

SetRank: A Setwise Bayesian Approach for Collaborative Ranking from Implicit Feedback

Chao Wang,^{1,2} Hengshu Zhu,^{2*} Chen Zhu,² Chuan Qin,^{1,2} Hui Xiong^{1,2,3*}

¹University of Science and Technology of China

²Baidu Talent Intelligence Center, Baidu Inc.

³Business Intelligence Lab, Baidu Research

wdyx2012@mail.ustc.edu.cn, zhuhengshu@baidu.com, {zc3930155, chuanqin0426, xionghui}@gmail.com

Abstract

The recent development of online recommender systems has a focus on collaborative ranking from implicit feedback, such as user clicks and purchases. Different from explicit ratings, which reflect graded user preferences, the implicit feedback only generates positive and unobserved labels. While considerable efforts have been made in this direction, the well-known pairwise and listwise approaches have still been limited by various challenges. Specifically, for the pairwise approaches, the assumption of independent pairwise preference is not always held in practice. Also, the listwise approaches cannot efficiently accommodate “ties” due to the precondition of the entire list permutation. To this end, in this paper, we propose a novel setwise Bayesian approach for collaborative ranking, namely SetRank, to inherently accommodate the characteristics of implicit feedback in recommender system. Specifically, SetRank aims at maximizing the posterior probability of novel setwise preference comparisons and can be implemented with matrix factorization and neural networks. Meanwhile, we also present the theoretical analysis of SetRank to show that the bound of excess risk can be proportional to $\sqrt{M/N}$, where M and N are the numbers of items and users, respectively. Finally, extensive experiments on four real-world datasets clearly validate the superiority of SetRank compared with various state-of-the-art baselines.

Introduction

Recommender systems have been widely deployed in many popular online services for enhancing user experience and business revenue (Wang, Wang, and Yeung 2015; Liu et al. 2019a; 2019b; Zhu et al. 2018). As one representative task of personalized recommendation, collaborative ranking aims at providing a user-specific item ranking for users based on their preferences learned from historical feedback. Indeed, in real-world scenarios, most of the user feedback is implicit (e.g., clicks and purchases) but not explicit (e.g., 5-star ratings). Different from explicit ratings, the implicit feedback only contains positive and unobserved labels instead of graded user preferences, which brings new research challenges for building recommender systems (Hsieh, Natarajan,

and Dhillon 2015). Therefore, collaborative ranking from implicit feedback has been attracting more and more attention in recent years (Rendle et al. 2009; Shi, Larson, and Hanjalic 2010; Huang et al. 2015; Xia 2019).

While considerable efforts have been made in this direction (Rendle et al. 2009; Chen et al. 2009; Weimer et al. 2008), represented by the well-known *pairwise* and *listwise* approaches, some critical challenges still exist. As for the family of pairwise approaches (Rendle et al. 2009; Freund et al. 2003; Chapelle and Keerthi 2010; Wu, Hsieh, and Sharpnack 2017; Krohn-Grimberghe et al. 2012), which take the item pair as the basic element to model the preference structure in implicit feedback, they are prone to the inconsistency problem between assumption and practice. For example, Bayesian Personalized Ranking (BPR) (Rendle et al. 2009), one of the most widely used collaborative pairwise approaches, tries to maximize the probability of binary comparison between positive and unobserved feedback. Such treatment requires the strict assumption of independent pairwise preference over two items as the basis for constructing pairwise loss. However, as shown in Figure 1, if there exist the item preference pairs “A>B” and “C>D” for user 1, the pairs “A>D” and “C>B” must also exist for user 1 due to the binary value of implicit feedback. In other words, we have $p(A>D, C>B | A>B, C>D) = 1$ in the practical pair construction process, which breaks the independence among pairs and thus influences the optimization result of the pairwise loss. Some follow-up studies chose to relax the independence assumption by considering group information. For example, GBPR (Pan and Chen 2013b) introduced richer users’ interactions and Cofiset (Pan and Chen 2013a) defined a user’s preference on the item group to consider the composition effect. However, the inconsistency problem remains to some extent.

