in \mathcal{I}(G)} \mathbf{S}_{ik} \hat{\mathbf{U}}_k \\
& + \lambda_2 \sum_{k \in \mathcal{I}(G)} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} \mathbf{U}_j) \Big(\mathbf{V}^\top \hat{\mathbf{D}}_i \mathbf{V} + \alpha \mathbf{I} \\
& + \lambda_1 \sum_{k \in \mathcal{I}(G)} \mathbf{S}_{ik} \mathbf{I} + \lambda_2 \sum_{k \in \mathcal{I}(G)} \sum_{u_j \in \mathcal{H}_{ik}} \mathcal{S}_{ijk} \mathbf{I} \Big)^{-1}
\end{aligned} \tag{10}$$

Similarly, taking the derivative of $\mathcal{J}_c$ with respect to $\mathbf{V}_j$, we have

$$\frac{\partial \mathcal{J}_c}{\partial \mathbf{V}_j} = \mathbf{V}_j (\mathbf{U}^\top \hat{\mathbf{H}}_j \mathbf{U} + \alpha \mathbf{I}) - ((\mathbf{R}^j)^\top \hat{\mathbf{H}}_j \mathbf{U}) \tag{11}$$

where $\hat{\mathbf{H}}_j \in \mathbb{R}^{m \times m}$ is a diagonal matrix with entries of j-th column in $\mathbf{W}$ on the diagonal. Setting $\frac{\partial \mathcal{J}_c}{\partial \mathbf{V}_j}$ to 0, we get

$$\mathbf{V}_j = ((\mathbf{R}^j)^\top \hat{\mathbf{H}}_j \mathbf{U})(\mathbf{U}^\top \hat{\mathbf{H}}_j \mathbf{U} + \alpha \mathbf{I})^{-1} \tag{12}$$

In Eq. (7), the only term containing $\hat{\mathbf{U}}_k$ is $\sum_{i=1}^{n} \sum_{k \in \mathcal{I}_G} \mathbf{S}_{ik} \|\mathbf{U}_i - \hat{\mathbf{U}}_k\|_2^2$. Taking the derivative of $\sum_{i=1}^{n} \sum_{k \in \mathcal{I}_G} \mathbf{S}_{ik} \|\mathbf{U}_i - \hat{\mathbf{U}}_k\|_2^2$ with respect to $\hat{\mathbf{U}}_k$, we can obtain the updating rule for $\hat{\mathbf{U}}_k$ as

$$\hat{\mathbf{U}}_k = \frac{\sum_{u_i \in \mathcal{G}_k} \mathbf{S}_{ik} \mathbf{U}_i}{\alpha + \sum_{u_i \in \mathcal{G}_k} \mathbf{S}_{ik}} \tag{13}$$

We use Eqs. (10), (12), and (13) to update $\mathbf{U}$, $\mathbf{V}$ and $\hat{\mathbf{U}}$ until convergence. After we learn $\mathbf{U}$ and $\mathbf{V}$, an unknown rating from $u_{i'}$ and $v_{j'}$ can be estimated as $\mathbf{U}_{i'} \mathbf{V}_{j'}^\top$.

Experiments

In this section, we conduct experiments to verify the following: (1) does exploiting heterogeneity and weak dependency connections via social dimensions help recommendation? and (2) if it does, where are these contributions from? After introducing experimental settings, we compare the proposed recommender systems SoDimRec with the state-of-the-art recommender systems to answer the first question, and then investigate the effects of the components on SoDimRec to answer the second question.

Table 1: Statistics of the datasets

	Epinions	Ciao
# of Users	21,882	7,252
# of Items	59,104	21,880
# of Ratings	632,663	183,749
# of Social Relations	348,197	110,536

Experimental Settings

We collect two datasets to evaluate our proposed recommender system, i.e., Epinions and Ciao[1], and these two datasets are publicly available via the homepage of the first author [2]. Epinions and Ciao are product review websites. Users in Epinions and Ciao are allowed to specify scores from 1 to 5 to rate items, and they can also establish relations with others. Some statistics of these two datasets are presented in Table 1.

For each dataset, we choose $x\%$ as the training set to learn parameters and the remaining $1 - x\%$ as the testing set where x is varied as $\{45, 65, 85\}$. We will repeat the experiments 5 times and report the average performance. The experimental settings are exactly the same as these in our previous paper (Tang et al. 2013). Two popular metrics, i.e., the Root Mean Square Error (RMSE) and the Mean Absolute Error (MAE), are chosen to evaluate the prediction performance. A smaller RMSE or MAE value means better performance. Note that (Koren 2008; 2009) demonstrated that *small improvement in RMSE or MAE terms can have a significant impact on the quality of the top-few recommendation*.

