Exploiting Implicit Item Relationships for Recommender Systems

Zhu Sun, Guibing Guo*, and Jie Zhang

School of Computer Engineering, Nanyang Technological University, Singapore
*School of Information Systems, Singapore Management University, Singapore
{sunzhu,zhangj}@ntu.edu.sg, *gbguo@smu.edu.sg

Abstract. Collaborative filtering inherently suffers from the *data sparsity* and *cold start* problems. Social networks have been shown useful to help alleviate these issues. However, social connections may not be available in many real systems, whereas implicit item relationships are lack of study. In this paper, we propose a novel matrix factorization model by taking into account implicit item relationships. Specifically, we employ an adapted association rule technique to reveal implicit item relationships in terms of item-to-item and group-to-item associations, which are then used to regularize the generation of low-rank user- and item-feature matrices. Experimental results on four real-world datasets demonstrate the superiority of our proposed approach against other counterparts.

1 Introduction

Recommender systems have become a prevalent tool to help satisfy users' need of personalization over the exponentially increasing amount of information on Web 2.0. Collaborative filtering (CF) is a widely accepted recommendation technique built upon the concept of user (or item) similarity. That is, a user's preference can be inferred by aggregating the taste of similar users. However, CF inherently suffers from the *data sparsity* and *cold start* problems [1].

To address these issues, trust-aware recommender systems [2–4, 1, 5] are emerging with the advent of social networks. Many recently proposed approaches are designed upon the matrix factorization technique [6]. The intuition behind is that social friends share similar preferences and influence each other by recommending items. It has been shown that such additional side information among users is useful to deal with the concerned issues and thus to improve recommendation performance. However, the reliance on social connections may restrict the application of trust-based approaches to other scenarios where social networks are not available or supported. The potential noise and weaker social ties (than trust) in social networks can further hinder the generality of these approaches [7].

Similarly, the side information of items is also exploited for recommender systems, given its effectiveness in improving recommendation performance [6]. The basic assumption is that users tend to have similar preferences towards a set of *associated* items. For example, a person is likely to enjoy the movie series of *The Lord of the Rings*, and possibly appreciates the associated background soundtracks. A number of approaches [8–10] have been proposed by making use of explicit item relationships such as category, genre, location, etc. However, similar as additional social information, items' side information may be unavailable for some real applications, or it is prohibitively expensive (or time-consuming) to extract the side information due to the large volume of items. Furthermore, only few works [6, 11] have considered and demonstrated the value of implicit item relationships for recommender systems.

In this paper, we propose a novel matrix factorization model by exploiting association rule-based implicit item relationships, called *IIR*. It is developed merely based on user-item rating information, and requires no reliance of additional user or item side information. We ascribes this feature to the essential difference from other literature studies. Specifically, we employ an adapted associate rule technique to reveal the implicit item relationships in the form of item-to-item and group-to-item associations, which are then used to regularize the generation of low-rank user- and item-feature matrices in the proposed IIR model. In addition, we design four different strategies to select the most reliable item associations to train the model. Experimental results on four real-world datasets show that our approach achieves superior performance against other counterparts, and that group-to-item associations are more effective than item-to-item associations.

2 Related Work

Additional side information is often incorporated in collaborative filtering to improve recommendation performance. We give a brief overview below regarding such kind of recommendation approaches from the perspectives of users and items, respectively. First, a notable research field is the trust-aware recommender systems which take into account additional user relationships. Many approaches have been proposed to date. Ma et al. [2] propose the *RSTE* approach by linearly combining a basic matrix factorization model and a trust-based neighborhood approach. The same authors later find that using social information as a regularizer works better than using it to decompose the user-item rating matrix [4]. This finding is endorsed by Jamali and Ester [3] where a user's latent feature vector is regularized by those of her trusted users. Therefore, to incorporate item relationships in the IIR model, we follow the same rule to utilize item relationships to regularize the generation of items' latent feature vectors. More recent works consider more aspects of social trust such as implicit trust influence [5, 7], etc. However, trust-based approaches may fail to work if being applied to the situations where social networks are not built-in or connected. Our work is intended for a more general case where only user-item ratings exist.

