Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence

A Transitivity Aware Matrix Factorization Model for Recommendation in Social Networks

Mohsen Jamali and Martin Ester
School of Computing Science
Simon Fraser University
Burnaby, BC, Canada
{mohsen_jamali, ester}@cs.sfu.ca

Abstract

Recommender systems are becoming tools of choice to select the online information relevant to a given user. Collaborative filtering is the most popular approach to building recommender systems and has been successfully employed in many applications. With the advent of online social networks, the social network based approach to recommendation has emerged. This approach assumes a social network among users and makes recommendations for a user based on the ratings of the users who have direct or indirect social relations with the given user. As one of their major benefits, social network based approaches have been shown to reduce the problems with cold start users. In this paper, we explore a model-based approach for recommendation in social networks, employing matrix factorization techniques. Advancing previous work, we incorporate the mechanism of trust propagation into the model in a principled way. Trust propagation has been shown to be a crucial phenomenon in the social sciences, in social network analysis and in trust-based recommendation. We have conducted experiments on two real life data sets. Our experiments demonstrate that modeling trust propagation leads to a substantial increase in recommendation accuracy, in particular for cold start users.

1 Introduction

With the rapidly growing amount of information available on the WWW, it becomes necessary to have tools to help users to select the relevant part of online information. To satisfy this need, recommender systems have emerged, e.g. there are popular recommenders for movies[1], books[2], music[3], etc.

Typically in a recommender, we have a set of *users* and a set of *items*. Each user u rates a set of items by some values. The recommender has the task to predict the rating for user u on a non-rated item i or to generally recommend some items for the given user u based on the ratings that already exist.

[1] http://www.netflix.com
[2] http://www.amazon.com
[3] http://www.last.fm

Generally two type of recommender systems have been investigated: Memory-based and Model-based. Memory based algorithms (collaborative filtering) explore the user-item rating matrix and make recommendations based on the ratings of item i by a set of users whose rating profiles are most similar to that of user u. Model-based approaches learn the parameters of a model and store only those parameters. Hence they do not need to explore the rating matrix. Model-based approaches are very fast after the parameters of the model are learnt. The bottleneck for model-based approaches is the training phase, while in memory-based approaches there is no training, but the prediction (test) phase is slower.

With the advent of online social networks, the social network based approach to recommendation has emerged. This approach assumes a social network among users and makes recommendations for a user based on the ratings of the users that have direct or indirect social relations with the given user.

Collaborative filtering is most effective when users have expressed enough ratings to have common ratings with other users, but it performs poorly for so-called cold start user. Cold start users are new users who have expressed only a few ratings. Users found to be similar based on few ratings expressed by cold start users are not a reliable indicator of similarity. Social network based recommenders, however, can make recommendations as long as a new user is connected to a large enough component of the social network.

Exploiting social networks in recommendation works because of the effects of selection and social influence that have been postulated by sociologists for a long time. Selection means that people tend to relate to people with similar attributes, and due to social influence related people in a social network influence each other to become more similar[Wasserman and Faust, 1994]. The increasing availability of online social network data has finally allowed a verification of these sociological models. The results of experiments in [Crandall *et al.*, 2008] and of similar work confirm that a social network provides an independent source of information which can be exploited to improve the quality of recommendations.

A social rating network (SRN) is a social network in which each user expresses ratings on some items besides creating social relations to other users. Note that the terms ”trust network” and ”social network” are used as synonyms throughout this paper.

Some memory based approaches have been proposed for recommendation in social rating networks [Golbeck, 2005][Jamali and Ester, 2009]. These methods typically explore the social network and find a neighborhood of users trusted (directly or indirectly) by a user and perform the recommendation by aggregating their ratings. These methods use the transitivity of trust and propagate trust to indirect neighbors in the social network.

Recently, the model based approach for recommendation in social rating networks has been investigated [Ma *et al.*, 2009a]. These methods exploit the matrix factorization technique to learn latent features for users and items from the observed ratings. The generative nature of these models allow better prediction of future behavior of users. Experimental results show better performance compared to state of the art memory-based approaches. However, existing generative models for recommendation do not consider the propagation of trust. In this paper, we also propose a matrix factorization based model for recommendation in social rating networks, called SocialMF. We incorporate the propagation of trust in our model to improve the quality of recommendation. To inject social influence in our model, we make the features of every user dependent on the feature vectors of the his direct neighbors in the social network. Using this idea, latent features of users indirectly connected in the social network will be dependent and hence the trust gets propagated.

