Leveraging Decomposed Trust in Probabilistic Matrix Factorization for Effective Recommendation

Hui Fang Yang Bao[†] Jie Zhang
School of Computer Engineering, Nanyang Technological University, Singapore
[†] Nanyang Business School, Nanyang Technological University, Singapore
{hfang@e.ntu.edu.sg, baoyang@ntu.edu.sg, zhangj@ntu.edu.sg}

Abstract

Trust has been used to replace or complement rating-based similarity in recommender systems, to improve the accuracy of rating prediction. However, people trusting each other may not always share similar preferences. In this paper, we try to fill in this gap by decomposing the original single-aspect trust information into four general trust aspects, i.e. benevolence, integrity, competence, and predictability, and further employing the support vector regression technique to incorporate them into the probabilistic matrix factorization model for rating prediction in recommender systems. Experimental results on four datasets demonstrate the superiority of our method over the state-of-the-art approaches.

Introduction

Trust has been extensively exploited to improve the predictive accuracy of recommendations by ameliorating the issues such as *data sparsity* and *cold start* that recommender systems inherently suffer from (Massa and Avesani 2007; Ma et al. 2008). Basically, trust provides additional information from which user preference can be better modeled, alternative or complementary to rating-based similarity. Both implicit (O'Donovan and Smyth 2005) and explicit (Yang et al. 2013) trust sources have been investigated in the literature. The former trust is usually inferred from user-item interactions (i.e. ratings) whereas the latter is directly specified by users indicating whom and to what extent they trust.

Most trust-aware recommender systems adopt explicit trust networks (e.g. Epinions.com) or social networks (e.g. Flixster.com) where users can refer to the information source of their trust neighbors or friends. As supported by the *social influence theory*, people are similar to each other in their networks for two reasons (Crandall et al. 2008): 1) they grow to ensemble their current friends or neighbors due to *social influence*; and 2) they tend to form new links to others who have already added them into the networks. However, a generally agreed proposition states that people trusting each other may not always share similar preferences (Jøsang, Quattrociocchi, and Karabeg 2011; Au Yeung and Iwata 2011). In other words, the explicit trust information involves too much noise regarding to user preference modeling, which consequently impacts the performance of recommender systems.

In this paper, in order to bridge the gap between trust and user preference-similarity, and adopt trust information more effectively in recommender systems, we propose a latent factor model that identifies more effective aspects of trust for recommender systems. Researchers in social science generally admit the multi-facet property of trust (McKnight and Chervany 2002). By decomposing the explicit trust value into finer-grained trust aspects, we can derive more effective information for recommendation. Specifically, we first identify four general trust aspects (i.e. benevolence, integrity, competence and predictability) that are modeled based on users' past ratings (user-item interactions). The four aspects are combined to a Support Vector Regression (SVR) model for trust value prediction between two users. The dependency between the trust aspects is captured by a Guassian radial basis kernel function. Then, we incorporate the trust information into the probabilistic matrix factorization model (Mnih and Salakhutdinov 2007) by modeling trust as jointly conditioning on the trust value obtained from the SVR model, as well as similarity between the corresponding latent user feature vectors factorized from rating matrix. Consequently, we re-interpret the trust value for recommendation, and also reasonably update a user's latent feature vector by considering the *social influence* of other users trusting and being trusted by the user. We further validate the effectiveness of our model by comparing with four state-of-the-art approaches. Experimental results on four real datasets clearly demonstrate the better effectiveness of our model for rating prediction.

Related Work

Trust is well-known as a heterogenous rather than homogenous concept in the fields of social science and computational trust, and consists of multiple aspects. Specifically, Mayer et al. (1995) report that the trust relationship between a trustor (who specifies trust statements) and a trustee (who receives trust statements) is mainly influenced by the trustor's *propensity* to trust others in terms of three aspects related with the trustee, namely *ability* (competence), *benevolence* and *integrity*. Mcknight and Chervany (2002) enrich this model by adding one more aspect of the trustee—

predictability. We present the formal definitions of these aspects in the following sections. Further, these frameworks are adopted to form the socio-cognitive trust theory in the area of computational trust (Castelfranchi and Falcone 2010; Fang, Bao, and Zhang 2013). Consistently, we also employ the four general aspects of the trustee as well as the trustor's propensity to formulate the trust relationship. However, most trust-aware recommender systems simply treat trust as a concept with a single aspect, i.e., the ability to provide accurate ratings (*competence*) (O'Donovan and Smyth 2005). One possible explanation is that limited information is available in the few and publicly accessible data sets. Although some previous works attempt to model multiple aspects of raters (who give ratings) in recommender systems (Kwon, Cho, and Park 2009), they are essentially distinct concepts from trust.