Second, some researchers also attempt to make use of item relationships to enhance recommender systems. In this paper, we classify two kinds of item relationships: *explicit* and *implicit*. Typical examples of explicit item relationships include items' extrinsic properties: category, location and tag, just to name a few. For example, Hu et al. [10] contend that the quality of a business shop has some implicit indication on that of other shops in a certain geographical neighborhood. Shi et al. [9] show that tags can be used to bridge cross-domain knowledge to provide better recommendation. Implicit item relationships refer to the relationships that cannot be explicitly observed between items. A classic example is that a man buying diaper is likely to buy beer as well, though the diaper and beer are distinct items. Kim and Kim [8] apply association rule mining techniques to reveal multi-level item associations in the light of item categories. Our approach also adopts association rule to identify implicit item relationships, but differs in that we do not classify item associations to multiple levels in terms of category. Instead, we consider simple rules (one item indicating another, or item-to-item) and then generalize to group-to-item (a group of items indicating another item) associations. Wang et al. [11] identify item relationships using a similarity measure, and claim that an item's feature vector can be influenced by those of other similar items. However, they build the item-item similarity matrix with some ad-hoc settings when item similarity equals 0 or uncomputable. Another issue is that their model is very time-consuming in training and thus prevents from being applied to large-scale datasets. In this paper, we are more interested in association rule-based item relationships rather than item similarity, the explanation of which is deferred to Section 3.2.

3 Recommendation with Implicit Item Relationships

In this section, we first introduce the IIR recommendation model, and then elaborate how item-to-item associations can be identified by an adapted association rule technique, followed by the generalization to group-to-item associations.

3.1 The IIR Model

Matrix factorization (MF) techniques have been widely applied in recommender systems. The basic assumption is that a user's preference can be characterized by a few number of latent features. In particular, MF models [6] factorize the user-item rating matrix $R \in \mathbb{R}^{m \times n}$ into two low-rank user-feature $U \in \mathbb{R}^{m \times d}$ and item-feature $V \in \mathbb{R}^{n \times d}$ matrices, where m, n are the number of users and items, respectively; and $d \ll \min(m, n)$ is the number of latent features. The rating matrix R is very sparse due to the fact that a user generally only rates a small portion of items. Let $r_{u,i}$ be a rating given by user u on item i, and $\hat{r}_{u,j}$ be a rating prediction for user u on a target item j. We preserve symbols u, v, p for users, and i, j, k for items. The rating prediction $\hat{r}_{u,j}$ can be estimated by the inner product of user-specific feature vector U_u and item-specific feature vector V_j, given by:

$$\hat{r}_{u,j} = U_u^\top V_j,$$

where the matrices U and V can be learned by minimizing the differences between the prediction and the ground truth over all the users and items.

In this paper, we propose the IIR model which focuses on the incorporation of implicit item relationships in order to better factorize the rating matrix R. Assume that two items j, k are implicitly associated with the strength of $s_{j,k}$. In this case, we contend that the two items should be close to each other in terms of feature vectors. Inspired by the usage of social relationships in [4], we devise the implicit item relationships as a regularizer to adjust the decomposition of rating matrix R. Specifically, the IIR model is given as follows.

$$\mathcal{L} = \frac{1}{2}\sum_{u=1}^{m}\sum_{j=1}^{n} I_{u,j}(r_{u,j} - \hat{r}_{u,j})^2 + \frac{\alpha}{2}\sum_{j=1}^{n}\sum_{k \in A_j} s_{k,j}\|V_j - V_k\|_F^2 + \frac{\lambda_u}{2}\|U\|_F^2 + \frac{\lambda_v}{2}\|V\|_F^2 \quad (1)$$

where $I_{u,j}$ is an indicator function that equals 1 if user u rated item j and equals 0 otherwise; $\alpha > 0$ is a regularization parameter to control the importance of regularization by implicit item relationships; $s_{k,j}$ indicates the extent to which item k is associated with item j; A_j is a set of reliable association rules for item j; $\|\cdot\|_F$ is the Frobenius norm; and λ_u, λ_v are regularization parameters to avoid over-fitting. Suppose that we have identified two implicit item relationships from i to j and from k to j, the IIR model will add the following two regularizers:

$$s_{i,j} \parallel V_j - V_i \parallel_F^2 \quad \text{and} \quad s_{k,j} \parallel V_j - V_k \parallel_F^2$$

In other words, the IIR model can indirectly minimize the differences between the feature vectors of related items i and k to some extent. We treat it as an advantage of our model to capture both the influence of direct and indirect implicit item relationships.

