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

GBPR: Group Preference Based Bayesian Personalized Ranking for One-Class Collaborative Filtering

Weike Pan and **Li Chen**
Department of Computer Science
Hong Kong Baptist University, Hong Kong
{wkpan, lichen}@comp.hkbu.edu.hk

Abstract

One-class collaborative filtering or collaborative ranking with implicit feedback has been steadily receiving more attention, mostly due to the "one-class" characteristics of data in various services, e.g., "like" in Facebook and "bought" in Amazon. Previous works for solving this problem include pointwise regression methods based on absolute rating assumptions and pairwise ranking methods with relative score assumptions, where the latter was empirically found performing much better because it models users' ranking-related preferences more directly. However, the two fundamental assumptions made in the pairwise ranking methods, (1) *individual pairwise preference over two items* and (2) *independence between two users*, may not always hold. As a response, we propose a new and improved assumption, *group Bayesian personalized ranking* (GBPR), via introducing richer interactions among users. In particular, we introduce *group preference*, to relax the aforementioned *individual* and *independence* assumptions. We then design a novel algorithm correspondingly, which can recommend items more accurately as shown by various ranking-oriented evaluation metrics on four real-world datasets in our experiments.

1 Introduction

In industry, recommender system as a critical engine in various online entertainments and shopping services has caught much attention and contributed significant revenue growth in recent years. Lots of internet, electronic and telecom giants embed recommendation technologies in their existing systems to increase user engagement and business revenues from more product or advertisement sales.

In academia, most research studies in recommender systems are somehow biased to the Netflix competition[1], partially due to the public availability of the data. The Netflix contest can be categorized as a "multi-class" recommendation problem, where the inputs are categorical scores. For example, the input data of Netflix are ratings with "1" for bad, "2" for fair, "3" for good, "4" for excellent, and "5" for perfect, for which various algorithms are proposed in order to predict the users' preference scores accurately. Such categorical data contain both positive feedback of "4" and "5", and negative feedback of "1" and "2", where some regression or ranking loss functions were designed to fit the score or to preserve the ordering, among which matrix-factorization based models have been shown to be the most effective solutions [Weimer *et al.*, 2007; Koren, 2008; Rendle, 2012; Liu *et al.*, 2013].

However, in most applications, the collected data of user behaviors are in "one-class" form rather than multi-class form, e.g., "like" in Facebook, "bought" in Amazon, and "click" in Google Advertisement. Such data are usually called implicit [Nichols, 1997; Rendle *et al.*, 2009] or one-class [Pan *et al.*, 2008] feedback. The one-class collaborative filtering problem is different from that of 5-star rating prediction in the Netflix competition, since the former only contains positive feedback rather than both positive feedback and negative feedback in the contest, and the goal is item ranking instead of rating prediction.

For solving the one-class collaborative filtering problem, previous matrix-factorization based algorithms can be roughly summarized in two manners, (1) pointwise regression methods, and (2) pairwise ranking methods. The former learn a latent representation of both users and items via minimizing a pointwise square loss [Pan *et al.*, 2008; Hu *et al.*, 2008] to approximate the absolute rating scores, while the latter take pairs of items as basic units and maximize the likelihood of pairwise preferences over observed items and unobserved items [Rendle *et al.*, 2009]. Empirically, the pairwise ranking method [Rendle *et al.*, 2009] achieves much better performance [Du *et al.*, 2011] than pointwise methods [Pan *et al.*, 2008; Hu *et al.*, 2008], and has been successfully adopted in many applications, e.g., tag recommendation [Rendle and Schmidt-Thieme, 2010], news recommendation [Yang *et al.*, 2011], and shopping items recommendation [Kanagal *et al.*, 2012].

In this paper, we study the two fundamental assumptions made in the seminal work of pairwise ranking method [Rendle *et al.*, 2009], and point out its limitations. As a response, we introduce *group preference* instead of *individual preference* alone in [Rendle *et al.*, 2009], in order to inject richer interactions among users and thus relax the *individual* and

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

independence assumptions. We then propose a new and improved assumption called *group Bayesian personalized ranking* (GBPR) and design an efficient algorithm correspondingly. Empirically, we find that our new assumption digests the one-class data more effectively and achieves better recommendation performance on all four real-world datasets in our experiments.

2 Related Work

In this section, we discuss some related works of one-class collaborative filtering in two branches of work, (1) pointwise methods with absolute preference assumptions, and (2) pairwise methods with relative preference assumptions.

Pointwise methods with absolute preference assumptions Pointwise methods take implicit feedback as *absolute* preference scores. For example, an observed user-item pair, (u, i), is interpreted as that user u likes item i with a high absolute score, e.g., 1. OCCF (one-class collaborative filtering) [Pan *et al.*, 2008] and iMF (implicit matrix factorization) [Hu *et al.*, 2008] are two typical pointwise approaches for solving this recommendation problem. OCCF [Pan *et al.*, 2008] proposes different sampling strategies for unobserved user-item pairs and take them as negative feedback to augment the observed positive feedback, so that existing matrix factorization methods can be applied. iMF [Hu *et al.*, 2008] introduces confidence weights on implicit feedback, which can then be approximated by two latent feature matrices. However, the limitation of OCCF [Pan *et al.*, 2008] is that taking unobserved user-item pairs as negative feedback may introduce errors. As for iMF [Hu *et al.*, 2008], it requires auxiliary knowledge of confidence for each observed feedback, which may not be available in real applications.