As for the listwise approaches, the key challenge is how to efficiently accommodate “ties” (items with the same rating value) due to the precondition of entire list permutation, since there is no clear sequential relationship but binary rating in implicit feedback. Besides, they measure the uncertainty between the top- P items on the observed and predicted list by calculating the cross-entropy (Cao et al. 2007; Shi, Larson, and Hanjalic 2010; Huang et al. 2015; Wang

*Hui Xiong and Hengshu Zhu are corresponding authors.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

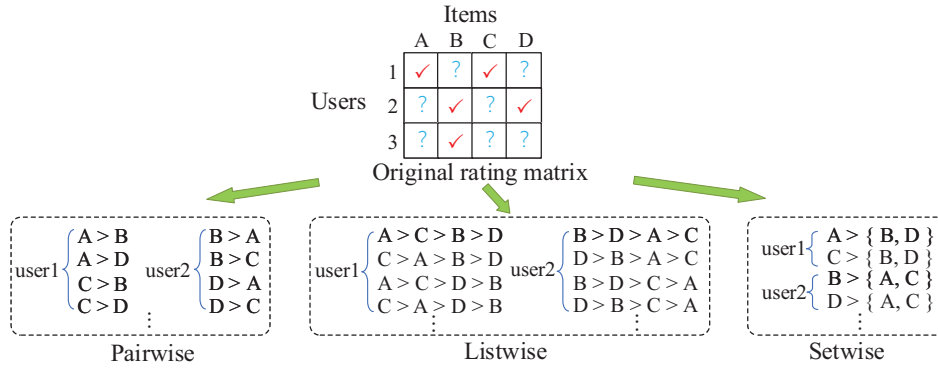


Figure 1: The diagrammatic sketch of preference structures in different collaborative ranking approaches, where the notation “>” represents the preference order.

et al. 2016), which would result in the exponential computational complexity to P (that is why often P is set as 1). Though Wu, Hsieh, and Sharpnack(2018) tried to propose a permutation probability based listwise model to solve the above challenges, only the upper bound rather than the original negative log-likelihood is optimized.

To avoid the limitations of the existing collaborative ranking approaches, in this paper, we propose a novel setwise Bayesian approach, namely SetRank, for collaborative ranking. SetRank has the ability in accommodating the characteristics of implicit feedback in recommender systems. Particularly, we first make a weaker independence assumption compared to pairwise approaches, that is, each user prefers every positive item over the set of unobserved items independently. Hence, we can transform the original rating records into the comparisons between each single positive item and the set of unobserved items, which could avoid the inconsistency problem in pairwise approaches, as the example shown in Figure 1. Moreover, since there is no ordering information between unobserved items, it is unnecessary to rank the set of unobserved items, which relaxes the permutation form in listwise approaches. Specifically, our approach is named as “setwise” because the preference order of a user is only defined between each positive item and the set of unobserved items. Consequently, SetRank is able to model the properties of implicit feedback in a more effective manner, with avoiding the disadvantages of both pairwise and listwise ranking approaches. The contributions of this work can be summarized as follows:

- We propose a novel setwise Bayesian collaborative ranking approach, namely SetRank, to provide a new research perspective for implicit feedback based recommendations. SetRank can inherently accommodate the characteristics of implicit feedback.
- We design two implementations for SetRank, namely MF-SetRank and Deep-SetRank based on matrix factorization and neural networks, respectively.
- We validate our approach by both theoretical analysis and experiments. Specifically, we prove that the bound of excess risk can be bounded by a $\sqrt{M/N}$ term, where M and

N are the numbers of items and users, respectively. Meanwhile, extensive experiments on four real-world datasets clearly demonstrate the advantages of our approach compared with various state-of-the-art baselines.

Setwise Bayesian Collaborative Ranking

Problem Formulation

Suppose there are N users and M items in the dataset. Let P_i and O_i denote the set of positive and unobserved items for each user i , respectively. User i has $J_i = |P_i|$ positive items and $K_i = |O_i|$ unobserved items. Then the rating matrix $R = \{R_{il}\}_{N \times M}$ is a binary matrix, i.e., $R_{ij} = 1$ for $j \in P_i$ and $R_{ik} = 0$ for $k \in O_i$. The goal of collaborative ranking is to recommend each user an ordered item list by predicting the preference score matrix $X = \{X_{il}\}_{N \times M}$.

SetRank Optimization Criterion

The target of SetRank is to maximize the posterior probability of preference structure to build the Bayesian formulation of collaborative ranking:

$$p(\Theta | >_{total}) \propto p(>_{total} | \Theta)p(\Theta), \quad (1)$$

where $>_{total} = \{>_i\}_{i=1}^N$ and $>_i$ is a random variable representing the preference structure of user i , which takes values from all possible preference structures. Θ is the model parameters to be learned.