There are two major ways of incorporating trust information in recommender systems. The first one is trust-related neighborhood models. For example, Massa and Avesani (2007) replace user similarity with explicitly specified trust relationships, and also allow trust relationships to propagate through the trust networks. Jamali and Ester (2009) design the *TrustWalker* approach to randomly select neighbors in the trust network, and then combine trust information of the selected neighbors with an item-based technique. Guo et al. (2012) empirically contend that the preferences of users can be better modeled by merging the ratings of trusted neighbors, which can lead to the improvement over recommendation performance. However, these approaches suffer from the following serious problems: 1) they combine trust information with rating information heuristically, without systematically revealing the relationship between them; 2) an underlying assumption that, a more trustworthy neighbor to a user also shares more similar preference with the user, usually could not hold in the real world; 3) the original explicit trust network might involve a considerable amount of noisy trust labels.

The other type is trust related latent factor models. For instance, Ma et al. (2008) design the *SoRec* approach by fusing the user-item rating matrix with user-user trust matrix. *SoRec* outperforms the basic matrix factorization model and other trust-related neighborhood models. However, this model suffers from the problem of low interpretability. To model trust information more realistically, they further propose *RSTE* (Ma, King, and Lyu 2009), which interprets one user's rating decision as the balance between this user's own taste and her trusted neighbors' favors. Jamali and Ester (2010) then enhance this model by enabling trust propagation in their *SocialMF* model. Yang et al. (2013) present a new way of fusing rating data with trust data (called *TrustMF*), which considers the influence of a user's trust network on her rating decision, and the user's rating decision on the rating decision of other users trusting her. To effectively use social network information when trust information is not available, Ma et al. (2011) propose a matrix factorization framework with social regularization. This work differs from trust-aware recommender systems by being aware of the difference of trust relationship and friend relationship. The latent factor models advance the neighborhood models as they systematically investigate the relationship between trust data and rating data when fusing them together for rating prediction. However, they still suffer from the second and third problems indicated in the neighborhood models.

Problem Definition and Preliminaries

In our scenario, we adopt two kinds of information sources for the recommendation: the user-item rating matrix and user-user trust network. To be specific, we have a set of users $U = \{u_1, \cdots, u_N\}$, and a set of items $I = \{i_1, \cdots, i_M\}$. The ratings given by users on items are represented in a rating matrix $R = [R_{ui}]_{N\times M}$, where R_{ui} denotes the rating of user u on item i. In our paper, we normalize the real ratings to the interval $[0, 1]$ to avoid the loss of generality. For the user-user trust network, each user u has specified N_u direct trust neighbors, and T_{uv} (ranged in $[0, 1]$) denotes the trust value of user u to user v. If user u has not specified a trust value to v, we assume $T_{uv} = 0$. All the trust values are denoted by a matrix $T = [T_{uv}]_{N\times N}$. The objective of a recommender system is: given a user $u \in U$ and an item $i \in I$, if R_{ui} is unknown, predict u's rating on item i.

In our paper, we learn the latent characteristics of users and items using the probabilistic matrix factorization (PMF) technique (Mnih and Salakhutdinov 2007) to factorize the user-item rating matrix. That is, let $U \in \mathbb{R}^{K\times N}$ and $V \in \mathbb{R}^{K\times M}$ be latent user and item feature matrices, where column vectors U_u and V_i representing K-dimensional user-specific and item-specific latent feature vectors of user u and item i, respectively. The objective of matrix factorization is to learn U and V for recommendation. According to (Mnih and Salakhutdinov 2007), we define the conditional probability of the observed ratings as:

$$p(R|U, V, \sigma_R^2) = \prod_{u=1}^{N}\prod_{i=1}^{M}\left[\mathcal{N}(R_{ui}|g(U_u^T V_i), \sigma_R^2)\right]^{C_{ui}^R} \quad (1)$$

where $N(x|\mu, \sigma^2)$ is the Gaussian (normal) distribution with mean μ and variance σ^2, and C_{ui}^R is the indicator function that is equal to 1 if user u rated item i and 0 otherwise. The function $g(x)$ is the logistic function $g(x) = 1/(1 + exp(-x))$ and bounds the range of $U_u^T V_i$ within $[0, 1]$. Zero-mean Gaussian priors are set for user and item feature vectors:

$$p(U|\sigma_U^2) = \prod_{u=1}^{N}\mathcal{N}(U_u|0, \sigma_U^2\mathbf{I}),\ p(V|\sigma_i^2) = \prod_{i=1}^{M}\mathcal{N}(V_i|0, \sigma_V^2\mathbf{I}) \quad (2)$$

Through the Bayesian inference, the posterior probability is:

$$\begin{aligned} p(U, V|R, \sigma_R^2, \sigma_U^2, \sigma_V^2) &\propto p(R|U, V, \sigma_R^2)p(U|\sigma_U^2)p(V|\sigma_V^2) \\ &= \prod_{u=1}^{N}\prod_{i=1}^{M}[\mathcal{N}(R_{ui}|g(U_u^T V_i), \sigma_R^2)]^{C_{ui}^R} \\ &\times \prod_{u=1}^{N}\mathcal{N}(U_u|0, \sigma_U^2\mathbf{I}) \times \prod_{i=1}^{M}\mathcal{N}(V_i|0, \sigma_V^2\mathbf{I}) \end{aligned} \quad (3)$$

Here, U and V can be learned purely based on user-item rating matrix using the gradient descent technique. In our model, we also consider the trust information. In the next sections, we first formally formulate the four trust aspects according to users' ratings, and then present our proposed approach by considering the trust aspects.