Pairwise methods with relative preference assumptions Pairwise methods take implicit feedback as *relative* preferences rather than absolute ones, e.g., a user u is assumed to prefer an item i to an item j if the user-item pair (u, i) is observed, and (u, j) is not observed [Rendle *et al.*, 2009]. The proposed algorithm, Bayesian personalized ranking (BPR) [Rendle *et al.*, 2009], is the first method with such pairwise preference assumption for addressing the one-class collaborative filtering problem.

Due to the great success of pairwise methods in various one-class collaborative filtering problems, some new algorithms have been proposed to combine BPR with some auxiliary data, such as BPR with temporal information [Rendle *et al.*, 2010], BPR with user-side social connections [Du *et al.*, 2011], and BPR with item-side taxonomy [Kanagal *et al.*, 2012], etc. There are also some work that extends BPR from two dimensions to three dimensions [Rendle and Schmidt-Thieme, 2010], or from one user-item matrix to multiple ones [Krohn-Grimberghe *et al.*, 2012].

However, the limitation of pairwise methods can be attributed to the two fundamental assumptions made in BPR, namely *individual pairwise assumption over two items* and *independence assumption between two users*. Most follow-up works do not refine the fundamental assumptions, but just directly adopt the BPR criterion in their own applications. A recent algorithm [Pan and Chen, 2013] generalizes BPR via proposing a new assumption that an individual user is likely to prefer a set of observed items to a set of unobserved items.

Compared with the aforementioned works, our proposed *group Bayesian personalized ranking* (GBPR) method, is a novel algorithm in one-class collaborative filtering. In particular, GBPR inherits the merit of pairwise methods, and improves the two fundamental assumptions in BPR via introducing *group preference*. We summarize GBPR and the aforementioned related works in Table 1.

Table 1: Summary of GBPR and other methods for one-class collaborative filtering w.r.t. different preference assumptions.

Preference assumption	Typical work
Absolute	*OCCF* [Pan *et al.*, 2008], etc.
Relative (individual)	*BPR* [Rendle *et al.*, 2009], etc.
Relative (group)	***GBPR*** (proposed in this paper)

3 Background

In this section, we first give the problem definition, and then introduce the likelihood of pairwise preferences and the two fundamental assumptions made in BPR [Rendle *et al.*, 2009].

3.1 Problem Definition

We use $\mathcal{U}^{tr} = \{u\}_{u=1}^{n}$ and $\mathcal{I}^{tr} = \{i\}_{i=1}^{m}$ to denote the sets of users and items, respectively. For each user $u \in \mathcal{U}^{tr}$, we have a set of items $\mathcal{I}_u^{tr} \subseteq \mathcal{I}^{tr}$ on which user u has expressed positive feedback, e.g., "like". Our goal is then to recommend each user u a personalized ranking list of items from $\mathcal{I}^{tr} \backslash \mathcal{I}_u^{tr}$. As mentioned before, this problem has been steadily receiving more attention, and usually called one-class collaborative filtering [Pan *et al.*, 2008] or collaborative ranking with implicit feedback [Rendle *et al.*, 2009].

3.2 Likelihood of Pairwise Preferences

In order to represent a user u's relative preference on two items i and j, Rendle et al. [Rendle *et al.*, 2009] use a binary random variable $\delta((u,i) \succ (u,j))$ to denote whether user u prefers item i to item j or not. The function $\delta(z)$ is a binary indicator with $\delta(z) = 1$ if the equation z is true, and $\delta(z) = 0$ otherwise. This representation is usually called a user's pairwise preference and has dominated in one-class collaborative filtering tasks in recent studies [Rendle *et al.*, 2009; Du *et al.*, 2011; Kanagal *et al.*, 2012].

For a typical user u, in order to calculate the overall likelihood of pairwise preferences (LPP) among all items, Bernoulli distribution over the binary random variable $\delta((u,i) \succ (u,j))$ is used in [Rendle *et al.*, 2009],

$$\begin{aligned} \text{LPP}(u) &= \prod_{i,j \in \mathcal{I}^{tr}} Pr(\hat{r}_{ui} > \hat{r}_{uj})^{\delta((u,i)\succ(u,j))} \\ &\quad \times [1 - Pr(\hat{r}_{ui} > \hat{r}_{uj})]^{[1-\delta((u,i)\succ(u,j))]} \\ &= \prod_{(u,i)\succ(u,j)} Pr(\hat{r}_{ui} > \hat{r}_{uj})[1 - Pr(\hat{r}_{uj} > \hat{r}_{ui})] \end{aligned}$$

where $(u, i) \succ (u, j)$ denotes that user u prefers item i to item j.

3.3 Bayesian Personalized Ranking

The two fundamental assumptions adopted by the method Bayesian personalized ranking (BPR) [Rendle *et al.*, 2009] are:

1. *Assumption of individual pairwise preference over two items*. It assumes that a user u prefers an item i to an item j, $(u,i) \succ (u,j)$, if the user-item pair (u,i) is observed and (u,j) is not observed. With this assumption, LPP(u) can be simplified to BPR(u) [Rendle *et al.*, 2009],

$$\text{BPR}(u) = \prod_{i \in \mathcal{I}_u^{tr}} \prod_{j \in \mathcal{I}^{tr} \backslash \mathcal{I}_u^{tr}} Pr(\hat{r}_{ui} > \hat{r}_{uj})[1 - Pr(\hat{r}_{uj} > \hat{r}_{ui})],$$

where $i \in \mathcal{I}_u^{tr}$ means that the user-item pair (u,i) is observed, and $j \in \mathcal{I}^{tr} \backslash \mathcal{I}_u^{tr}$ means that the user-item pair (u,j) is not observed.