Before modeling the setwise preference structure probability, we first give a new independence assumption:

Assumption 1 Every user i prefers the positive item $j \in P_i$ to unobserved item set O_i independently.

In this setwise assumption, we ignore the direct comparisons among positive or unobserved items to better reflect the nature of implicit feedback, since there is no explicit item-level preference information. Supposing there is only one user, we have no reason to decide which positive item is better than another positive one, or which unobserved item is better than another unobserved one. Only when there are many users, we can then exploit collaborative information to derive entire ranking results.

By comparison, pairwise approaches like BPR (Rendle et al. 2009) establish the individual binary comparison over each positive item and each unobserved item, which need the strict assumption that item comparisons are independent for optimization. However, in the pair construction process, the pairs are bound to be dependent due to the characteristic of implicit feedback. When calculating the pairwise loss, pairwise approaches still assume pairs are independent and optimize the improper loss. By contrast, setwise approach has no such inconsistency problem owing to the weaker independence assumption.

Moreover, setwise permutation form is weaker than listwise approaches. In the comparisons, we do not care about the ranking of items in O_i , since the ordering information of unobserved items is naturally missing in implicit data. As a result, all the unobserved items are treated equally in the preference comparison, and thus the setwise ranking approach is inherently suitable for handling implicit data.

According to our assumption, we can transform Equation 1 into the following form:

$$p(>_{total} | \Theta) = \prod_{i=1}^N p(>_i | \Theta) = \prod_{i=1}^N \prod_{j \in P_i} p(j >_i O_i | \Theta), \quad (2)$$

where $j >_i O_i$ denotes the user i prefers item j to item set O_i . Therefore, we turn to collect the preference comparison between a single positive item and an unobserved item set. For example, in Figure 1, there are two comparisons for user 1, $A > \{B, D\}$ and $C > \{B, D\}$.

In the setwise preference structure, the positive item j and the unobserved set O_i compose a new item list L_{ij} . Hence, it is convenient to draw the concept of permutation probability (Cao et al. 2007) from listwise approaches for further specifying the preference structure probability $p(j >_i O_i)$. Review that in listwise approach, a permutation $\pi = \{\pi_1, \pi_2, \dots, \pi_m\}$ is a list in descending order of the m items (Cao et al. 2007). Denote the scores assigned to items as a vector $s = (s_1, s_2, \dots, s_m)$ and $\phi(x)$ is an increasing and strictly positive function. Then the permutation probability is defined as:

$$p_s(\pi) := \prod_{d=1}^m \frac{\phi(s_{\pi_d})}{\sum_{l=d}^m \phi(s_{\pi_l})}. \quad (3)$$

It is easy to verify that $p_s(\pi)$ is a valid probability distribution. In the literature, permutation probability has been widely used in many listwise approaches to calculate the cross entropy due to many beneficial properties (Xia et al. 2008). These properties guarantee that items with higher scores are more likely to be ranked higher. However, a serious problem of the definition is that we have to calculate $P!$ permutation probabilities to obtain the top- P probability of the list. Fortunately, in our case, we only need to place the positive item j at the top of List L_{ij} , which means we just concentrate on the top-1 probability. Actually, Cao et al.(2007) had proved that the top-1 probability $p_{s,1}(d)$ of item d can be efficiently calculated under the definition of Equation 3 as follows:

$$p_{s,1}(d) = \frac{\phi(s_d)}{\sum_{l=1}^m \phi(s_l)}. \quad (4)$$

With the help of Equation 4 and the preference score matrix X , now we can give the detailed formulation of the setwise preference probability over all users:

$$p(>_{total} | \Theta) = \prod_{i=1}^N \prod_{j \in P_i} \frac{\phi(X_{ij})}{\phi(X_{ij}) + \sum_{k \in O_i} \phi(X_{ik})}. \quad (5)$$

As one can see, Equation 5 indicates that if positive items have higher scores and unobserved items have lower scores, this preference structure will be more likely to be true.