Trust Modeling

Trust has been well recognized as a multi-faceted concept that involves *benevolence*, *integrity*, *competence*, and *predictability* in social science. We justify the connections between the aspects and trust in the original trust model in (Mayer, Davis, and Schoorman 1995). Specifically, we regard the combination of each aspect of a trustee and the propensity of a trustor as *an aspect of the trustee perceived by the trustor*, or a *trusting belief* of the trustor that the trustee has the corresponding characteristic in her favor. Therefore, trust in our model is connected with four different trust beliefs, each of which is regarded as an aspect of a trustee perceived by a trustor. Together with the trustor's propensity[1], they are known as the antecedents of trust (McKnight, Choudhury, and Kacmar 2002), and formulated in detail as follows.

Formulating Trust Aspects

As belief can be modeled by evidence (Shafer 1976), we proceed to formulate the trust aspects in the light of users' historical experience from rating matrix R. Suppose there are two users: a trustor u and a trustee v, and each user has a set of experience denoted by E_u and E_v, respectively. A piece of experience is denoted by a 4-tuple $e_u = (u, i, R_{ui}, t)$, indicating that user u rated item i as R_{ui} at time t. Hence, users' experience can be represented as $E_u = \{e_{u1}, \dots, e_{ux}\}$ and $E_v = \{e_{v1}, \dots, e_{vy}\}$, where x and y are the number of experience of users u and v, respectively. Based on user experience, we then model the four general trust aspects of trustee v from the viewpoint of trustor u as follows.

Benevolence, $Be(u, v)$, refers to the extent to which v cares about the preferences of u, i.e. v's willingness to do good deed for u. If u has a high benevolence belief to v, v's preference is more likely to be similar with u. It means, in our case, that both u and v report similar ratings on many items. Hence, benevolence is modeled as the user similarity which can be computed by the Pearson correlation coefficient (Adomavicius and Tuzhilin 2005):

$$Be(u,v) = \frac{\sum\limits_{i \in E_{u,v}} (R_{ui} - \bar{R}_u)(R_{vi} - \bar{R}_v)}{\sqrt{\sum\limits_{i \in E_{u,v}} (R_{ui} - \bar{R}_u)^2} \sqrt{\sum\limits_{i \in E_{u,v}} (R_{vi} - \bar{R}_v)^2}}, \tag{4}$$

where $E_{u,v} = E_u \cap E_v$ is the set of shared experience on the commonly rated items between u and v, and $\bar{R}_u, \bar{R}_v$ are the average of the ratings reported by u and v, respectively.

Integrity, $In(v)$, refers to the extent to which v conforms to a norm or code of moral or artistic values. It stresses v's characteristic to follow the norm or rules of an organization. In contrast to benevolence, integrity is independent of the trustor-trustee relationship. Therefore, it is formulated merely based on the past experience of v regardless of u's actions. Specifically, the behaviors of the majority are treated as the norm or the code when evaluating the integrity of v, i.e., the similarity between the trustee's behaviors and the majority's. Hence, integrity is computed by the similarity between the preferences of the trustee and the average:

$$In(v) = \frac{\sum\limits_{i \in E_v} (R_{vj} - \bar{R}_v)(\bar{R}_i - \bar{R})}{\sqrt{\sum\limits_{i \in E_v} (R_{vi} - \bar{R}_v)^2} \sqrt{\sum\limits_{i \in E_v} (\bar{R}_i - \bar{R})^2}}, \tag{5}$$

where $\bar{R}_i$ refers to the average of the ratings on item $i \in E_v$, and $\bar{R}$ is the average of the ratings on all items.

Competence, $Co(u, v)$, refers to the ability or the power of v to conduct the actions that are expected by u. The more experience trustee v has, the more competent she will be in the view of trustor u. Two factors are taken into account, i.e., the number of v's experience (see Equation 7), and the ratio of satisfactory opinions given by v to all the other users in the system (see Equation 6), employing the basic idea of O'Donovan and Smyth (2005). The competence is computed by integrating both factors:

$$Co(u,v) = \gamma \, \frac{\sum\limits_{i \in E_v} \underset{j \in U_i}{count}(|R_{vi} - R_{ji}| < \varepsilon_c)}{\sum\limits_{i \in E_v} \parallel U_i \parallel}, \tag{6}$$

where U_i represents the set of users who rated item i, and ε_c is a predefined error tolerance threshold below which a rating R_{vi} of v is treated as satisfactory for item i relative to the other's real preference R_{ji}, and γ is defined by:

$$\gamma = \begin{cases} \frac{y}{N^u} & \text{if } y \le N^u; \\ 1 & \text{otherwise;} \end{cases} \tag{7}$$

where N^u is the minimal number of experience required by u such that a user can be regarded as a reliable rater.