2. *Assumption of independence among users*. It assumes that the joint likelihood of pairwise preferences of two users, u and w, can be decomposed as BPR(u,w) = BPR(u)BPR(w), which means that the likelihood of pairwise preferences of user u is independent of that of user w. With this assumption, the overall likelihood among the users can be represented as follows [Rendle *et al.*, 2009],

$$\text{BPR} = \prod_{u \in \mathcal{U}^{tr}} \text{BPR}(u).$$

However, the above two assumptions may not always hold in real applications. First, a user u may potentially prefer an item j to an item i, though the user u has expressed positive feedback on item i instead of on item j. For example, John may like *Prince of Egypt* more than *Forrest Gump* though we have only observed positive feedback on the latter yet. Second, two users, u and w, may be correlated, and their joint likelihood can not be decomposed into two independent likelihoods.

4 Our Solution

4.1 Group Bayesian Personalized Ranking

As a response to the possible violations of the two fundamental assumptions made in BPR [Rendle *et al.*, 2009], we propose a new assumption and introduce richer interactions among users via *group preference*. And for this reason, we call our assumption *group Bayesian personalized ranking* (GBPR). In the following, we first describe two definitions before introducing our new assumption.

Definition (**Individual Preference**) The *individual preference* is a preference score of a user on an item. For example, the individual preference of user u on item i is denoted as $\hat{r}_{ui}$.

Definition (**Group Preference**) The *group preference* is an overall preference score of a group of users on an item. For example, the group preference of users from group $\mathcal{G}$ on item i can be estimated from individual preferences, $\hat{r}_{\mathcal{G}i} = \frac{1}{|\mathcal{G}|}\sum_{w \in \mathcal{G}} \hat{r}_{wi}$. Note that our primary goal is to recommend items for a single user, not for a group of users [Amer-Yahia *et al.*, 2009].

We assume that the group preference of $\mathcal{G} \subseteq \mathcal{U}_i^{tr}$ on an item i is more likely to be stronger than the individual preference of user u on item j, if the user-item pair (u,i) is

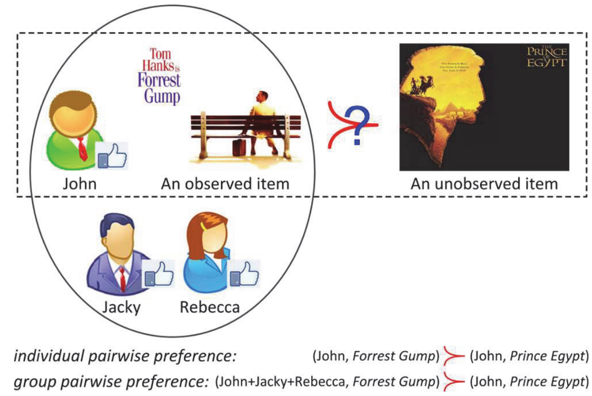

Figure 1: Illustration of *group pairwise preference* over two items. We assume that the *group preference* (shown in oval) of John, Jacky and Rebecca on movie *Forrest Gump* is stronger than the individual preference of John on movie *Prince of Egypt*, since there is positive feedback on movie *Forrest Gump* from all those three guys.

observed and the user-item pair (u,j) is not observed. The *group pairwise preference* can then be written conceptually,

$$(\mathcal{G}, i) \succ (u, j), \quad (1)$$

where $\mathcal{G} \subseteq \mathcal{U}_i^{tr}$ and $u \in \mathcal{G}$. We illustrate our assumption of *group pairwise preference* via a toy example in Figure 1. Our assumption can be interpreted from two aspects,

- For items, it is more likely to be true if user u can find some other users' support on his pairwise preference on item i and item j. This is reflected by the replacement of the individual pairwise relationship $(u,i) \succ (u,j)$ with a new one $(\mathcal{G},i) \succ (u,j)$ that involves the group preference.
- For users, it is natural to introduce interactions and collaborations among users who are all with positive feedback on a specific item, since that implies common interests of those users. This is reflected in the group of like-minded users, $\mathcal{G} \subseteq \mathcal{U}_i^{tr}$, who share the same positive feedback to item i.

To explicitly study the unified effect of group preference and individual preference, we combine them linearly,

$$(\mathcal{G}, i) + (u,i) \succ (u,j) \ \text{ or } \ \hat{r}_{\mathcal{G}ui} > \hat{r}_{uj} \quad (2)$$

where $\hat{r}_{\mathcal{G}ui} = \rho \hat{r}_{\mathcal{G}i} + (1-\rho)\hat{r}_{ui}$ is the fused preference of group preference $\hat{r}_{\mathcal{G}i}$ and individual preference $\hat{r}_{ui}$. Note that $0 \le \rho \le 1$ is a tradeoff parameter used to fuse the two preferences, which can be determined via empirically testing a validation set.