Equation 1 indicates that it is necessary for our model to effectively identify the set of item associations A_j and the corresponding strength with other items $s_{k,j}$. We proceed to describe how to achieve them in the next subsection.

3.2 Mining Implicit Item Relationships

Item Similarity. A straightforward method to define implicit item relationships is item similarity. If many users like both items, it indicates that the two items have some similarity in common. This intuition underpins the well-known item-based collaborative filtering. The most popular similarity measures are the Pearson correlation coefficient (PCC) and cosine similarity (COS). However, with the following concerns, we believe that item similarity measures are not suitable for our work. First, a key characteristic of similarity measures is symmetry, i.e., $s_{j,k} = s_{k,j}$. In our case, we would like to distinguish the influence of items j to k from that of items k to j. For example, a man buying beer may not buy diaper, though a man buying diaper is likely to buy beer. Second, the computation of similarity measures is generally based on the overlapping ratings between two rating vectors. However, we argue that the non-overlapping ratings may help define the differences between the two items. Third, PCC and COS may produce misleading similarity measurements as pointed out by Guo et al. [12], especially when the size of overlapping ratings is small. Lastly, similarity measures often consider the correlation between two individual items, whereas we intend to measure more generalized correlations between a group of items and a single item (i.e., group-to-item) other than item-to-item implicit relationships.

Item-to-Item Associations. We define the implicit item relationships as the item associations between a target item and another item or a set of other items. This definition directs us to adopt association rule techniques to measure item associations. The association rule mining is to search the associated item pairs that often co-occur in transaction events. Assume that item i appears frequently together with item j, and an association rule can then be denoted by $l_{i,j} : i \rightarrow j$. Note that the occurrence of item j may not frequently indicate the occurrence of item i, i.e., the association is asymmetric as we demand. Generally, an association rule is valid only if its *support* and *confidence* (indicating the usefulness and certainty) are greater than a user-specified minimum support (denoted by *minSup*) and confidence (denoted by *minCon*) thresholds, respectively.

By regarding the rating matrix as a user-item transaction matrix, we can apply association rules to reveal implicit relationships among items. Generally, we need to determine proper values for thresholds *minSup* and *minCon*. The settings are not trivial considering that: (1) higher thresholds will decrease the available number of association rules due to the sparseness of rating matrix; and (2) lower thresholds will result in too many association rules, which may significantly slow the model training and take into account many unreliable association rules. Both issues can greatly deteriorate recommendation performance. Instead of empirically tuning the two parameters, we define a new measure *reliability* as the extent to which a mined association rule is reliable in terms of both support and confidence measures. The reliability of an association rule $l_{i,j}$ is defined as:

$$\text{reliability}(l_{i,j}) = \frac{\text{support}(l_{i,j})}{\text{support}(l_{i,j}) + C} * \text{confidence}(l_{i,j}), \tag{2}$$

where $\text{support}(l_{i,j})$ and $\text{confidence}(l_{i,j})$ are the support and confidence measures of an association rule $l_{i,j}$, respectively; and C is a constant to adapt the importance of association rule support. This formulation produces high reliability only if both support and confidence values are high. Then, we sort all the association rules in the descending order of computed reliability values, and select the top-K most reliable association rules to form the association set A_j for target item j.

Group-to-Item Associations. A natural generalization to item-to-item associations is group-to-item associations. That is, we further consider whether a set of items can be associated with a specific item. We represent such kind of association rules as $l_{G,j} : G \rightarrow j$, where G denotes a set of associated items. Similarly, this kind of association rules can be also identified by applying association rule techniques as well as the computation of association rule reliability. Hence, the objective function in Equation 1 can be rewritten as follows:

$$\begin{aligned}\mathcal{L} =& \frac{1}{2}\sum_{u=1}^{m}\sum_{j=1}^{n} I_{u,j}(r_{u,j} - \hat{r}_{u,j})^2 + \frac{\alpha}{2}\sum_{j=1}^{n}\sum_{G\in A_j} s_{G,j}\|V_j - |G|^{-0.5}\sum_{k\in G} V_k\|_F^2 \\ &+ \frac{\lambda_u}{2}\|U\|_F^2 + \frac{\lambda_v}{2}\|V\|_F^2,\end{aligned} \tag{3}$$

where $s_{G,j}$ is the strength (i.e., reliability) of association rule $l_{G,j}$. Note that we represent the characteristic of a group G by the average of all group items' feature vectors. In other words, we constrain that a target item's feature vector should be close to the majority of its associated group.