Predictability, $Pr(u, v)$, refers to the consistency of v's actions (good or bad, negative or positive) such that u can make a prediction in a given situation. Different from other aspects, the value of predictability is neutral. Predictability is able to alleviate the problem of behaviors changing strategically, that is, a user may first act honestly but conduct dishonest behaviors later. In this sense, predictability of v is defined as the degree to which the (positive, neutral or negative) trend of v's rating behaviors is distinct relative to the behaviors of trustor u:

$$\begin{aligned} n_1 &= \underset{i \in E_{u,v}}{\text{count}} (|R_{ui} - R_{vi}| \le \theta); \\ n_2 &= \underset{i \in E_{u,v}}{\text{count}} (R_{ui} - R_{vi} > \theta); \\ n_3 &= \underset{i \in E_{u,v}}{\text{count}} (R_{ui} - R_{vi} < -\theta); \\ Pr(u,v) &= 1 + \sum_{i=1}^{3} \frac{n_i}{\parallel E_{u,v} \parallel} \log_3 \frac{n_i}{\parallel E_{u,v} \parallel} \end{aligned} \tag{8}$$

where n_1, n_2 and n_3 refer to the neutral, negative and positive trends of user v's rating behaviors compared to trustor u's behaviors, respectively; θ is a threshold predefined by trustor u. The intuition is that for a user who is highly predictable, the difference in trends should be significant. In case of $n_1 = n_2 = n_3$, we obtain the lowest predictability since it is hard to predict the next behavior of the trustee.

Trust Prediction

T_{uv} will be influenced by the set of the five aspects (including the trustor's propensity) that we investigated, denoted by F_{uv}=$\{b_u, Be(u,v), Co(u,v), In(v), Pr(u,v)\}$.

[1]This refers to the trustor's inherent propensity to trust other users. It could be treated as a continuous constant (in the range of $[0, 1]$) subject to each trustor.

In practice, users may specify other users as trusted ($t \in (0,1]$) neighbors, or not ($t = 0$). The trust and absence of trust connections will help build a useful model of the trust aspects and the overall trust. We hence adopt support vector regression (Chang and Lin 2011) for trust prediction between two users. We choose SVR mainly because 1) it is recognized as a moderate method to cope with the imbalanced dataset problem. That is, in our scenario, each user only labels a small amount of trust neighbors, whereas majority relationships between users are unidentified; 2) by using kernel function, we can catch the nonlinearly dependent relationships among trust aspects. Given a training set of N_t data points $\{y_k, x_k\}_{k=1}^{N_t}$, where x_k represents the value of the trust aspects of the k-th input data point, and y_k is the corresponding trust value. We thus aim at constructing a trust prediction model of the form:

$$y = f(x) = \sum_{k=1}^{N_t} (\bar{\alpha}_k - \bar{\alpha}_k^*)\mathcal{K}(x, x_k) + \bar{b} \qquad (9)$$

where $\mathcal{K}(x, x_k)$ is the Gaussian radial basis function (RBF), denoted as $\mathcal{K}(x, x_k) = exp(-\|x - x_k\|^2/2\sigma^2)$, and α_k and α_k^* are given by the Lagrange multipliers:

$$\begin{aligned} \max_{\alpha,\alpha^*} W(\alpha, \alpha^*) &= \max_{\alpha,\alpha^*} \sum_{k=1}^{N_t} \alpha_k^*(y_k - \epsilon) - \alpha_k(y_k + \epsilon) \\ &- \frac{1}{2}\sum_{k=1}^{N_t}\sum_{j=1}^{N_t} (\alpha_k^* - \alpha_k)(\alpha_j^* - \alpha_j)\mathcal{K}(x_k, xj) \\ \text{with constraints} &\begin{cases} 0 \le \alpha_k, \alpha_k^* \le C, k = 1, \cdots, N_t \\ \sum_{k=1}^{N_t} (\alpha_k - \alpha_k^*) = 0 \end{cases} \end{aligned} \qquad (10)$$

where ϵ and C are well defined parameters that influence the performance of the SVR. In the next section, we will illustrate how to incorporate the predicted trust value into the PMF model for rating prediction.

Trust Decomposed Matrix Factorization

To reasonably interpret and adopt trust information for recommendation, we also consider the association between trust and the latent user feature matrix. That is, the trust between two users is related to trust aspects indicated above, and the latent user vectors of the two users simultaneously. This is accordance with social influence theory that *a user will become more similar to other users trusting and being trusted by the user.* Hence, the trust between two users is modeled as jointly conditioning on trust aspects, and similarity between the corresponding latent user feature vectors (see Figure 1). Specifically, we model the conditional distribution over the observed trust network as $p(T|F, U, \sigma_T^2)$

$$= \prod_{u=1}^{N}\prod_{v=1}^{N}\left[\mathcal{N}\left(T_{uv}|(\lambda s(U_u, U_v) + (1-\lambda)f(F_{uv})), \sigma_T^2\right)\right]^{C_{uv}^T} \qquad (11)$$

where $f(F_{uv})$ is from Equation 9, and $s(U_u, U_v) = e^{-(U_u - U_v)^T(U_u - U_v)}$ captures the similarity of the latent user feature vectors. The influence of the two perspectives on trust is balanced by the parameter λ. C_{uv}^T is the indicator function, and equal to 1 if trust is specified from u to v and 0 otherwise.