Cold start users are one of the most important challenges in recommender systems. Since cold start users are more dependent on the social network compared to users with more ratings, the effect of using trust propagation gets more important for cold start users. Moreover, in many real life SRNs a very large portion of users do not express any ratings, and they only participate in the social network. Hence, using only the observed ratings does not allow to learn the user features. The SocialMF model forces the user feature vectors to be close to those of their neighbors to be able to learn the latent user features for users with no or very few ratings.

2 Matrix Factorization for Recommendation

In recommender systems we have a set of users $\mathbb{U} = \{u_1, \ldots u_N\}$ and a set of items $\mathbb{I} = \{i_1, \ldots i_M\}$. The ratings expressed by users on items are given in a rating matrix $R = [R_{u,i}]_{N \times M}$. In this matrix $R_{u,i}$ denotes the rating of user u on item i. $R_{u,i}$ can be any real number, but often ratings are integers in the range $[1, 5]$. In this paper, without loss of generality, we map the ratings *1, ..., 5* to the interval [0,1] by normalizing the ratings. In a social rating network, each user u has a set N_u of direct neighbors and $t_{u,v}$ denotes the value of social trust u has on v as a real number in $[0, 1]$. Zero means no trust and one means full trust. Binary trust networks are the most common trust networks (Amazon[4], eBay[5], ...). The trust values are given in a matrix $T = [T_{u,v}]_{N \times N}$. Non-zero cells $T_{u,v}$ in T denote the existence of a social relation from u to v. Note that T is asymmetric in general. The task of a recommender is as follows: Given a user $u \in \mathbb{U}$ and an item $i \in \mathbb{I}$ for which $R_{u,i}$ is unknown, predict the rating for u on item i using R and T.

[4]www.amazon.com

[5]www.ebay.com

Matrix factorization (MF) techniques have been widely employed for recommendation [Salakhutdinov and Mnih, 2008][Koren, 2008][Koren *et al.*, 2009]. The underlying assumption is that the observed rating behavior of users is governed by latent features associated with both users and items. In order to learn the latent features of users and items, MF learns a factorization of the rating matrix into a product of user features and item feature. To predict the unknown rating of a user for an item, MF uses the product of the latent features of the given user and item. Let $U \in \mathbb{R}^{K \times N}$ and $V \in \mathbb{R}^{K \times M}$ be latent user and item feature matrices, with column vectors U_u and V_i representing K-dimensional user-specific and item-specific latent feature vectors of users u and item i, respectively. The conditional probability of the observed ratings is defined as:

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

where $\mathcal{N}(x|\mu, \sigma^2)$ is the normal distribution with mean μ and variance σ^2, and $I_{u,i}^R$ is the indicator function that is equal to 1 if u has rated i and equal to 0 otherwise. The function $g(x)$ is the logistic function $g(x) = 1/(1 + e^{-x})$, which bounds the range of $U_u^T V_i$ within [0,1]. The corresponding graphical model is presented in figure 1. Using equation 3, we can learn the latent feature vectors of users and items purely based on the user-item rating matrix.

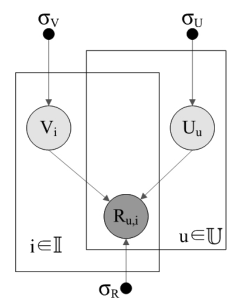

Figure 1: Graphical Model of the baseline factorization of user-item rating matrix.

3 Related Work

In this section we review some related work on recommendation in social networks. Trust propagation in recommendation has been widely investigated in the memory based approaches. Hence, we first review some of the memory based methods for recommendation in social networks. Matrix factorization has been widely used in the model based recommendation [Koren, 2008][Salakhutdinov and Mnih, 2008][Koren *et al.*, 2009]. However these models do not take into account the social network among users. Recently, some model-based approaches have been proposed which use matrix factorization for recommendation in social networks [Ma *et al.*, 2009a], however, most of these works do not consider the propagation of trust. The only method that can potentially incorporate trust propagation is not a generative model and defines a loss function that is not intuitive. In this section, after reviewing memory based approaches we discuss some model based approaches for recommendation in social networks.