For simplicity, hereafter we only consider group-to-item associations with group size 2, i.e., $|G| = 2$.[1] This is due to that, $|G| = 0$ indicates no items are associated with item j while $|G| = 1$ implies that group-to-item associations are equivalent with item-to-item associations. Now that there are two kinds of item association rules, we propose the following four strategies to select the association neighborhood of item j, i.e., A_j.

Half: select half a number of group-to-item and the other half a number of item-to-item association rules separately as the baseline strategy.
Mix: select the top-K most reliable association rules after sorting all kinds of association rules in term of reliability values.
Group: select the top-K most reliable group-to-item association rules only, and ignore all item-to-item association rules.
Group+: select the top-K most reliable group-to-item association rules. In case of insufficient rules, we select item-to-item association rules to complement.

Note that we do not incorporate the strategy of *selecting item-to-item association rules only* by setting $|G| = 1$. The reason is that, by definition the strength of item-to-item association rules is generally weaker than that of group-to-item association rules. Hence, incorporating weaker rules will have smaller effect in regularizing the objective function, and result in less-performing recommendations. The experimental results on real-world datasets have also confirmed our intuition (see Section 4.2).

3.3 Model Learning

The stochastic gradient descent (SGD) method is widely used to achieve a local minimum of the objective function given by Equation 3. Specifically, the SGD update rules for variables U_u and V_j of the IIR[2] model are given as follows.

$$
\begin{aligned}
\frac{\partial \mathcal{L}}{\partial U_u} =& \sum_{j=1}^{n} I_{u,j}\big(g(U_u^\top V_j) - r_{u,j}\big) g'(U_u^\top V_j) V_j + \lambda_u U_u, \\
\frac{\partial \mathcal{L}}{\partial V_j} =& \sum_{i=1}^{m} I_{u,j}\big(g(U_u^\top V_j) - r_{u,j}\big) g'(U_u^\top V_j) U_u + \lambda_v V_j \\
&+ \alpha \sum_{G \in A_j} s_{G,j} \Big(V_j - |G|^{-0.5} \sum_{i \in G} V_i\Big) - \alpha \sum_{\substack{M \,\in\, A_k, \\ j \,\in\, M}} \frac{s_{M,k}}{\sqrt{|M|}} \Big(V_k - |M|^{-0.5} \sum_{g \in M} V_g\Big),
\end{aligned}
$$

where $g(x) = 1/(1 + exp(-x))$ is a logistic function used to bound the value rang of rating prediction into $[0, 1]$, and $g'(x)$ is the derivative of function $g(x)$.

[1] We empirically noted that groups with size greater than 2 did not provide visibly better performance or even provide worse performance sometimes. We are aware that the observed effect may not be the same on other datasets we did not use.

[2] Source code is included in the Librec library at www.librec.net.

Table 1. Statistics of the used datasets

Dataset	#Users	#Items	#Ratings	Sparsity
FilmTrust	1058	2071	35,497	98.86%
MoiveLens	943	1682	100,000	93.70%
Ciao	8000	9749	38,591	99.95%
Epinions	7941	10,000	107,552	99.86%

To be consistent, we adopt the max-min normalization approach to convert the observed ratings to the same value range $[0, 1]$.

4 Experiments and Results

4.1 Experimental Setup

Datasets. Four real-world datasets are used in our experiments, namely FilmTrust[3], MoiveLens[4], Ciao[3] and Epinions[5]. FilmTrust is a movie sharing website that allows users to assign numerical ratings (scaled from 0.5 to 4.0 with step 0.5) to movies. MoiveLens is a personalized movie recommendation website, where users can rate movies with integers from 1 to 5. The data set has been preprocessed such that each user has rated at least 20 items. Both Ciao and Epinions are product review sites, where consumers can review various products with ratings from 1 to 5 stars. The statistics of our datasets is presented in Table 1.