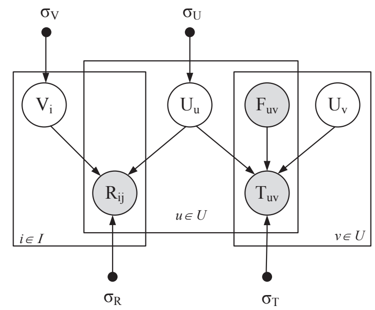

Figure 1: The Proposed Model

Through Bayesian inference in Figure 1, we model the conditional distribution of U and V over the observed ratings and trust information as:

$$\begin{aligned} &p(U, V|R, T, F, \sigma_T^2, \sigma_R^2, \sigma_U^2, \sigma_V^2) \propto \\ &\quad p(R|U, V, \sigma_R^2)p(U|\sigma_U^2)p(V|\sigma_V^2)p(T|F, U, \sigma_T^2) \end{aligned} \qquad (12)$$

Accordingly, the log of the posterior distribution for the recommendation is given by:

$$\begin{aligned} &\ln p(U, V|R, T, F, \sigma_T^2, \sigma_R^2, \sigma_U^2, \sigma_V^2) = -\frac{1}{2\sigma_R^2}\sum_{u=1}^{N}\sum_{i=1}^{M} C_{ui}^R(g(U_u^T V_i) - R_{ui})^2 \\ &- \frac{1}{2\sigma_U^2}\sum_{u=1}^{N} U_u^T U_u - \frac{1}{2\sigma_V^2}\sum_{i=1}^{M} V_i^T V_i \\ &- \frac{1}{2\sigma_T^2}\sum_{u=1}^{N}\sum_{v=1}^{N} C_{uv}^T(\lambda s(U_u, U_v) + (1-\lambda)f(F_{uv}) - T_{uv})^2 \\ &- \frac{1}{2}(\sum_{u=1}^{N}\sum_{i=1}^{M} C_{ui}^R)\ln\sigma_R^2 - \frac{1}{2}(\sum_{u=1}^{N}\sum_{v=1}^{N} C_{uv}^T)\ln\sigma_T^2 \\ &- \frac{1}{2}(NK\ln\sigma_U^2 + MK\ln\sigma_V^2) + \text{Constant} \end{aligned} \qquad (13)$$

Maximizing the log-posterior over latent features of users and items is equivalent to minimizing the following objective equation by keeping the hyperparameters (the observation noise variance and prior variance) fixed:

$$\begin{aligned} &\mathcal{L}(U, V, R, T) = \frac{1}{2}\sum_{u=1}^{N}\sum_{i=1}^{M} C_{ui}^R(U_u^T V_i - R_{ui})^2 + \frac{\lambda_U}{2}\sum_{u=1}^{N} U_u^T U_u + \\ &\frac{\lambda_V}{2}\sum_{i=1}^{M} V_i^T V_i + \frac{\lambda_T}{2}\sum_{u=1}^{N}\sum_{v=1}^{N} C_{uv}^T(\lambda s(U_u, U_v) + (1-\lambda)f(F_{uv}) - T_{uv})^2 \end{aligned} \qquad (14)$$

where λ_U=σ_R^2/σ_U^2, λ_V=σ_R^2/σ_V^2, and λ_T=σ_R^2/σ_T^2. The local minimum of Equation 14 can thus be obtained by conducting gradient descent on U_u and V_i for all users and items:

$$\begin{aligned} \frac{\partial\mathcal{L}}{\partial U_u} &= \sum_{i=1}^{M} C_{ui}^R V_i g'(U_u^T V_i)(g(U_u^T V_i) - R_{ui}) + \lambda_U U_u \\ &+ \lambda_T\sum_{v=1}^{N} C_{uv}^T s'(U_u, V_v)(\lambda s(U_u, U_v) + (1-\lambda)f(F_{uv}) - T_{uv}) \\ &+ \lambda_T\sum_{v=1}^{N} C_{vu}^T s'(U_v, V_u)(\lambda s(U_v, U_u) + (1-\lambda)f(F_{uv}) - T_{uv}) \\ \frac{\partial\mathcal{L}}{\partial V_i} &= \sum_{u=1}^{N} C_{ui}^R U_v g'(U_u^T V_i)(g(U_u^T V_i) - R_{ui}) + \lambda_V V_i \end{aligned} \qquad (15)$$

where g' and s' are the derivatives of the corresponding functions, respectively.

Our goal is to simultaneously optimize the parameters associated with ratings (i.e. U and V) and parameters associated with trust network (i.e. α_i and α_i^*). As indicated above

U and V are fit by gradient descent in Equation 15, whereas α_i and α_i^* is updated through Equation 10. Note that α_i and α_i^* can also influence the rating prediction. In this case, we design an alternative optimization procedure, where each iteration consists of the following two steps: 1) fix α_i and α_i^*, update U and V using Equation 15; 2) fix U and V, update α_i and α_i^* using Equation 10.