TidalTrust[Golbeck, 2005] performs a modified breadth first search in the trust network to compute a prediction. Basically, it finds all raters with the shortest path distance from the source user and aggregates their ratings weighted by the trust between the source user and these raters. To compute the trust value between user u and v who are not directly connected, TidalTrust aggregates the trust value between u's direct neighbors and v weighted by the direct trust values of u and its direct neighbors.

In order to consider enough ratings without suffering from noisy data, [Jamali and Ester, 2009] proposes a random walk method (TrustWalker) which combines trust-based and item-based recommendation. TrustWalker considers not only ratings of the target item, but also those of similar items. The probability of using the rating of a similar item instead of a rating for the target item increases with increasing length of the walk. Their framework contains both trust-based and item-based collaborative filtering recommendations as special cases. Their experiments show that their method outperforms other existing memory based approaches. The random walk model allows them to compute the confidence in the predictions.

The authors of [Ma *et al.*, 2009a] proposed a matrix factorization approach for social network based recommendation, called STE. Their method is a linear combination of basic matrix factorization approach and a social network based approach. The graphical model for their proposed model is illustrated in figure 2[6]. The predicted rating of user u on item i is as follows:

$$\hat{R}_{u,i} = g(\alpha U_u^T V_i \ + \ (1-\alpha) \sum_{v \in N_u} T_{u,v} U_v^T V_i) \qquad (2)$$

where parameter α controls the effects of neighbors on the estimated rating.

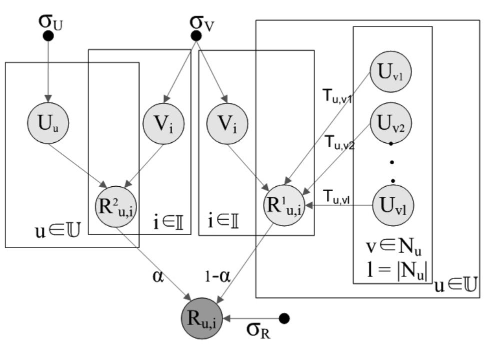

Figure 2: The STE model [Ma et al., 2009a].

Experiments show that their model outperforms the basic matrix factorization based approach and existing trust based based approaches. However, in their model, the feature vectors of direct neighbors of u affect the ratings of u instead of affecting the feature vector of u. This model does not handle trust propagation. We use this method as our main comparison partners in our experiments and call it as *STE* model.

[6]It should be noted that model in figure 2 is different from the graphical model presented in [Ma *et al.*, 2009a], but it correctly represents the joint probability distribution actually computed for the STE model.

The same authors proposed another method for recommendation in social networks [Ma *et al.*, 2009b]. This method is not a generative model. They define a loss function and try to minimize the loss function. The loss function introduced in [Ma *et al.*, 2009b] looks like the loss function of our proposed model. However, it is not intuitive since it punishes the users with lots of social relations more than other users. Besides, the loss function does not correspond to any underlying probabilistic model. Their loss function tries to make the latent features of a user close to sum of the latent features of his friends which leads to extreme punishment for users with many social relations. This model could potentially incorporate trust propagation into the recommendation model.

4 The SocialMF Model

Traditional recommender systems, like collaborative filtering approaches [Salakhutdinov and Mnih, 2008][Goldberg *et al.*, 1992], only utilize the information of the user-item rating matrix for recommendations but ignore the social relations among users. With the exponential growth of online social networks, incorporating social networks into recommender systems is becoming more and more important. In this section, we introduce our proposed model that incorporates trust propagation into a matrix factorization model for recommendation in social networks.