Comparisons of Different Recommender Systems

In this subsection, we aim to evaluate the effectiveness of the proposed recommender system by comparing it to the following representative recommender systems:

- **MF**: This method performs matrix factorization on the user-item rating matrix (Salakhutdinov and Mnih 2008). It only utilizes rating information;
- **SoRec**: This method is based on matrix factorization and performs a co-factorization on the user-term rating matrix and user-user social relation matrix (Ma et al. 2008);
- **SoReg**: This method is also based on matrix factorization and social regularization is defined to capture strong dependency connections (Ma et al. 2011);
- **STE**: This method models a rating from one user by combining ratings from the user and her social network under the matrix factorization framework (Ma, King, and Lyu 2009); and
- **LOCABAL**: This method models social networks from both local and global perspectives under the matrix factorization framework (Tang et al. 2013).

Note that LOCABAL treats a user's relations homogeneously and use her/his direct connections while SoDimRec considers heterogeneity and weak ties, which results in

[1] http://www.ciao.co.uk/

[2] http://www.jiliang.xyz/trust.html

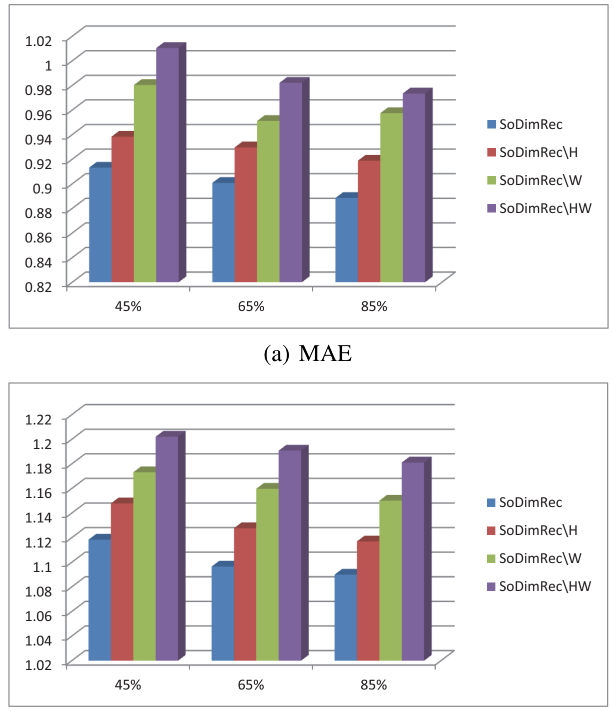

(b) RMSE

Figure 2: Effects of Heterogeneity and Weak Dependence Connections on SoDimRec.

that SoDimRec is substantially different from LOCABAL in terms of both key ideas and techniques. For all baseline methods, we use cross validation to determine their parameters. For SoDimRec, we set $\{K = 20, c = 100, \lambda_1 = 5, \lambda_2 = 100\}$ and $\{K = 30, c = 500, \lambda_1 = 10, \lambda_2 = 100\}$ for Ciao and Epinions, respectively. α is empirically set to 0.1. Comparison results are demonstrated in Table 2, and we make the following observations:

- Exploiting social relations can improve recommendation performance in terms of both RMSE and MAE.
- The proposed recommender system SoDimRec always outperforms baseline methods. The major reason is that the proposed framework exploits heterogeneity of social relations and weak dependency connections via social dimensions. More details about the effects of heterogeneity and weak dependency connections on the proposed recommender system will be discussed in the following subsection.

We conduct t-test on all comparisons and the t-test results suggest that all improvements are significant. With these observations, we can draw an answer to the first question - the proposed recommender system based on social dimensions outperforms the state-of-the-art recommender systems.

Table 2: Comparisons of different recommender systems.