With this assumption, we have a new criterion called *group Bayesian personalized ranking* (GBPR) for user u,

$$\text{GBPR}(u) = \prod_{i \in \mathcal{I}_u^{tr}} \prod_{j \in \mathcal{I}^{tr} \backslash \mathcal{I}_u^{tr}} Pr(\hat{r}_{\mathcal{G}ui} > \hat{r}_{uj})[1 - Pr(\hat{r}_{uj} > \hat{r}_{\mathcal{G}ui})],$$

where $\mathcal{G} \subseteq \mathcal{U}_i^{tr}$ is a user group, and item i is observed by user u and item j is not observed. For any two users, u and w, the joint likelihood can then be approximated via multiplication, GBPR$(u,w) \approx$ GBPR(u)GBPR(w), since the user correlations have been introduced via the user group $\mathcal{G}$ already. More specifically, for any two users u and w who have same positive feedback to an item i, the corresponding user groups $\mathcal{G}(u,i) \subseteq \mathcal{U}_i^{tr}$ and $\mathcal{G}(w,i) \subseteq \mathcal{U}_i^{tr}$ are likely to be overlapped, $\mathcal{G}(u,i) \cap \mathcal{G}(w,i) \neq \emptyset$. Then, we have the following overall likelihood for all users and all items,

$$\text{GBPR} = \prod_{u\in\mathcal{U}^{tr}} \prod_{i\in\mathcal{I}_u^{tr}} \prod_{j\in\mathcal{I}^{tr}\backslash\mathcal{I}_u^{tr}} Pr(\hat{r}_{\mathcal{G}ui} > \hat{r}_{uj})[1 - Pr(\hat{r}_{uj} > \hat{r}_{\mathcal{G}ui})], \quad (3)$$

where $\mathcal{G} \subseteq \mathcal{U}_i^{tr}$.

Following [Rendle *et al.*, 2009], we use $\sigma(\hat{r}_{\mathcal{G}ui} - \hat{r}_{uj}) = \frac{1}{1+\exp(-\hat{r}_{\mathcal{G}ui}+\hat{r}_{uj})}$ to approximate the probability $Pr(\hat{r}_{\mathcal{G}ui} > \hat{r}_{uj})$, and have $Pr(\hat{r}_{\mathcal{G}ui} > \hat{r}_{uj})[1 - Pr(\hat{r}_{uj} > \hat{r}_{\mathcal{G}ui})] = \sigma(\hat{r}_{\mathcal{G}ui} - \hat{r}_{uj})[1-\sigma(\hat{r}_{uj} - \hat{r}_{\mathcal{G}ui})] = \sigma^2(\hat{r}_{\mathcal{G}ui} - \hat{r}_{uj})$. Finally, we reach the objective function of our GBPR,

$$\min_{\Theta} -\frac{1}{2}\ln \text{GBPR} + \frac{1}{2}\mathcal{R}(\Theta) \quad (4)$$

where $\Theta = \{U_{u\cdot} \in \mathbb{R}^{1\times d}, V_{i\cdot} \in \mathbb{R}^{1\times d}, b_i \in \mathbb{R}, u \in \mathcal{U}^{tr}, i \in \mathcal{I}^{tr}\}$ is a set of model parameters to be learned, $\ln \text{GBPR} = \sum_{u\in\mathcal{U}^{tr}} \sum_{i\in\mathcal{I}_u^{tr}} \sum_{j\in\mathcal{I}^{tr}\backslash\mathcal{I}_u^{tr}} 2\ln\sigma(\hat{r}_{\mathcal{G}ui} - \hat{r}_{uj})$ is the log-likelihood of GBPR, $\mathcal{R}(\Theta) = \sum_{u\in\mathcal{U}^{tr}} \sum_{i\in\mathcal{I}_u^{tr}} \sum_{j\in\mathcal{I}^{tr}\backslash\mathcal{I}_u^{tr}} [\alpha_w \sum_{w\in\mathcal{G}} \|U_{w\cdot}\|^2 + \alpha_v\|V_{i\cdot}\|^2 + \alpha_v\|V_{j\cdot}\|^2 + \beta_v\|b_i\|^2 + \beta_v\|b_j\|^2]$ is the regularization term used to avoid overfitting, and $\mathcal{G} \subseteq \mathcal{U}_i^{tr}$ is a group of like-minded users who share the same positive feedback to item i. We show the graphical model of GBPR in Figure 2, where the individual preference is generated via $\hat{r}_{ui} = U_{u\cdot}V_{i\cdot}^T + b_i$, $\hat{r}_{uj} = U_{u\cdot}V_{j\cdot}^T + b_j$, and group preference $\hat{r}_{\mathcal{G}i} = \bar{U}_{\mathcal{G}\cdot}V_{i\cdot}^T + b_i$ with $\bar{U}_{\mathcal{G}\cdot} = \sum_{w\in\mathcal{G}} U_{w\cdot}/|\mathcal{G}|$. The main difference between GBPR and BPR [Rendle *et al.*, 2009] is the first term in Eq.(4), which introduces richer interactions among users via the user group $\mathcal{G}$, and as a consequence relax the two fundamental assumptions made in BPR [Rendle *et al.*, 2009].

Therefore, once we have learned the model parameters Θ, we can predict the preference of user u on item j via $\hat{r}_{uj} = U_{u\cdot}V_{j\cdot}^T + b_j$, which can then be used to generate a personalized ranking list for user u via picking up the top-k items with largest preference scores.