Due to social influence [Friedkin, 1998], the behavior of a user u is affected by his direct neighbors N_u. In other words, the latent feature vector of u is dependent on the latent feature vectors of all his direct neighbors $v \in N_u$. We formulate this influence as follows:

$$\widehat{U}_u = \frac{\sum_{v \in N_u} T_{u,v} U_v}{\sum_{v \in N_u} T_{u,v}} = \frac{\sum_{v \in N_u} T_{u,v} U_v}{|N_u|} \qquad (3)$$

where $\widehat{U}_u$ is the estimated latent feature vector of u given the feature vectors of his direct neighbors. Since the social networks we are working with are all binary social networks, all none-zero values of $T_{u,v}$ are 1. We normalize each row of the trust matrix so that $\sum_{v=1}^{N} T_{u,v} = 1$. Now, we have:

$$\widehat{U}_u = \sum_{v \in N_u} T_{u,v} U_v \qquad (4)$$

The above equation indicates that the estimate of the latent feature vector of a user is the weighted average of the latent feature vectors of his direct neighbors. Note that taking the social network into account does not change the equation for the conditional distribution of the observed ratings. It only affects the user latent feature vectors. So the conditional probability of observed rating is the same as the conditional probability in equation 1:

$$p(R|U,V,\sigma_R^2) = \prod_{u=1}^{N}\prod_{i=1}^{M}\Big[\mathcal{N}\Big(R_{u,i}|g(U_u^T V_i),\sigma_R^2\Big)\Big]^{I_{u,i}^R} \quad (5)$$

For the user latent features, we have two factors: The zero-mean Gaussian prior to avoid over-fitting, and the conditional distribution of user latent features given the latent features of his direct neighbors. Therefore,

$$p(U|T,\sigma_U^2,\sigma_T^2) \propto p(U|\sigma_U^2) \times p(U|T,\sigma_T^2)$$
$$= \prod_{u=1}^{N}\mathcal{N}\Big(U_u|0,\sigma_U^2\mathbf{I}\Big) \times \prod_{u=1}^{N}\mathcal{N}\Big(U_u|\sum_{v\in N_u} T_{u,v}U_v,\sigma_T^2\mathbf{I}\Big) \quad (6)$$

The graphical model corresponding to the SocialMF model is shown in figure 3. Note that the trust matrix in the above equation is not explicitly shown in the figure. However, the edges among the latent feature vectors of users are representatives of the trust network among users and the degree of trust of user u on user v is $T_{u,v}$.

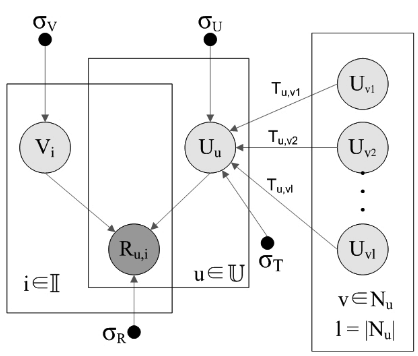

Figure 3: Proposed Graphical Model to consider the social network in the factorization of user-item rating matrix.

Through a Bayesian inference, we computed the log-posterior of the latent variables[7]. Maximizing the log-posterior is equivalent to minimizing the following objective function, which is a sum of squared errors with quadratic regularization terms:

$$\mathcal{L}(R,T,U,V) = \frac{1}{2}\sum_{u=1}^{N}\sum_{i=1}^{M} I_{u,i}^R (R_{u,i} - g(U_u^T V_i))^2$$
$$+ \frac{\lambda_U}{2}\sum_{u=1}^{N} U_u^T U_u + \frac{\lambda_V}{2}\sum_{i=1}^{M} V_i^T V_i$$
$$+ \frac{\lambda_T}{2}\sum_{u=1}^{N}\Big((U_u - \sum_{v\in N_u} T_{u,v}U_v)^T (U_u - \sum_{v\in N_u} T_{u,v}U_v)\Big) \quad (7)$$

In the above equation, $\lambda_U = \sigma_R^2/\sigma_U^2$, $\lambda_V = \sigma_R^2/\sigma_V^2$, and $\lambda_T = \sigma_R^2/\sigma_T^2$. We use gradient descent to find the optimum latent vectors. The objective function of the generative model underlying SocialMF looks similar to the loss function introduced in [Ma *et al.*, 2009b]. However, the loss function in [Ma *et al.*, 2009b] considers the sum of the latent features of neighbor instead of the average of them which leads to punishment of users with more social relations. Also, the loss function in [Ma *et al.*, 2009b] has no underlying probabilistic model.