Training	Metrics	Algorithms					
		MF	**SoRec**	**STE**	**SoReg**	**LOCABAL**	SoDimRec
Ciao (45%)	MAE	0.9962	0.9625	0.9598	0.9573	0.9385	0.9256
	RMSE	1.1801	1.1411	1.1357	1.1303	1.1132	1.0981
Ciao (65%)	MAE	0.9851	0.9457	0.9402	0.9331	0.9263	0.9023
	RMSE	1.1569	1.1172	1.1137	1.1107	1.0905	1.0697
Ciao (85%)	MAE	0.9623	0.9439	0.9358	0.9264	0.9112	0.8878
	RMSE	1.1405	1.1033	1.1019	1.1000	1.0821	1.0519
Epinions (45%)	MAE	1.0104	0.9647	0.9532	0.9492	0.9299	0.9131
	RMSE	1.2021	1.1672	1.1612	1.1545	1.1364	1.1184
Epinions (65%)	MAE	0.9821	0.9555	0.9475	0.9349	0.9116	0.9007
	RMSE	1.1908	1.1563	1.1386	1.1319	1.1148	1.0965
Epinions (85%)	MAE	0.9734	0.9487	0.9321	0.9292	0.9039	0.8885
	RMSE	1.1811	1.1479	1.11392	1.1281	1.1097	1.0900

Impact of Heterogeneity and Weak Dependence Connections

In this subsection, we investigate the effects of heterogeneity and weak dependence connections on the proposed framework SoDimRec to answer the second question. In detail, we systematically eliminate their effects from SoDimRec by defining its variants as follows,

- *SoDimRec*$\backslash H$ - Eliminating the effect of heterogeneity by setting $\lambda_1 = 0$ in Eq. (8);
- *SoDimRec*$\backslash W$ - Eliminating the effect of weak dependence connections by setting $\lambda_2 = 0$ in Eq. (8); and
- *SoDimRec*$\backslash HW$ - Eliminating the effects of both heterogeneity and weak dependence connections by setting $\lambda_1 = 0$ and $\lambda_2 = 0$;

The results in Epinions are shown in Figures 2(a) and 2(b) for MAE and RMSE, respectively. Since we have similar observations in Ciao, we only show results in Epinions to save space. When eliminating the effect of heterogeneity, the performance of *SoDimRec*$\backslash H$ degrades. We have the similar observation for *SoDimRec*$\backslash W$ when eliminating the effect of weak dependence connections. When eliminating both effects, the performance of *SoDimRec*$\backslash HW$ further decreases. These results suggest that both heterogeneity and weak dependence connections can help improve the recommendation performance.

Conclusion

Connections in online social networks are intrinsically heterogeneous and various relations are mixed together. Users are likely to interact with other users with similar interests more frequently, which results in forming groups in online social networks, and share similar interests with other users in the same group although they may not connect directly. However, heterogeneity of social relations and weak dependency connections are overlooked by most existing social recommender systems. In this paper, we study how to exploit heterogeneity of social relations and weak dependency connections for recommendation. We first adopt social dimensions to capture heterogeneity of social relations and weak dependency connections simultaneously, and then propose a novel recommendation framework which models heterogeneity and weak dependency connections based on social dimensions. Experimental results on real-world datasets show that the proposed recommender system outperforms the state-of-the-art recommender systems.

There are several directions to investigate in the future. First, in our current work, we only consider user connections within social dimensions. Users from different social dimensions may have different interests, thus we will investigate how to take advantage of connections between social dimensions for recommendation. Second, aside from positive connections such as friendships and trust relations, users in social media can also have negative relations such as foes and distrust relations (Yang et al. 2012). Exploiting negative relations for recommendation is rarely studied by existing social recommender systems (Tang et al. 2015). Therefore investigating ways to exploit negative relations in recommendation will be a promising direction. Finally, user preferences may change over time. For example, a user currently interested in “Sports” may prefer “Electronics” in the future. Also social dimensions are likely to evolve such as new members being added and old members leaving. Therefore, we want to investigate temporal information of both ratings and social relations in recommender systems.

Acknowledgments

This material is based upon work supported by, or in part by, the National Science Foundation (NSF) under grant number IIS-1217466 and Office of Naval Research (ONR) under grant number N000141410095.

References

Golbeck, J. 2006. Generating predictive movie recommendations from trust in social networks. *Trust Management* 93–104.

Granovetter, M. 1973. The strength of weak ties. *American Journal of Socialogy* 78(6):1360–1380.

Granovetter, M. 1983. The strength of weak ties: A network theory revisited. *Sociological theory* 1(1):201–233.

Herlocker, J.; Konstan, J.; Borchers, A.; and Riedl, J. 1999. An algorithmic framework for performing collaborative filtering. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 230–237. ACM.