4.2 Learning the GBPR

We follow the widely used stochastic gradient descent (SGD) algorithm to optimize the objective function in Eq.(4). We go one step beyond SGD, and randomly sample a subset of like-minded users to construct the user group $\mathcal{G}$. For each randomly sampled record, it includes a user u, an item i, an item j, and a user group $\mathcal{G}$, where $u \in \mathcal{G}$. The tentative objective function can be written as $f(\mathcal{G},u,i,j) = -\ln\sigma(\hat{r}_{\mathcal{G}ui} - \hat{r}_{uj}) + \frac{\alpha_u}{2}\sum_{w\in\mathcal{G}}\|U_{w\cdot}\|^2 + \frac{\alpha_v}{2}\|V_{i\cdot}\|^2 + \frac{\alpha_v}{2}\|V_{j\cdot}\|^2 + \frac{\beta_v}{2}\|b_i\|^2 + \frac{\beta_v}{2}\|b_j\|^2 = \ln[1+\exp(-\hat{r}_{\mathcal{G}ui;uj})] + \frac{\alpha_u}{2}\sum_{w\in\mathcal{G}}\|U_{w\cdot}\|^2 + \frac{\alpha_v}{2}\|V_{i\cdot}\|^2 + \frac{\alpha_v}{2}\|V_{j\cdot}\|^2 + \frac{\beta_v}{2}\|b_i\|^2 + \frac{\beta_v}{2}\|b_j\|^2$, where $\hat{r}_{\mathcal{G}ui;uj} = \hat{r}_{\mathcal{G}ui} - \hat{r}_{uj}$ is the difference between the fused preference $\hat{r}_{\mathcal{G}ui}$ and individual preference $\hat{r}_{uj}$.

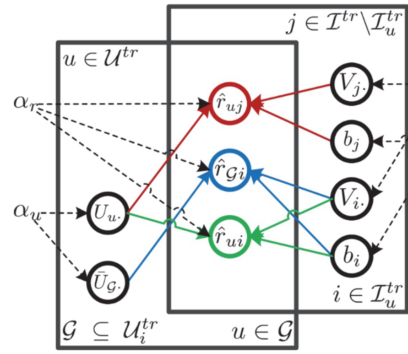

Figure 2: Graphical model of *group Bayesian personalized ranking* (GBPR).

We then have the gradients of the user-specific parameters w.r.t. the tentative objective function $f(\mathcal{G},u,i,j)$,

$$\frac{\partial f(\mathcal{G},u,i,j)}{\partial U_{w\cdot}} = \frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}} \times ((1-\rho)\delta(w=u)V_{i\cdot} + \rho\frac{V_{i\cdot}}{|\mathcal{G}|} - \delta(w=u)V_{j\cdot}) + \alpha_u U_{w\cdot},$$

and the gradients of the item-specific parameters,

$$\frac{\partial f(\mathcal{G},u,i,j)}{\partial V_{i\cdot}} = \frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}}[(1-\rho)U_{u\cdot} + \rho\bar{U}_{\mathcal{G}\cdot}] + \alpha_v V_{i\cdot},$$

$$\frac{\partial f(\mathcal{G},u,i,j)}{\partial V_{j\cdot}} = \frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}}(-U_{u\cdot}) + \alpha_v V_{j\cdot},$$

$$\frac{\partial f(\mathcal{G},u,i,j)}{\partial b_i} = \frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}} + \beta_v b_i,$$

$$\frac{\partial f(\mathcal{G},u,i,j)}{\partial b_j} = \frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}}(-1) + \beta_v b_j,$$

where $\frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}} = -\frac{\exp(-\hat{r}_{\mathcal{G}ui;uj})}{1+\exp(-\hat{r}_{\mathcal{G}ui;uj})} = -\frac{1}{1+\exp(\hat{r}_{\mathcal{G}ui;uj})}$, and $w \in \mathcal{G}$. With the above gradients, we can update the model parameters as follows,

$$\theta = \theta - \gamma\frac{\partial f(\mathcal{G},u,i,j)}{\partial \theta}, \quad (5)$$

where θ can be $U_{w\cdot}, w \in \mathcal{G}$, $V_{i\cdot}$, $V_{j\cdot}$, b_i or b_j, and $\gamma > 0$ is the learning rate. Note that when $\mathcal{G} = \{u\}$ or $\rho = 0$, GBPR is reduced to BPR [Rendle *et al.*, 2009], which does not explicitly incorporate interactions among like-minded users.

The complete steps to learn the model parameters are depicted in Figure 3. The time complexity of the update rule in Eq.(5) is $O(|\mathcal{G}|d)$, where $|\mathcal{G}|$ is the size of user group, and d is the number of latent features. The total time complexity is then $O(Tn|\mathcal{G}|d)$, where T is the number of iterations and n is the number of users. We can see that introducing interactions in GBPR does not increase the time complexity much, because $|\mathcal{G}|$ is usually small, e.g., $|\mathcal{G}| \leq 5$ in our experiments. For predicting a user's preference on an item, the time complexity is $O(d)$, which is the same as that of BPR. Thus,

Input: Training data $T_R = \{(u, i)\}$ of observed feedback and the size of user group $|\mathcal{G}|$.
Output: The learned model parameters $\Theta = \{U_{u\cdot}, V_{i\cdot}, b_{i\cdot}, u \in \mathcal{U}^{tr}, i \in \mathcal{I}^{tr}\}$, where $U_{u\cdot}$ is the user-specific latent feature vector of user u, $V_{i\cdot}$ is the item-specific latent feature vector of item i, and b_i is the bias of item i.
For $t_1 = 1, \ldots, T$.
  For $t_2 = 1, \ldots, n$.
  Step 1. Randomly pick a user $u \in \mathcal{U}^{tr}$.
  Step 2. Randomly pick an item $i \in \mathcal{I}_u^{tr}$.
  Step 3. Randomly pick an item $j \in \mathcal{I}^{tr} \backslash \mathcal{I}_u^{tr}$.
  Step 4. Randomly pick $|\mathcal{G}| - 1$ users from $\mathcal{U}_i^{tr} \backslash \{u\}$ to form the user group $\mathcal{G}$.
  Step 5. Calculate $\frac{\partial f(\mathcal{G},u,i,j)}{\partial \hat{r}_{\mathcal{G}ui;uj}}$ and $\bar{U}_{\mathcal{G}\cdot}$.
  Step 6. Update $U_{w\cdot}, w \in \mathcal{G}$ via Eq.(5).
  Step 7. Update $V_{i\cdot}$ via Eq.(5).
  Step 8. Update $V_{j\cdot}$ via Eq.(5).
  Step 9. Update b_i via Eq.(5).
  Step 10. Update b_j via Eq.(5).
  End
End

Figure 3: The algorithm of *group Bayesian personalized ranking* (GBPR).