Comparison Methods. We compare our **IIR** model with the following methods: (1) **PMF** [13] is a basic matrix factorization method without any additional side information. (2) **IR-P** [11] is the item relationship-based approach where PCC is used to compute item similarity.[6] For fair comparison, we remove the influence of social networks from the original model. (3) **IR-I** is a substitute of IR-P by replacing PCC with item-to-item associations. (4) **IIR-C** is a variant of the IIR model which adopts COS to identify implicit item relationships. (5) **IIR-I** is our approach merely based on item-to-item relationships. (6) **IIR-G** is our approach with the incorporation of group-to-item relationships.

Evaluation Metrics. We perform 5-fold cross validation in our experiments. Specifically, we randomly split each dataset into five folds and in each iteration four folds are used as the training set and the remaining fold as the test set. All folds will be tested, and the average results are reported as the final performance. Predictive performance is evaluated by two widely used measures: the mean absolute error (MAE) and root mean square error (RMSE), defined by:

$$\text{MAE} = \frac{\sum_{u,j} |r_{u,j} - \hat{r}_{u,j}|}{N}, \qquad \text{RMSE} = \sqrt{\frac{\sum_{u,j}(r_{u,j} - \hat{r}_{u,j})^2}{N}}$$

[3] `http://www.librec.net/datasets.html`
[4] `http://www.cs.umn.edu/Research/GroupLens`
[5] `http://www.trustlet.org/wiki/Epinions`
[6] We convert the original PCC value from [-1,1] to [0,1] by function $f(x) = (x + 1)/2$.

where N is the number of test ratings. Smaller MAE and RMSE values imply better predictive accuracy.

Parameter Settings. The number of latent features d is selected in $\{5, 10, 20, 50\}$. We empirically find that the following parameter settings can help achieve the best performance for each comparison method. For **IR-P** and **IR-I**, the importance of item relationship-based rating prediction is set $0.005, 0.01, 0.1, 0.005$ corresponding to FilmTrust, MovieLens, Ciao and Epinions, respectively.[7] All these item relationship-based approaches reach the best performance when the size of association neighborhood is set 50, i.e., $K = 50$. For all the methods, we apply grid search in $\{0.00001, 0.0001, 0.001, 0.01, 0.1\}$ for regularization parameters λ_u, λ_v, and in $\{0.0001, 0.001, 0.01\}$ for the learning rate.

4.2 Results and Analysis

Effect of Parameters C and K. In our approach, C controls the importance of association rule support. We apply a grid search in {0,10,20,50,100,150,200} to find the optimal setting for parameter C. The results are shown in Figure 1, where the best settings are around 100.[8] To select the top-K most reliable association neighbors for each item j (see Equation 3), we tune the value of K from 0 to 100 stepping by 10, where $K = 0$ indicates that no implicit item relationships are considered, i.e., our model is degrading to the basic matrix factorization model. The performance is illustrated in Figure 2.[8] The results show that adding implicit item relationships can help improve predictive performance ($K > 0$), but further improvements tend to be negligible when $K > 50$ across all the datasets. In other words, the top-50 association rules have the most important influence. For the association rules after that, their strength of associations with the target item may be too small to have a visible impact.

Effect of Parameter α. The parameter α in Equation 3 controls the importance of item relationship regularization. We apply a grid search in {0.0001,0.001,0.01, 0.1,0.2,0.5,1.0} to find the optimal setting for parameter α. The results are plotted in Figure 3, where the best settings for parameter α are around 0.5 on the Epinions dataset and $0.1 \sim 0.2$ on the other datasets. Note that for the IIR-G method, similar trends are obtained when any one of the four strategies is adopted to select group-to-item association rules (see Section 3.2).

Effect of Four Strategies. We have identified four different association rule selection strategies for our approach IIR-G in Section 3.2, namely Half, Mix, Group and Group+. Table 2 summarizes the performance obtained by applying the four strategies to all the datasets and across the different number of latent features d. The results are consistent across all the cases, and demonstrate that (1) Half reaches the poorest performance; (2) Group works better than Mix; and (3) Group+ achieves the best performance. These results confirm our previous

[7] Due to space limitation, we do not present the results of tuning this parameter.

[8] For simplicity, we only present the results when $d = 5$, and similar trends are observed on other settings including d values and IIR variants.