Experiments

We carry out experiments to evaluate the performance of our approach for recommendation, and conduct comparisons with four competing approaches on four real datasets.

Dataset Description

Four available real-world datasets are used in the experiments, namely Epinions[2], FilmTrust[3], Flixster.com, and Ciao[4]. All the four datasets enable their users to provide ratings to items and also specify other users as their trust links. Note that in Epinions and Ciao, the trust links are directed: user u_1 trusts user u_2 does not imply u_2 also trusts u_1, while the links in Flixster and FilmTrust are undirected. We use the original FilmTrust, Epinions, and Ciao datasets, and randomly sample a small portion of original Flixster dataset. The statistic information is summarized in Table 1.

Table 1: Statistic Information about the Datasets

Aspect	Epinions	Ciao	Flixster	FilmTrust
Num. of Users	49,289	7,375	1,000	1,508
Num. of Items	139,738	106,797	2,867	2,071
Avg. Ratings/Item	4.76	2.66	2.78	17.14
Avg. Trust Links/User	9.88	15.16	3.34	1.75

Benchmark Approaches

We compare our approach with four state-of-the-art approaches, including two approaches without using trust information, i.e. PMF (Mnih and Salakhutdinov 2007) and SVD++ (Koren 2008), and two approaches using trust information, i.e. SocialMF (Jamali and Ester 2010), and TrustMF (Yang et al. 2013). For all these methods, we set optimal parameters recommended in the literature, as indicated in Table 2, and set the dimension (K) of the latent space as 5 and 10, respectively. For our method, we use the SVR component provided by LIBSVM C# package[5]. For PMF, SocialMF, and SVD++, we adopt the source codes provided by MyMediaLite recommender system library[6]. Besides, we empirically set ε_c=0.1 for competence (see Equation 6) and θ=0.1 for predictability (see Equation 8) computations in our approach.

Evaluation Metrics

We use the standard 5-fold cross-validation for learning and testing, where for each dataset, 80% of the data is randomly selected as the training set and the rest as the testing set. The performance of approaches is evaluated by two commonly used measures: the mean square error (MSE) and mean absolute error (MAE). They both refer to the differences between the predictions and the ground truth, but differently, MAE simply tells how much (the absolute value) the predictions deviate from the ground-truth, while MSE penalizes strong deviations from the ground-truth (Ferri, Hernández-Orallo, and Modroiu 2009). Generally, smaller MSE and MAE values indicate better predictive accuracy.

Table 2: Parameters Setting of Methods

Methods	Parameters Setting
PMF	$\lambda_U = \lambda_V = 0.001$
SVD++	recommended in (Koren 2008)
SocialMF	$\lambda_U = \lambda_V = 0.001, \lambda_T = 1$
TrustMF	$\lambda = 0.001, \lambda_T = 1$
Our method	$\lambda_U = \lambda_V = 0.001, \lambda_T = 1$

Results and Discussion

We first compare the performance of our method with benchmarks, and then examine the effectiveness of the SVR component, and sensitivity of parameters in our method.

Model Comparison Table 3 presents the performance comparison between our method and other approaches on the four datasets. As demonstrated, our method (with λ=0.5) achieves better performance than other approaches in terms of MSE and MAE on all the four datasets (only slightly worse than SVD++ with respect to MSE on the Ciao dataset while K=10). "$*$" in the table indicates the corresponding approach with the best performance other than our method, respectively, and the percentages in the eighth and fifteenth rows present our improvements over them. Although the improvements may look small, as argued in (Koren 2010), slight improvements in terms of MSE and MAE value can lead to significant improvements in real applications. We also conduct t-test based on the results of each approach for 10-runs, and the results (the ninth and sixteenth rows) show that our improvements are statistically significant at the 5% level (p-value < 0.05). Besides, the improvement of the second best approach over the third best approach is also about 2.56% in average on MSE, indicating that further improvement on the best existing approaches becomes not easy and slight improvements should count as a big contribution.

As shown in Table 3, SVD++ performs better than TrustMF on the Flixster dataset, while the performance of SocialMF is worse than PMF on the Flixster and FilmTrust datasets. These might be due to: 1) adopting other kinds of information (e.g. hidden feedback), and combing model-based with latent-factor based approaches can greatly enhance the recommendation performance, especially when rating and trust information are relatively sparse (e.g. Flixster, see Table 1); 2) there exists substantial noisy trust relationships on the two datasets (i.e. FilmTrust and Flixster); and 3) trust neighbors of a user might not share the same preference with the user. In this case, these results, on the other hand, also partially support our research that by decomposing single-aspect trust into multiple fine-grained aspects, we can well deal with the noisy trust relationships