At last, to complete the Bayesian inference, we introduce a general prior probability $p(\Theta)$. Following BPR (Rendle et al. 2009), $p(\Theta)$ is set as a normal distribution with zero mean and variance-covariance matrix $\lambda_{\Theta} I$. Hence, maximizing the posterior probability is equivalent to minimizing the following function:

$$L = \sum_{i=1}^N \sum_{j \in P_i} -\log p(j >_i O_i | \Theta) + \lambda_{\Theta} \|\Theta\|^2. \quad (6)$$

Note that though some listwise approaches also exploit Equation 4 (Cao et al. 2007; Shi, Larson, and Hanjalic 2010), it is actually quite different from SetRank. First, listwise approaches are essentially based on the top- P probability since they consider the order in a list composed of multiple positive and unobserved items. In fact, using a larger P tends to improve the performance of listwise approaches (Cao et al. 2007). They use the top-1 probability mainly due to the compromise of exponential computational complexity for calculating the top- P probability. However, our setwise assumption, which is more appropriate for implicit feedback, is naturally based on the top-1 probability. Second, they could only employ cross-entropy loss for calculation while the cross-entropy loss may rank worse scoring permutations higher (Wu, Hsieh, and Sharpnack 2018). On the contrary, our loss is strictly obtained by Bayesian inference without the adoption of cross-entropy.

Implementation

It is quite flexible to apply many well-known models to learn the score matrix X . In the literature, matrix factorization (Mnih and Salakhutdinov 2008) and neural network (NN) (Xue et al. 2017) have demonstrated their effectiveness and practicability for recommender systems. Therefore, in this paper, we introduce two implementations for SetRank, namely MF-SetRank and Deep-SetRank, based on the above two models, respectively.

MF-SetRank. MF-SetRank is based on the popular collaborative model, Probabilistic Matrix Factorization (PMF) (Mnih and Salakhutdinov 2008). PMF factorizes the score matrix into two factor matrices representing user and item latent features. Along this line, we have $X = U^T V$, where $U \in \mathbb{R}^{r \times N}$ and $V \in \mathbb{R}^{r \times M}$ are latent user and item matrices, respectively. Then the prior probabilities over

Algorithm 1 Gradient update for V when fixing U

Require: $V, U, \gamma, decay, \lambda, \{P_i, \tilde{O}_i | 1 \leq i \leq N\}$
Ensure: V

```

1:  $grad = \lambda \cdot V$ 
2: for  $i = 1$  to  $N$  do
3:   Precompute  $g_l = u_i^T v_l$ , for  $\forall l \in P_i \cup \tilde{O}_i$ 
4:   Initialize  $tmp = 0, totalsum = 0, sum = 0, s[l] = 0$  for  $\forall l \in P_i, c[l] = 0$  for  $\forall l \in P_i \cup \tilde{O}_i$ 
5:   for  $l \in \tilde{O}_i$  do
6:      $sum += \exp(g_l)$ 
7:   end for
8:   for  $l \in P_i$  do
9:      $c[l] -= g_l \cdot (1 - g_l)$ 
10:     $s[l] = sum + \exp(g_l)$ 
11:     $totalsum += 1/s[l]$ 
12:   end for
13:   for  $l \in \tilde{O}_i$  do
14:      $c[l] += \exp(g_l) \cdot g_l \cdot (1 - g_l) \cdot totalsum$ 
15:   end for
16:   for  $l \in P_i$  do
17:      $c[l] += \exp(g_l) \cdot g_l \cdot (1 - g_l) / s[l]$ 
18:   end for
19:   for  $l \in P_i \cup \tilde{O}_i$  do
20:      $grad[:, l] += c[l] \cdot u_i$ 
21:   end for
22: end for
23:  $V -= \gamma \cdot grad$ 
24:  $\gamma *= decay$ 
25: Return  $V$ 

```

columns of U, V are assumed to be the normal distribution, i.e., $p(u_i) \sim \mathcal{N}(0, \lambda^{-1}I)$ and $p(v_l) \sim \mathcal{N}(0, \lambda^{-1}I)$, where λ is the regularization parameter. In this way, we can transform Equation 6 to the following form:

$$L = \sum_{i=1}^N \sum_{j \in P_i} -\log \frac{\phi(u_i^T v_j)}{\phi(u_i^T v_j) + \sum_{k \in O_i} \phi(u_i^T v_k)} + \frac{\lambda}{2} \left(\sum_{i=1}^N \|u_i\|^2 + \sum_{l=1}^M \|v_l\|^2 \right). \quad (7)$$

For the ease of calculation, we let $\log \phi(x)$ be the sigmoid function, i.e., $\log \phi(x) = \sigma(x) = 1/(1 + e^{-x})$. It is easy to verify that such $\phi(x)$ is an increasing and strictly positive function. Besides, this choice is also beneficial for bounding the excess risk which we will discuss in the next subsection.