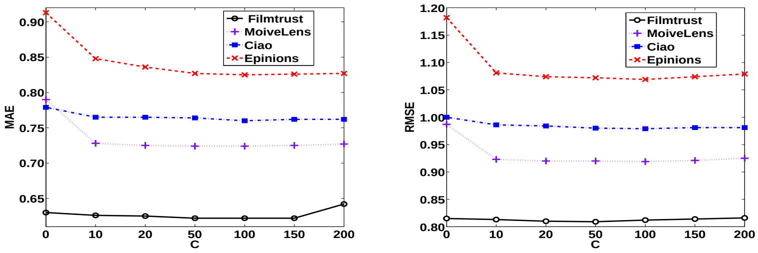

(a) IIR-I (MAE) (b) IIR-I (RMSE)

Fig. 1. The effect of parameter C in our approaches IIR-I ($d = 5$)

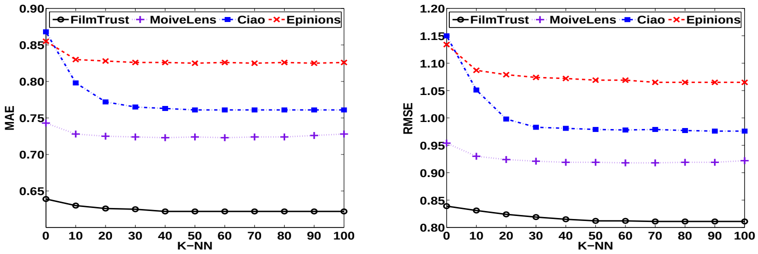

(a) MAE (b) RMSE

Fig. 2. The effect of number of association rules K in our approach IIR-I ($d = 5$)

claim, that is, the strength of group-to-item associations is stronger than that of item-to-item associations. When only half of group-to-item associations are used, the performance is the worst; as more group-to-item associations are selected (by Mix, Group and Group+), the performance is improved accordingly. It suggests that we should always first select stronger group-to-item associations, and adopt relatively weaker associations only if stronger ones do not suffice.

Comparison with Other Methods. Table 3 presents the results of all comparison methods on the four real-world datasets, where the best performance is highlighted in bold and the second best performance among the first four methods[9] is denoted by * symbol. A number of interesting observations can be noted from the present results. First, all the methods exploiting item relationships perform better than the basic PMF method, indicating the usefulness of incorporating item relationships for recommender systems. Second, IR-I consistently obtains lower values of MAE and RMSE than IR-P, implying that item-to-item associations are more effective than item similarity computed by similarity measures. This is further confirmed by the fact that IIR-I outperforms IIR-C. In Section 3.2, we explained in detail why item similarity may not be suitable to reveal implicit item relationships. Third, among the three variants of the IIR

[9] We tend to treat IIR-C as a baseline as it uses item similarity-based rather than association rule-based implicit item relationships.

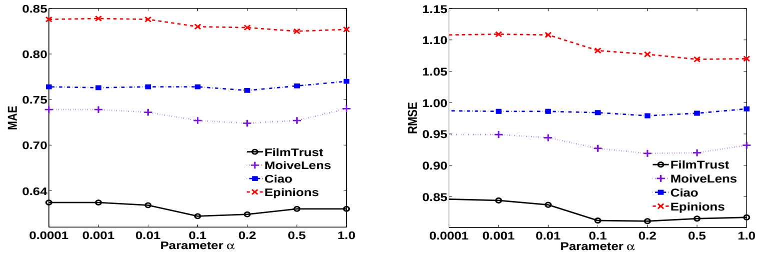

(a) IIR-I (MAE) (b) IIR-I (RMSE)

Fig. 3. The effect of regularization parameter α in our approaches IIR-I

Table 2. The effect of four strategies to select association rules for our approach IIR-G

d	**Metrics**	**FilmTrust**				**MovieLens**			
		Half	**Mix**	**Group**	**Group+**	**Half**	**Mix**	**Group**	**Group+**
5	MAE	0.621	0.620	0.619	**0.615**	0.724	0.721	0.721	**0.718**
	RMSE	0.809	0.810	0.809	**0.807**	0.914	0.916	0.914	**0.911**
10	MAE	0.615	0.614	0.614	**0.611**	0.720	0.717	0.716	**0.712**
	RMSE	0.804	0.805	0.804	**0.802**	0.909	0.908	0.907	**0.903**
20	MAE	0.611	0.611	0.610	**0.608**	0.718	0.714	0.714	**0.707**
	RMSE	0.798	0.800	0.796	**0.792**	0.906	0.905	0.904	**0.899**
50	MAE	0.610	0.606	0.605	**0.603**	0.716	0.713	0.711	**0.703**
	RMSE	0.794	0.792	0.790	**0.789**	0.904	0.902	0.901	**0.896**