GBPR can be comparable to the seminal work BPR [Rendle *et al.*, 2009] in terms of efficiency. In the experiments, we mainly assess whether GBPR would be more accurate than BPR.

5 Experimental Results

5.1 Datasets

We use four real-world datasets in our empirical studies, including MovieLens100K[2], MovieLens1M, UserTag [Pan *et al.*, 2008] and a subset of Netflix. MovieLens100K contains $100,000$ ratings assigned by 943 users on $1,682$ movies, MovieLens1M contains $1,000,209$ ratings assigned by $6,040$ users on $3,952$ movies, and UserTag contains $246,436$ user-tag pairs from $3,000$ users and $2,000$ tags. We randomly sample $5,000$ users from the user pool and $5,000$ items from the item pool of the Netflix dataset, and obtain $282,474$ ratings by those $5,000$ users on those $5,000$ items. We call this subset of Netflix dataset Netflix5K5K. We use "item" to denote movie (for MovieLens100K, MovieLens1M and Netflix5K5K) or tag (for UserTag). For MovieLens100K, MovieLens10M and Netflix5K5K, we take a pre-processing step [Sindhwani *et al.*, 2009], which only keeps the ratings larger than 3 as the observed positive feedback (to simulate the one-class feedback).

For all four datasets, we randomly sample half of the observed user-item pairs as training data, and the rest as test data; we then randomly take 1 user-item pair for each user from the training data to construct a validation set. We repeat the above procedure for three times, so we have three copies of training data and test data. The final datasets used in the experiments are shown in Table 2. The experimental results are averaged over the performance on those three copies of test data.

[2]http://www.grouplens.org/node/73

Table 2: Description of the datasets used in the experiments.

Data set		user-item pairs
MovieLens100K	training	$27,688$
	test	$27,687$
MovieLens1M	training	$287,641$
	test	$287,640$
UserTag	training	$123,218$
	test	$123,218$
Netflix5K5K	training	$77,936$
	test	$77,936$

5.2 Evaluation Metrics

Because users usually only check a few top-ranked items [Chen and Pu, 2011], we use top-k evaluation metrics to study the recommendation performance, including top-k results of precision, recall, F1, NDCG and 1-call [Chen and Karger, 2006]. Since BPR [Rendle *et al.*, 2009] optimizes the AUC criterion, we also include it in our evaluation. For each evaluation metric, we first calculate the performance for each user from the test data, and then obtain the averaged performance over all users.

5.3 Baselines and Parameter Settings

We use two popular baseline algorithms in our experiments, PopRank and BPR [Rendle *et al.*, 2009]. PopRank is a basic algorithm in one-class collaborative filtering, which ranks the items according to their popularity in the training data. BPR [Rendle *et al.*, 2009] is a seminal work for this problem and is also a very strong baseline, which is shown to be much better than two well-known pointwise methods [Du *et al.*, 2011], i.e., iMF [Hu *et al.*, 2008] and OCCF [Pan *et al.*, 2008]. In this paper, we extend BPR via introducing richer interactions, and thus, we concentrate our empirical study on comparisons between BPR and our GBPR, which are both implemented in the same code framework as shown in Figure 3.

For all experiments, the tradeoff parameters are searched as $\alpha_u = \alpha_v = \beta_v \in \{0.001, 0.01, 0.1\}$ and $\rho \in \{0.2, 0.4, 0.6, 0.8, 1\}$, and the iteration number is chosen from $T \in \{1000, 10000, 100000\}$. The $NDCG@5$ performance on the validation data is used to select the best parameters α_u, α_v, β_v and ρ, and the best iteration number T for both BPR and GBPR. The learning rate in BPR and GBPR is fixed as $\gamma = 0.01$. The initial value of $U_{u\cdot}, V_{i\cdot}, b_i$ in BPR and GBPR are set the same as in [Pan *et al.*, 2012]. For the user group $\mathcal{G}$ in GBPR, we first fix the size as $|\mathcal{G}| = 3$, and then change it as $|\mathcal{G}| \in \{1, 2, 3, 4, 5\}$ in order to study the effect of different levels of interactions as introduced in our proposed GBPR algorithm. Note that when $\rho = 0$ or $|\mathcal{G}| = 1$, GBPR reduces to BPR.

Table 3: Recommendation performance of PopRank, BPR and GBPR on MovieLens100K, MovieLens1M, UserTag and Netflix5K5K. Note that the parameter $\rho \in \{0.2, 0.4, 0.6, 0.8, 1\}$ in GBPR is determined via the performance on the validation set, which is shown in parentheses. The size of the user group is fixed as $|\mathcal{G}| = 3$, and the latent dimension is fixed as $d = 20$. Numbers in boldface (e.g., **0.4051**) are the best results among all methods.