Another notable thing is that we do not have to go through all the unobserved items for every user in each epoch, considering that the positive feedback is much more influential than the unobserved feedback. Following Wu, Hsieh, and Sharpnack(2018), we can randomly sample $\tilde{K}_i = \tau \cdot J_i$ unobserved items in each epoch to compose the set \tilde{O}_i for replacing O_i in Equation 7.

In each epoch, we update the latent factors U and V by the gradients $\nabla_U L$ and $\nabla_V L$, respectively. The speed efficiency of recommender system is quite important (Wu,

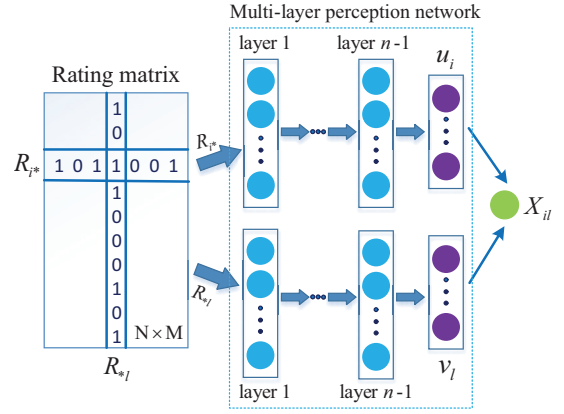


Figure 2: The modeling process of Deep-SetRank.

Hsieh, and Sharpnack 2017). Though a direct way to calculate the gradients costs $O(NJKr) = O(NJ^2r)$ time, where $J = \max\{J_i, 1 \leq i \leq N\}$ and $\tilde{K} = \tau \cdot J$, there is actually numerous repeated calculations here. We provide a clever approach in Algorithm 1 to rearrange the computation so that it only requires $O(N(J + \tilde{K})r) = O(NJr)$ time. Let us take the process of updating V as an example. Specifically, given a fixed latent user matrix U , the regularization parameter λ , the positive item set P_i , the unobserved item set \tilde{O}_i and the decaying rate $decay$ of the step size γ , Algorithm 1 shows how to update the gradients for latent item matrix V . Thus, MF-SetRank could run with a linear computational complexity, which is same as the efficient rating prediction methods based on matrix factorization (Mnih and Salakhutdinov 2008; Hu, Koren, and Volinsky 2008).

Deep-SetRank. In recent years, neural networks have shown good capacity on non-linear projection and embedding in recommender systems (Qin et al. 2019; He et al. 2017). Inspired by Deep Matrix Factorization (DeepMF) (Xue et al. 2017), we design a NN based setwise model called Deep-SetRank.

As shown in Figure 2, Deep-SetRank transforms the row R_{i*} and column R_{*l} of rating matrix R to obtain latent user and item matrices by user and item embedding networks, respectively. Then we still employ Equation 7 as the setwise loss function. Following DeepMF, we choose the multi-layer perception network (MLP) as the embedding network. Take user network as an example, we have

$$\begin{aligned} h_1 &= f_1(W_1 R_{i*} + b_1), \\ h_t &= f_t(W_t h_{t-1} + b_t), \quad t \in [2, n-1], \\ u_i &= f_n(W_n h_{n-1} + b_n), \end{aligned} \quad (8)$$

where h_t is the t -th hidden layer with weight matrix W_t and bias term b_t . For the activation function $f_t(\cdot)$, we employ the *sigmoid* function for the first $n-1$ layers and the *tanh* function for the last layer. Hence, we can predict the scores by the product of these two matrices. For each user, we simultaneously calculate the scores of items in both P_i and \tilde{O}_i in a batch for optimizing setwise loss. Different from MF-

SetRank, Deep-SetRank needs to train two neural networks rather than latent matrices.

Theoretical Analysis

In this subsection, we aim at giving the theoretical bound for the excess risk, i.e., the expected difference between the estimate and the truth, of SetRank. Without loss of generality, we assume that all the users have the same number of positive items and unobserved items for the sake of convenience. Hence, we have $J = J_i, K = K_i$ for $\forall i$. Note that the result can be readily generalized to the individual setting.