d	**Metrics**	**Ciao**				**Epinions**			
		Half	**Mix**	**Group**	**Group+**	**Half**	**Mix**	**Group**	**Group+**
5	MAE	0.757	0.754	0.754	**0.750**	0.827	0.824	0.823	**0.820**
	RMSE	0.977	0.976	0.977	**0.969**	1.067	1.066	1.064	**1.063**
10	MAE	0.746	0.746	0.744	**0.741**	0.826	0.823	0.821	**0.817**
	RMSE	0.968	0.970	0.962	**0.960**	1.064	1.063	1.058	**1.056**
20	MAE	0.738	0.739	0.738	**0.732**	0.823	0.820	0.817	**0.816**
	RMSE	0.959	0.960	0.957	**0.952**	1.059	1.056	1.056	**1.053**
50	MAE	0.731	0.731	0.730	**0.728**	0.819	0.816	0.815	**0.812**
	RMSE	0.954	0.956	0.950	**0.946**	1.054	1.050	1.050	**1.050**

model, IIR-G achieves the best performance from which we may conclude that: group-to-item associations are stronger than item-to-item associations which are then preferred to similarity-based item associations. Last, our approach IIR-G reaches the superior performance to the other counterparts, and the percentages of improvements relative to other baselines are put to the last column in Table 3. On the average, the percentages of relative improvements across different number of latent features are summarized as follows in terms of (MAE, RMSE) pair: (1.97%, 1.30%) on FilmTrust, (2.48%, 2.38%) on MovieLens, (3.25%, 2.50%) on Ciao, (1.44%, 2.06%) on Epinions and (2.29%, 2.06%) over all the datasets. Koren [6] has pointed out that even small improvements in predictive accuracy can have great impact on real applications. Hence, we claim that our approach obtains important improvements by exploiting implicit item relationships.

Table 3. The experimental results on the four datasets, where * indicates the best performance among the first four methods, and the column "Improve" indicates the relative improvements that our approaches achieve relative to the * results.

Dataset	d	Metrics	PMF	IR-P	IR-I	IIR-C	IIR-I	IIR-G	Improve
FilmTrust	5	MAE	0.639	0.630	0.627	0.627*	0.622	**0.615**	1.91%
		RMSE	0.839	0.835	0.825	0.818*	0.812	**0.807**	1.34%
	10	MAE	0.638	0.629	0.622*	0.623	0.617	**0.611**	1.77%
		RMSE	0.837	0.832	0.812	0.811*	0.805	**0.802**	1.11%
	20	MAE	0.636	0.625	0.620*	0.621	0.613	**0.608**	1.94%
		RMSE	0.830	0.826	0.807	0.805*	0.799	**0.792**	1.61%
	50	MAE	0.632	0.620	0.617*	0.618	0.618	**0.603**	2.27%
		RMSE	0.824	0.814	0.798*	0.798	0.790	**0.789**	1.13%
MoiveLens	5	MAE	0.743	0.737	0.735	0.734*	0.724	**0.718**	2.18%
		RMSE	0.954	0.951	0.947	0.930*	0.919	**0.911**	2.04%
	10	MAE	0.742	0.737	0.730*	0.732	0.720	**0.712**	2.47%
		RMSE	0.944	0.946	0.935	0.927*	0.913	**0.903**	2.59%
	20	MAE	0.735	0.728	0.725*	0.726	0.715	**0.707**	2.48%
		RMSE	0.935	0.931	0.924*	0.925	0.906	**0.899**	2.71%
	50	MAE	0.733	0.725	0.723*	0.724	0.713	**0.703**	2.77%
		RMSE	0.921	0.924	0.919	0.916*	0.903	**0.896**	2.18%
Ciao	5	MAE	0.868	0.788	0.779	0.770*	0.760	**0.750**	2.60%
		RMSE	1.150	1.101	1.081	0.991*	0.979	**0.969**	2.22%
	10	MAE	0.832	0.785	0.775	0.766*	0.749	**0.741**	3.26%
		RMSE	1.134	1.096	1.078	0.984*	0.970	**0.960**	2.44%
	20	MAE	0.828	0.770	0.765	0.760*	0.740	**0.732**	3.68%
		RMSE	1.126	1.074	1.060	0.978*	0.959	**0.952**	2.66%
	50	MAE	0.823	0.769	0.762	0.754*	0.732	**0.728**	3.45%
		RMSE	1.124	1.067	1.050	0.972*	0.952	**0.946**	2.67%
Epinions	5	MAE	0.855	0.845	0.838	0.836*	0.825	**0.820**	1.91%
		RMSE	1.134	1.130	1.123	1.090*	1.069	**1.063**	2.48%
	10	MAE	0.848	0.843	0.833	0.833*	0.824	**0.817**	1.92%
		RMSE	1.128	1.120	1.112	1.082*	1.063	**1.056**	2.40%
	20	MAE	0.846	0.836	0.826*	0.830	0.822	**0.816**	1.21%
		RMSE	1.117	1.094	1.090	1.075*	1.058	**1.053**	2.05%
	50	MAE	0.839	0.826	0.818*	0.828	0.816	**0.812**	0.73%
		RMSE	1.106	1.070	1.064*	1.065	1.051	**1.050**	1.32%