Dataset	Method	$Prec@5$	$Rec@5$	$F1@5$	$NDCG@5$	1-$call@5$	AUC
MovieLens100K	PopRank	0.2724±0.0094	0.0549±0.0028	0.0821±0.0036	0.2915±0.0072	0.6520±0.0201	0.8526±0.0006
	BPR	0.3709±0.0066	0.0950±0.0014	0.1308±0.0026	0.3885±0.0107	0.8156±0.0015	0.9033±0.0007
	GBPR ($\rho = 1$)	**0.4051**±0.0038	**0.1046**±0.0016	**0.1445**±0.0015	**0.4201**±0.0031	**0.8414**±0.0058	**0.9140**±0.0008
MovieLens1M	PopRank	0.2822±0.0019	0.0407±0.0004	0.0634±0.0003	0.2935±0.0010	0.6676±0.0006	0.8771±0.0002
	BPR	0.4410±0.0008	0.0744±0.0003	0.1135±0.0003	0.4540±0.0009	0.8496±0.0047	0.9339±0.0004
	GBPR ($\rho = 0.6$)	**0.4494**±0.0020	**0.0781**±0.0009	**0.1188**±0.0010	**0.4636**±0.0014	**0.8670**±0.0022	**0.9354**±0.0005
UserTag	PopRank	0.2647±0.0012	0.0405±0.0003	0.0640±0.0004	0.2730±0.0014	0.5221±0.0062	0.6810±0.0016
	BPR	0.2969±0.0025	0.0476±0.0008	0.0740±0.0006	0.3072±0.0017	0.6172±0.0008	0.7711±0.0013
	GBPR ($\rho = 0.8$)	**0.3011**±0.0008	**0.0491**±0.0014	**0.0766**±0.0012	**0.3104**±0.0009	**0.6226**±0.0019	**0.7892**±0.0011
Netflix5K5K	PopRank	0.1728±0.0012	0.0563±0.0013	0.0683±0.0001	0.1794±0.0004	0.4472±0.0059	0.9147±0.0015
	BPR	0.2318±0.0006	0.0945±0.0012	0.1046±0.0002	0.2508±0.0006	0.5683±0.0016	0.9268±0.0019
	GBPR ($\rho = 0.8$)	**0.2411**±0.0027	**0.0979**±0.0013	**0.1095**±0.0013	**0.2611**±0.0025	**0.5844**±0.0015	**0.9321**±0.0014

5.4 Summary of Experimental Results

The recommendation performance of GBPR and other baselines are shown in Table 3, from which we can have the following observations,

1. both BPR and GBPR are much better than the PopRank algorithm, which demonstrates the effectiveness of *pairwise preference* assumptions, and
2. GBPR further improves BPR on all evaluation metrics on all four datasets, which shows the effect of the injected interactions among users via *group preference*.

We can thus see that the assumption that combines *pairwise preference* and *group preference* in GBPR is indeed more effective than that of simple *pairwise preference* assumed in BPR [Rendle *et al.*, 2009].

To have a deep understanding of the effect of *group preference* in GBPR, we adjust the group size $|\mathcal{G}| \in \{1, 2, 3, 4, 5\}$ and show the results of $Pre@5$, $NDCG@5$ and AUC in Figure 4. The performance on $Rec@5$, $F1@5$ and 1-$call@5$ are similar, so they are not included for the sake of saving space. From Figure 4, we can see that using a relatively larger user group (e.g., $|\mathcal{G}| = 3$ or 4) improves the recommendation performance on all four datasets. This can be explained by the effect of introducing the user group $\mathcal{G}$ for modeling the pairwise preference in Eq.(2) and learning the model parameters in Eq.(5).

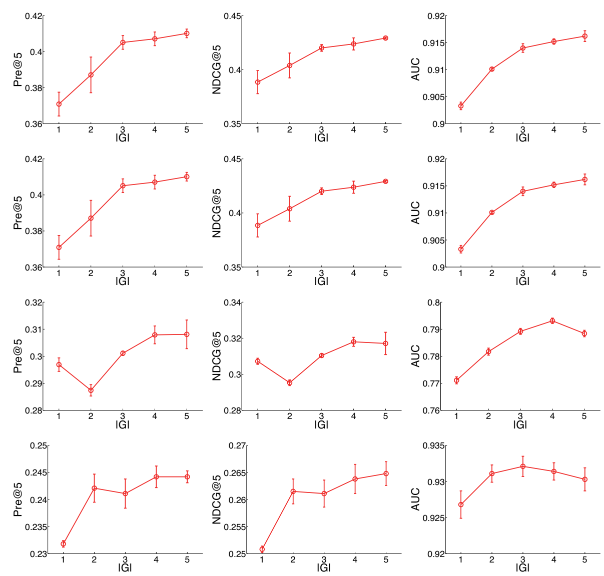

Figure 4: Recommendation performance of GBPR with different sizes of user group (from top row to bottom row: MovieLens100K, MovieLens1M, UserTag, and Netflix5K5K).

6 Conclusions and Future Work

In this paper, we study the one-class collaborative filtering problem and design a novel algorithm called *group Bayesian personalized ranking* (GBPR). GBPR introduces richer interactions among users in order to improve the *individual* and *independence* assumptions as made in Bayesian personalized ranking (BPR) [Rendle *et al.*, 2009], a seminal work for addressing this problem. GBPR is comparable to BPR in terms of time complexity. Experimental results on four real-world datasets show that GBPR can recommend items more accurately than BPR regarding various evaluation metrics.

For future works, we are interested in extending GBPR via (1) optimizing the user group construction process, such as incorporating time, location, taxonomy and other possible contexture information to refine the like-minded user groups, and (2) adaptively changing group size for different segmentations of users and items during the learning process.

Acknowledgement

We thank the support of Hong Kong RGC under the project ECS/HKBU211912.