[7] For detailed posterior inference, please refer to the original paper [Jamali and Ester, 2010]

4.1 Desirable properties of the proposed model

In this section, we discuss some desirable properties of SocialMF and compare it against the closely related STE model [Ma *et al.*, 2009a].

The SocialMF model addresses the transitivity of trust in social networks. In other words, our model takes the trust propagation into account. According to the graphical model, the feature vector of any user is dependent on the feature vectors of his direct neighbors. Recursively, the feature vector of each direct neighbor is dependent on the feature vector of his direct neighbors. This effect is shown in the conditional distributions by considering the feature vector of a user being a normal distribution around the average of the feature vectors of his neighbors. On the other hand, the STE model [Ma *et al.*, 2009a] does not support trust propagation and they list trust propagation as future work.

In the baseline MF approach [Salakhutdinov and Mnih, 2008] and the STE model [Ma *et al.*, 2009a], the features are being learnt based only on the observed ratings. However, in real life SRNs, a huge portion of users have expressed no ratings and they participate only in the social network. So their features can not be learnt based on their observed ratings. However, our model can handle these users very well. The SocialMF model learns to tune the latent features of these users close to their neighbors. So, despite not having any expressed ratings, the feature vectors of these users will be learnt to be close to their neighbors. Basically, the social trust relations among users is an observed dependency among the feature vectors of users. It should be noted that since evaluating the learnt features is typically based on the withheld observed ratings, we are currently not able to evaluate the features learnt for users with no expressed ratings.

4.2 Complexity analysis of parameter learning

The main cost in learning the parameters is computing $\mathcal{L}$ and its gradients against feature vectors of users and items. Assuming the average number of ratings per user is $\overline{r}$, and the average number of direct neighbors per user is $\overline{t}$, the complexity of evaluation of $\mathcal{L}$ is $O(N\overline{r}K + N\overline{t}K)$. Since both the rating matrix R and trust matrix T are very sparse, $\overline{t}$ and $\overline{r}$ are relatively small. So the computation of the objective function $\mathcal{L}$ is very fast and linear with respect to the number of users in the social rating network. The computational complexity of computing the gradients is $O(N\overline{r}K + N\overline{t}^2K)$ which is linear with respect to the number of users in the social rating network. Note that the cost of computing the gradient in STE [Ma *et al.*, 2009a] is $O(N\overline{r}\overline{t}^2K)$. So SocialMF is $\frac{\overline{r}\overline{t}^2}{\overline{r}+\overline{t}^2}$ times faster than STE in computing the gradient in each iteration of parameter learning process.

For each rating estimation, the model proposed in [Ma *et al.*, 2009a] needs to take the average of estimated ratings for direct neighbors which makes it slower in prediction compared to SocialMF proposed in this paper.

5 Experiments

We performed experiments on two real life data sets from Epinions.com and Flixster.com.. In this section, we report our experimental results and compare the results with existing methods. Also we present the results for different settings of model parameters.

5.1 Datasets

Epinions is an online product review website. We used the version of the Epinions dataset[8] published by the authors of [Richardson and Domingos, 2002]. Flixster is social networking service for rating movies. The Flixster dataset was first published by the authors of [Jamali and Ester, 2010]. Table 1 presents the general statistics of these two datasets.

Statistics	**Flixster**	**Epinions**
Users	1M	71K
Social Relations	26.7M	508K
Ratings	8.2M	575K
Items	49K	104K
Users with Rating	150K	47K
Users with Friend	980K	60K

Table 1: General statistics of the Flixster and Epinions

A large portion of users in Flixster have no expressed ratings, but most of them have social relations. Users without any ratings are also important. They may not be useful to compute the prediction for other users based on their own ratings, but they may allow us to connect indirectly to other users who have rated items. The distribution of the number of ratings per user follows a power law. It should also be noted that unlike Flixster, the items in Epinions are from different categories such as cameras, dvd players, music, etc, while all the items in the Flixster dataset are movies.