[2] www.trustlet.org/wiki/Epinions_datasets

[3] www.trust.mindswap.org/FilmTrust

[4] www.public.asu.edu/~jtang20/datasetcode/truststudy.htm

[5] www.csie.ntu.edu.tw/~cjlin/libsvm/

[6] www.mymedialite.net/

Table 3: Performance Comparison of Different Approaches

Dimension	Approach \ Dataset	Epinions MSE	Epinions MAE	Ciao MSE	Ciao MAE	Flixster MSE	Flixster MAE	FilmTrust MSE	FilmTrust MAE
$K = 5$	PMF	1.3193	0.8856	1.1209	0.8157	1.0648	0.8071	0.6861	0.6353
	SVD++	1.1553	0.8388	0.9684*	0.7557	0.8122*	0.6861*	0.6798	0.6523
	SocialMF	1.1913	0.8354	1.0416	0.7710	1.1439	0.7709	0.7347	0.6564
	TrustMF	1.1281*	0.8257*	0.9696	0.7546*	0.8469	0.6990	0.6494*	0.6268*
	Our Method	**1.0962**	**0.8064**	**0.9453**	**0.7366**	**0.7950**	**0.6747**	**0.6295**	**0.6155**
	(improve)	2.83%	2.34%	2.39%	2.44%	2.12%	1.66%	3.06%	1.80%
	p-value	0.0000	0.0000	0.0000	0.0003	0.0326	0.0012	0.0294	0.0220
$K = 10$	PMF	1.3764	0.9032	1.1384	0.8200	1.0431	0.8006	0.7001	0.6392
	SVD++	1.1553*	0.8389	0.9690*	0.7560	0.8127*	0.6865*	0.6799	0.6523
	SocialMF	1.2001	0.8389	1.0520	0.7737	1.1470	0.7722	0.7431	0.6597
	TrustMF	1.1688	0.8162*	1.0193	0.7488*	0.8963	0.7018	0.6667*	0.6308*
	Our Method	**1.1221**	**0.8137**	0.9700	**0.7454**	**0.8032**	**0.6796**	**0.6560**	**0.6253**
	(improve)	2.87%	0.31%	-0.10%	0.45%	1.17%	1.01%	1.60%	0.87%
	p-value	0.0000	0.1100	0.7434	0.0654	0.0222	0.0122	0.0221	0.0311

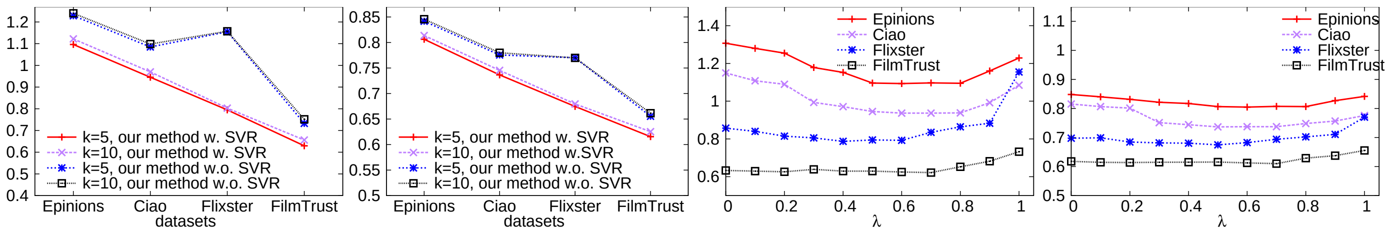

(a) **(b)** **(c)** **(d)**

Figure 2: The Effectiveness of the SVR Component in Our Method on (a) MSE; and (b) MAE; Performance Change of Our Method by Varying the λ ($K = 5$) on (c) MSE; and (d) MAE.

in recommender systems, and thus derive more effective information for more accurate recommendation.

Model Effectiveness and Sensitivity In Figures 2(a) and (b), we analyze the effectiveness of the SVR component (i.e. trust decomposition process) in our model by comparing the performance of our method with the SVR component (*w. SVR*), and that without the SVR component (*w.o. SVR*). As can be seen, our method with SVR performs much better than that without SVR when $K = 5$ and 10, further indicating the effectiveness of the trust decomposition using SVR technique, which can significantly contribute to the performance improvement for recommendation.

We also investigate the sensitivity of K parameter in our method. As shown in Table 3 and Figure 2, similar to other trust-aware latent factor models, smaller K values (e.g. $K = 5$) leads to more accurate rating prediction than relatively larger K values (e.g. $K = 10$). This is in accordance with the findings in (Yang et al. 2013), suggesting that, in the recommender context, users trust each other (or specify others as their trust neighbors) might because they are similar in a few particularly important aspects (not all).

Figures 2(c) and (d) present the performance change of our method in terms of λ parameter in Equation 11. As can be seen, our method obtains worse performance when λ approaches to 0 or 1. This is mainly because, $\lambda = 0$ means that there is no connection between user similarity and derived trust value from SVR, while $\lambda = 1$, user similarity is equivalent to the original trust value specified by users. None of the two conditions can reasonably explain the real scenario. This also well justifies our choice of jointly modeling trust as the combination of user similarity and re-predicted trust value from SVR. As presented in Figures 2(c) and (d), our method can achieve satisfactory performance when $\lambda \in [0.3, 0.8]$.