References

[Amer-Yahia *et al.*, 2009] Sihem Amer-Yahia, Senjuti Basu Roy, Ashish Chawlat, Gautam Das, and Cong Yu. Group recommendation: semantics and efficiency. *Proc. VLDB Endow.*, 2(1):754–765, August 2009.

[Chen and Karger, 2006] Harr Chen and David R. Karger. Less is more: probabilistic models for retrieving fewer relevant documents. In *Proceedings of the 29th annual international ACM SIGIR conference on research and development in information retrieval*, SIGIR '06, pages 429–436, New York, NY, USA, 2006. ACM.

[Chen and Pu, 2011] Li Chen and Pearl Pu. Users' eye gaze pattern in organization-based recommender interfaces. In *Proceedings of the 16th international conference on intelligent user interfaces*, IUI '11, pages 311–314, New York, NY, USA, 2011. ACM.

[Du *et al.*, 2011] Liang Du, Xuan Li, and Yi-Dong Shen. User graph regularized pairwise matrix factorization for item recommendation. In *Proceedings of the 7th international conference on advanced data mining and applications*, ADMA '11, pages 372–385, Berlin, Heidelberg, 2011. Springer-Verlag.

[Hu *et al.*, 2008] Yifan Hu, Yehuda Koren, and Chris Volinsky. Collaborative filtering for implicit feedback datasets. In *Proceedings of the 2008 8th IEEE international conference on data mining*, ICDM '08, pages 263–272, Washington, DC, USA, 2008.

[Kanagal *et al.*, 2012] Bhargav Kanagal, Amr Ahmed, Sandeep Pandey, Vanja Josifovski, Jeff Yuan, and Lluis Garcia-Pueyo. Supercharging recommender systems using taxonomies for learning user purchase behavior. *Proceedings of the VLDB Endowment*, 5(10):956–967, June 2012.

[Koren, 2008] Yehuda Koren. 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 '08, pages 426–434, New York, NY, USA, 2008. ACM.

[Krohn-Grimberghe *et al.*, 2012] Artus Krohn-Grimberghe, Lucas Drumond, Christoph Freudenthaler, and Lars Schmidt-Thieme. Multi-relational matrix factorization using bayesian personalized ranking for social network data. In *Proceedings of the 5th ACM international conference on web search and data mining*, WSDM '12, pages 173–182, New York, NY, USA, 2012. ACM.

[Liu *et al.*, 2013] Nathan N. Liu, Luheng He, and Min Zhao. Social temporal collaborative ranking for context aware movie recommendation. *ACM Transactions on Intelligent Systems and Technology (ACM TIST)*, 4(1):15:1–15:26, February 2013.

[Nichols, 1997] David M. Nichols. Implicit rating and filtering. In *Proceedings of the 5th DELOS workshop on filtering and collaborative filtering*, November 1997.

[Pan and Chen, 2013] Weike Pan and Li Chen. Cofiset: Collaborative filtering via learning pairwise preferences over item-sets. In *Proceedings of SIAM International Conference on Data Mining*, SDM '13, Austin, Texas, USA, 2013.

[Pan *et al.*, 2008] Rong Pan, Yunhong Zhou, Bin Cao, Nathan Nan Liu, Rajan M. Lukose, Martin Scholz, and Qiang Yang. One-class collaborative filtering. In *Proceedings of the 8th IEEE international conference on data mining*, ICDM '08, pages 502–511, Washington, DC, USA, 2008. IEEE Computer Society.

[Pan *et al.*, 2012] Weike Pan, Evan Wei Xiang, and Qiang Yang. Transfer learning in collaborative filtering with uncertain ratings. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, AAAI '12, pages 662–668, July 2012.

[Rendle and Schmidt-Thieme, 2010] Steffen Rendle and Lars Schmidt-Thieme. Pairwise interaction tensor factorization for personalized tag recommendation. In *Proceedings of the 3rd ACM international conference on web search and data mining*, WSDM '10, pages 81–90, New York, NY, USA, 2010. ACM.

[Rendle *et al.*, 2009] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th conference on uncertainty in artificial intelligence*, UAI '09, pages 452–461, Arlington, Virginia, United States, 2009. AUAI Press.

[Rendle *et al.*, 2010] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. Factorizing personalized markov chains for next-basket recommendation. In *Proceedings of the 19th international conference on world wide web*, WWW '10, pages 811–820, New York, NY, USA, 2010. ACM.

[Rendle, 2012] Steffen Rendle. Factorization machines with libfm. *ACM Transactions on Intelligent Systems and Technology (ACM TIST)*, 3(3):57:1–57:22, May 2012.

[Sindhwani *et al.*, 2009] Vikas Sindhwani, S.S. Bucak, J. Hu, and A. Mojsilovic. A family of non-negative matrix factorizations for one-class collaborative filtering. In *The 1st international workshop on recommendation-based industrial applications held in the 3rd ACM conference on recommender systems*, RecSys: RIA '09, 2009.

[Weimer *et al.*, 2007] Markus Weimer, Alexandros Karatzoglou, Quoc V. Le, and Alex J. Smola. Cofi rank - maximum margin matrix factorization for collaborative ranking. In *Annual conference on neural information processing systems 19*, NIPS '07. MIT Press, 2007.

[Yang *et al.*, 2011] Shuang-Hong Yang, Bo Long, Alexander J. Smola, Hongyuan Zha, and Zhaohui Zheng. Collaborative competitive filtering: learning recommender using context of user choice. In *Proceedings of the 34th international ACM SIGIR conference on research and development in information retrieval*, SIGIR '11, pages 295–304, New York, NY, USA, 2011. ACM.