5.2 Experimental Setup

We perform 5-fold cross validation in our experiments. In each fold we have 80% of data as the training set and the remaining 20% as the test data.

The evaluation metric we use in our experiments is RMSE which is defined as follows:

$$RMSE = \sqrt{\frac{\sum_{(u,i)|R_{test}} (r_{u,i} - \widehat{r}_{u,i})^2}{|R_{test}|}} \tag{8}$$

where R_{test} is the set of all pairs (u, i) in the test data.

To evaluate the performance of our method we consider three comparison partners:

- BaseMF: This method is the baseline matrix factorization approach proposed in [Salakhutdinov and Mnih, 2008], which does not take the social network into account.
- STE: This is the model proposed in [Ma *et al.*, 2009a], which takes into account the social network in a way different from SocialMF. We set $\alpha = 0.4$ for STE in our experiments which is the optimum value according to the results of experiments in [Ma *et al.*, 2009a].
- CF: This is the well-known user based collaborative filtering method which is a memory based approach.

In our experiments, we refer to our proposed model as *SocialMF*. In all our experiments, we set $\lambda_U = \lambda_V = 0.1$.

5.3 Experimental Results

Table 2 reports the RMSE values of all comparison partners on Epinions and Flixster. The parameter λ_T is set to 5 for experiments on Epinions and $\lambda_T = 1$ for Flixster[9]. Table 2 shows that SocialMF outperforms existing methods. Note that since collaborative filtering has no latent features, there is no dimensionality K associated with it and hence the result for different values of K are the same.

SocialMF improves the RMSE of the state-of-the-art method STE by 6.2%. for K=5 and by 5.7% for K=10. To show how significant our gain is, note that the gain of STE over the baseline MF method is 2.5% and the gain of SocialMF over STE is more than 2 times that gain. As another evidence for the significance of these RMSE reductions, note that in the Netflix prize competition[10], there was a $1 Million reward for a reduction of the RMSE by 10%.

Method	Epinions		Flixster	
	K=5	K=10	K=5	K=10
CF	1.180	1.180	0.911	0.911
BaseMF	1.175	1.195	0.878	0.863
STE	1.145	1.150	0.864	0.852
SocialMF	1.075	1.085	0.821	0.815

Table 2: RMSE values for comparison partners on Epinions and Flixster with different settings of dimensionality K.

RMSE values for Flixster are also presented in table 2. Again, SocialMF clearly outperforms existing methods. In Flixster, the improvement of the RMSE for SocialMF over STE is 5% which more than 3 times of the gain of STE over baseline MF (1.5%).

It should be noted that the results for Flixster are generally better than the results for Epinions for all methods, possibly because of the fact that the items in Epinions are from multiple categories such as DVD players, cameras, printers, laptops, ..., while the items in Flixster are all movies, which makes the recommendation easier in general. Another explanation for the better results on Flixster could be that Flixster is a richer dataset since there are more social relations and ratings per user in Flixster compared to Epinions.

Intuitively, increasing K should add more flexibility to the model and hence should improve the results. However, com-

[8]http://alchemy.cs.washington.edu/data/epinions/

[9]It should be noted that λ_T has been tuned using sensitivity analysis for both data sets. For more details, please refer to [Jamali and Ester, 2010].

[10]http://www.netflixprize.com

paring results of tables 2 for different values of K shows that increasing K in Epinions did not improve the results, while increasing K in Flixster improved the results. We believe that these counter-intuitive results for Epinions are due to the fact that the Epinions data set is smaller than Flixster and increasing K leads to more parameters in the model which leads to overfitting. Flixster, on the other hand, is a huge data set, and increasing K to 10 does not lead to overfitting.

Some users in an SRN express a lot of ratings, but most users express a few ratings. We consider users who have expressed less than 5 ratings as cold start users [Jamali and Ester, 2009]. In both Flixster and Epinions more than 50% of users are cold start users[11]. Hence efficiency of any recommendation algorithm for cold start users becomes very important. We performed experiments on cold start users. According to the results [Jamali and Ester, 2010], the improvement of the RMSE for cold start users compared to STE is 11.5% for Epinions and 8.5% for Flixster. The gain for cold start users is more than the gain for all users which we discussed in previous subsection. This implies that SocialMF handles cold start users better than STE. We believe this is mainly due to the consideration of trust propagation and transitivity in our model.