Jamali, M., and Ester, M. 2009. Trustwalker: a random walk model for combining trust-based and item-based recommendation. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 397–406. ACM.

Jamali, M., and Ester, M. 2010. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM conference on Recommender systems*, 135–142. ACM.

Jiang, M.; Cui, P.; Liu, R.; Yang, Q.; Wang, F.; Zhu, W.; and Yang, S. 2012. Social contextual recommendation. In *CIKM*. ACM.

Koren, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 426–434. ACM.

Koren, Y. 2009. Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 447–456. ACM.

Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010. Predicting positive and negative links in online social networks. In *ACM WWW International conference on World Wide Web (WWW)*.

Levin, D. Z., and Cross, R. 2004. The strength of weak ties you can trust: The mediating role of trust in effective knowledge transfer. *Management science* 50(11):1477–1490.

Li, Y.; Hu, J.; Zhai, C.; and Chen, Y. 2010. Improving one-class collaborative filtering by incorporating rich user information. In *CIKM*, 959–968. ACM.

Ma, H.; Yang, H.; Lyu, M.; and King, I. 2008. Sorec: social recommendation using probabilistic matrix factorization. In *Proceeding of the 17th ACM conference on Information and knowledge management*, 931–940. ACM.

Ma, H.; Zhou, D.; Liu, C.; Lyu, M.; and King, I. 2011. Recommender systems with social regularization. In *Proceedings of the fourth ACM international conference on Web search and data mining*, 287–296. ACM.

Ma, H.; King, I.; and Lyu, M. 2009. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*, 203–210. ACM.

Marsden, P., and Friedkin, N. 1993. Network studies of social influence. *Sociological Methods and Research* 22(1):127–151.

Massa, P., and Avesani, P. 2007. Trust-aware recommender systems. In *Proceedings of the 2007 ACM conference on Recommender systems*, 17–24. ACM.

Miller McPherson, L. S.-L., and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27(1):415–444.

Newman, M. E. J. 2005. Power laws, pareto distributions and zipf's law. *Contemporary Physics* 46:323–351.

Pan, R.; Zhou, Y.; Cao, B.; Liu, N. N.; Lukose, R.; Scholz, M.; and Yang, Q. 2008. One-class collaborative filtering. In *ICDM*, 502–511. IEEE.

Salakhutdinov, R., and Mnih, A. 2008. Probabilistic matrix factorization. *Advances in neural information processing systems* 20:1257–1264.

Sarwar, B.; Karypis, G.; Konstan, J.; and Riedl, J. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*, 285–295. ACM.

Su, X., and Khoshgoftaar, T. 2009. A survey of collaborative filtering techniques. *Advances in Artificial Intelligence* 2009:4.

Sun, Y., and Han, J. 2012. Mining heterogeneous information networks: Principles and methodologies. *Synthesis Lectures on Data Mining and Knowledge Discovery* 3(2):1–159.

Tang, L., and Liu, H. 2009a. Relational learning via latent social dimensions. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 817–826.

Tang, L., and Liu, H. 2009b. Scalable learning of collective behavior based on sparse social dimensions. In *CIKM'09: Proceeding of the 18th ACM conference on Information and knowledge management*, 1107–1116. New York, NY, USA: ACM.

Tang, J.; Hu, X.; Gao, H.; and Liu, H. 2013. Exploiting local and global social context for recommendation. In *IJCAI*.

Tang, J.; Chang, Y.; Aggarwal, C.; and Liu, H. 2015. A survey of signed network mining in social media. *arXiv:1511.07569*.

Tang, J.; Chang, Y.; and Liu, H. 2014. Mining social media with social theories: A survey. *SIGKDD Explorations*.

Tang, J.; Gao, H.; and Liu, H. 2012. mTrust: Discerning multi-faceted trust in a connected world. In *Proceedings of the fifth ACM international conference on Web search and data mining*, 93–102. ACM.

Wang, F.; Li, T.; Wang, X.; Zhu, S.; and Ding, C. 2011. Community discovery using nonnegative matrix factorization. *Data Mining and Knowledge Discovery* 22(3):493–521.

Yang, S.-H.; Smola, A. J.; Long, B.; Zha, H.; and Chang, Y. 2012. Friend or frenemy?: predicting signed ties in social networks. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, 555–564. ACM.

Zhou, Y.; Wilkinson, D.; Schreiber, R.; and Pan, R. 2008. Large-scale parallel collaborative filtering for the netflix prize. *Algorithmic Aspects in Information and Management* 337–348.