5 Conclusion and Future Work

This paper proposed a novel matrix factorization model that incorporated the influence of implicit item relationships for recommender systems. We introduced an adapted association rule technique to reveal implicit item relationships, and justified that item similarity may not be a good measure for implicit item relationships. A new measure *reliability* of an association rule was defined to help sort and select item associations. We investigated not only item-to-item associations but also generalized group-to-item associations. Four strategies were designed to choose the most reliable set of association rules, which were used to regularize the generation of low-rank user- and item-feature matrices. Empirical results on four real-world datasets demonstrated that our approach gained important improvements relative to other comparison methods. For future work, we intend to incorporate both explicit and implicit item relationships to further improve recommendation performance.

References

1. Guo, G., Zhang, J., Thalmann, D.: A simple but effective method to incorporate trusted neighbors in recommender systems. In: Proceedings of the 20th International Conference on User Modeling, Adaptation and Personalization (UMAP). 114–125 (2012)
2. Ma, H., King, I., Lyu, M.: 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 (2009)
3. Jamali, M., Ester, M.: A matrix factorization technique with trust propagation for recommendation in social networks. In: Proceedings of the 4th ACM Conference on Recommender Systems (RecSys). 135–142 (2010)
4. Ma, H., Zhou, D., Liu, C., Lyu, M., King, I.: Recommender systems with social regularization. In: Proceedings of the 4th ACM International Conference on Web Search and Data Mining (WSDM). 287–296 (2011)
5. Yang, B., Lei, Y., Liu, D., Liu, J.: Social collaborative filtering by trust. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI). 2747–2753 (2013)
6. Koren, Y.: 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 (2008)
7. Guo, G., Zhang, J., Yorke-Smith, N.: Trustsvd: Collaborative filtering with both the explicit and implicit influence of user trust and of item ratings. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI). (2015)
8. Kim, C., Kim, J.: A recommendation algorithm using multi-level association rules. In: Proceedings of 2003 IEEE/WIC International Conference on Web Intelligence (WI). 524–527 (2003)
9. Shi, Y., Larson, M., Hanjalic, A.: Tags as bridges between domains: Improving recommendation with tag-induced cross-domain collaborative filtering. In: Proceedings of the 19th International Conference on User Modeling, Adaptation and Personalization (UMAP). 305–316 (2011)
10. Hu, L., Sun, A., Liu, Y.: Your neighbors affect your ratings: on geographical neighborhood influence to rating prediction. In: Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR). 345–354 (2014)
11. Wang, D., Ma, J., Lian, T., L., G.: Recommendation based on weighted social trusts and item relationships. In: Proceedings of the 29th Annual ACM Symposium on Applied Computing (SAC). 2000–2005 (2014)
12. Guo, G., Zhang, J., Yorke-Smith, N.: A novel Bayesian similarity measure for recommender systems. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI). 2619–2625 (2013)
13. Salakhutdinov, R., Mnih, A.: Probabilistic matrix factorization. In: Advances in Neural Information Processing Systems (NIPS). Volume 20. 1257–1264 (2008)