It should also be noted that we performed thorough analysis on the actual learning time of SocialMF and STE, showing that our proposed SocialMF is much faster than STE. For the detailed discussion on the runtime experiments , please refer to the original paper [Jamali and Ester, 2010].

6 Conclusions and Future Work

Recommender systems are emerging as tools of choice to select the online information relevant to a given user. Collaborative filtering is the most popular approach to building recommender systems and has been successfully employed in many applications. With the advent of online social networks, exploiting the information hidden in the social network to predict the behavior of users has become very important.

In this paper we proposed a novel model based approach for recommendation in social networks. Our model is a matrix factorization based approach. Similar to the STE model presented in [Ma *et al.*, 2009a], SocialMF learns the latent feature vectors of users and items. Different from STE, the feature vector of each user is dependent on the feature vectors of his direct neighbors in the social network.

This allows SocialMF to handle the transitivity of trust and trust propagation, which is not captured by the STE model. Trust propagation has been shown to be a crucial phenomenon in the social sciences, in social network analysis and in trust-based recommendation. Also if a user has not expressed any ratings, his feature vectors can be learnt as long as he is connected to the social network via a social relation. Thus SocialMF deals better with cold start users than existing methods. Note that if a cold start user is not connected to the social network, then social network based approaches have no additional information to improve the quality of recommendation for that user.

[11] In Flixster, we do not take into account the users with no ratings in this statistics.

Experiments on two real life data sets from Epinions and Flixster demonstrate that SocialMF outperforms existing methods for social network based recommendation. This work suggests several interesting directions for future work. We want to extend the model to handle negative trust relations, since some social networks allow users to express distrust towards other users. Also, cold start items have not been addressed in this paper, and it should be explored how the model can be extended so that the feature vectors of cold start items are also learnt efficiently.

References

[Crandall *et al.*, 2008] David Crandall, Dan Cosley, Daniel Huttenlocher, Jon Kleinberg, and Siddharth Suri. Feedback effects between similarity and social influence in online communities. In *KDD'08*, 2008.

[Friedkin, 1998] Noah E. Friedkin. *A Structural Theory of Social Influence*. Cambridge University Press, 1998.

[Golbeck, 2005] Jennifer Golbeck. *Computing and Applying Trust in Web-based Social Networks*. PhD thesis, University of Maryland College Park, 2005.

[Goldberg *et al.*, 1992] David Goldberg, David Nichols, Brian M. Oki, and Douglas Terry. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12), 1992.

[Jamali and Ester, 2009] Mohsen Jamali and Martin Ester. Trustwalker: A random walk model for combining trust-based and item-based recommendation. In *KDD'09*, 2009.

[Jamali and Ester, 2010] Mohsen Jamali and Martin Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM conference on Recommender systems*, RecSys '10, pages 135–142, 2010.

[Koren *et al.*, 2009] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *IEEE Computer*, 42:30–37, 2009.

[Koren, 2008] Yehuda Koren. Factorization meets the neighborhood a multifaceted collaborative filtering model. In *KDD'08*, Las Vegas, USA, 2008.

[Ma *et al.*, 2009a] Hao Ma, Irwin King, and Michael R. Lyu. Learning to recommend with social trust ensemble. In *SIGIR*, pages 203–210, 2009.

[Ma *et al.*, 2009b] Hao Ma, Michael R. Lyu, and Irwin King. Learning to recommend with trust and distrust relationships. In *Proceedings of the third ACM conference on Recommender systems*, RecSys '09, pages 189–196, 2009.

[Richardson and Domingos, 2002] Matthew Richardson and Pedro Domingos. Mining knowledge-sharing sites for viral marketing. In *KDD'02*, 2002.

[Salakhutdinov and Mnih, 2008] Ruslan Salakhutdinov and Andriy Mnih. Probabilistic matrix factorization. In *NIPS'08*, volume 20, 2008.

[Wasserman and Faust, 1994] Stanley Wasserman and Katherine Faust. *Social Network Analysis*. Cambridge Univ. Press, 1994.