Conclusion and Future Work

We presented a trust-aware recommendation approach to mitigate the research gap between user similarity and trust concepts in recommender systems. Specifically, we first decomposed the original single-aspect trust information into four general trust aspects, i.e. benevolence, integrity, competence and predictability. Then, we incorporated these aspects into the probabilistic matrix factorization model for rating prediction with the support vector regression technique. We conducted comparisons with four state-of-the-art approaches, PMF, SVD++, SocialMF and TrustMF. Experimental results on four real datasets demonstrated that our model can improve the performance of recommender systems, and also validated the effectiveness of trust decomposition process in our approach, indicating that the more valuable trust information is derived for recommendation.

The contributions of our research mainly lie in: (1) we recognize the gap between trust and user preference that trust can only partially explain user similarity, but not all; (2) we address the multi-facet property of trust, and re-model trust information by jointly considering the trust aspects and user similarity; and (3) we combine a user's influence to other users trusting the user, and that to other users trusted by the user, where their trust information is utilized to update the user latent feature vectors. For future work, we will extend our model by considering other properties (i.e. trust transitivity) of trust networks to further improve recommendation accuracy, and well interpret the real-world intuitions.

Acknowledgements

This work is supported by the MoE AcRF Tier 2 Grant M4020110.020 and the ACI Seed Funding M4080962.C90 awarded to Dr. Jie Zhang.

References

Adomavicius, G., and Tuzhilin, A. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering* 17(6):734–749.

Au Yeung, C.-m., and Iwata, T. 2011. Strength of social influence in trust networks in product review sites. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM)*, 495–504. ACM.

Castelfranchi, C., and Falcone, R. 2010. *Trust theory: A socio-cognitive and computational model*, volume 18. John Wiley & Sons.

Chang, C.-C., and Lin, C.-J. 2011. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)* 2(3):27.

Crandall, D.; Cosley, D.; Huttenlocher, D.; Kleinberg, J.; and Suri, S. 2008. Feedback effects between similarity and social influence in online communities. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 160–168. ACM.

Fang, H.; Bao, Y.; and Zhang, J. 2013. Misleading opinions provided by advisors: Dishonesty or subjectivity. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 1983–1989.

Ferri, C.; Hernández-Orallo, J.; and Modroiu, R. 2009. An experimental comparison of performance measures for classification. *Pattern Recognition Letters* 30(1):27–38.

Guo, G.; Zhang, J.; and Thalmann, D. 2012. A simple but effective method to incorporate trusted neighbors in recommender systems. In *User Modeling, Adaptation, and Personalization (UMAP)*. Springer. 114–125.

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 (KDD)*, 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 (RecSys*, 135–142. ACM.

Jøsang, A.; Quattrociocchi, W.; and Karabeg, D. 2011. Taste and trust. In *Trust Management V*. Springer. 312–322.

Koren, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 426–434. ACM.

Koren, Y. 2010. Factor in the neighbors: Scalable and accurate collaborative filtering. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 4(1):1–23.

Kwon, K.; Cho, J.; and Park, Y. 2009. Multidimensional credibility model for neighbor selection in collaborative recommendation. *Expert Systems with Applications* 36(3):7114–7122.

Ma, H.; Yang, H.; Lyu, M. R.; and King, I. 2008. SoRec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management (CIKM)*, 931–940. ACM.

Ma, H.; Zhou, D.; Liu, C.; Lyu, M. R.; and King, I. 2011. Recommender systems with social regularization. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM)*, 287–296. ACM.

Ma, H.; King, I.; and Lyu, M. R. 2009. Learning to recommend with social trust ensemble. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 203–210. ACM.

Massa, P., and Avesani, P. 2007. Trust-aware recommender systems. In *Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys)*, 17–24. ACM.

Mayer, R. C.; Davis, J. H.; and Schoorman, F. D. 1995. An integrative model of organizational trust. *Academy of Management Review* 20(3):709–734.

McKnight, D. H., and Chervany, N. L. 2002. What trust means in e-commerce customer relationships: an interdisciplinary conceptual typology. *International Journal of Electronic Commerce* 6:35–60.

McKnight, D. H.; Choudhury, V.; and Kacmar, C. 2002. Developing and validating trust measures for e-commerce: an integrative typology. *Information Systems Research* 13(3):334–359.

Mnih, A., and Salakhutdinov, R. 2007. Probabilistic matrix factorization. In *Advances in neural information processing systems*, 1257–1264.

O'Donovan, J., and Smyth, B. 2005. Trust in recommender systems. In *Proceedings of the 10th International Conference on Intelligent User Interfaces*, 167–174. ACM.

Shafer, G. 1976. *A mathematical theory of evidence*, volume 1. Princeton university press Princeton.

Yang, B.; Lei, Y.; Liu, D.; and Liu, J. 2013. Social collaborative filtering by trust. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2747–2753. AAAI Press.