Considering the following constrained optimization of a general setwise method:

$$\hat{X} := \arg \min_X -\log p(>_{total} | X) \text{ such that } X \in \mathcal{X}, \quad (9)$$

where \mathcal{X} is the feasible set. Usually, X is constrained by the norm regularization to satisfy the low-rank condition. For example, in the personalized collaborative setting, $\mathcal{X} = \{X | X = U^T V, \|U\|_F \leq c_u, \|V\|_F \leq c_v\}$. Here $\|\cdot\|_F$ is the Frobenius norm. Supposing there is a $X^* \in \mathcal{X}$ such that $>_{total}$ is generated from $p(>_{total} | X^*)$. Then the excess risk is given in the form of KL divergence between the real and estimated probability: $D(X^*, \hat{X}) := \frac{1}{N} \sum_{i=1}^N \mathbb{E} \log \frac{p(>_i | X_i^*)}{p(>_i | \hat{X}_i)}$.

So far, the state-of-the-art listwise method could bound the excess risk by $O_{\mathbb{P}}\left(\sqrt{rM/N} \ln M\right)$ in the personalized collaborative setting (Wu, Hsieh, and Sharpnack 2018). Here we will show that the bound of SetRank is $O_{\mathbb{P}}\left(\sqrt{rM/N}(1 + J/K)\right)$ owing to the weaker precondition. In practice, the positive feedback always accounts for merely a tiny fraction of the total items. So, we have $J/K \ll 1$, which makes the result sound.

First, we give another statistical interpretation of Equation 5 from the generative perspective:

Theorem 1 *Suppose there is a matrix $Y = \{Y_{il}\}_{N \times M}$. Each entry Y_{il} is independently drawn from an exponential distribution with rate $\phi(X_{il})$. For each row Y_i , let the J smallest entries form the set P_i and others form the set O_i . Then the ranking structure probability $p(>_{total} | X)$, i.e., the probability that entries in P_i are less than those in O_i , is exactly equal to the RHS of Equation 5.*

The proof for Theorem 1 can be found in the Appendix. From Theorem 1, we know that the setwise preference probability can also be seen as the probability of a ranking structure over the matrix Y composed of $N \times M$ independent exponential random variables. Thus, we could give the following theorem according to McDiarmid's inequality (McDiarmid 1989) and Dudley's chaining (Talagrand 2006):

Theorem 2 *Let $\mathcal{Z} := \{\log \phi(X) | X \in \mathcal{X}\}$ be the image of element-wise function $\log \phi(X)$ and $\|\cdot\|_{\infty, 2}$ be the $\infty, 2$ norm defined as $\|Z\|_{\infty, 2} := \sqrt{\sum_{i=1}^N \|Z_i\|_{\infty}^2}$, $Z \in \mathbb{R}^{N \times M}$. Denote $\mathcal{N}(\epsilon, \mathcal{Z}, \|\cdot\|_{\infty, 2})$ as the ϵ -covering number of \mathcal{Z} in $\infty, 2$ norm, which represents the fewest number of spherical*

balls of radius ϵ needed to completely cover \mathcal{Z} in the condition of $\infty, 2$ norm. Hence, if $\|Z_i\|_{\infty} \leq \alpha$ for $\forall i$, we have

$$D(X^*, \hat{X}) = O_{\mathbb{P}}\left(g(\mathcal{Z})\sqrt{M}/N(1 + J/K)\right), \quad (10)$$

where $g(\mathcal{Z}) = \int_0^{\infty} \sqrt{\ln \mathcal{N}(u, \mathcal{Z}, \|\cdot\|_{\infty, 2})} du$.

By Theorem 2, we are able to obtain a bound in the general setting of X . Particularly, in the personalized collaborative setting, we can obtain the further result as follows:

Theorem 3 *Suppose that $\|Z_i\|_{\infty} \leq \alpha$ for $\forall i$ and $\log \phi(x)$ is 1-Lipschitz, then in the personalized collaborative setting, we have*

$$D(X^*, \hat{X}) = O_{\mathbb{P}}\left(\sqrt{rM/N}(1 + J/K)\right). \quad (11)$$

The detailed proofs for Theorem 2 and 3 are in the Appendix. Theorem 3 shows that when fixing the rank of latent factors, we will have a better estimate with a larger number of users and a smaller number of items, which is accord with the intuition. Besides, a smaller J with a larger K is beneficial for bounding the excess risk.

Experiments

Experimental Settings

Datasets. We evaluated the performance of our SetRank method on four real-world datasets, i.e., *MovieLens*¹, *Kindle*², *Yahoo*³ and *CiteULike*⁴. *MovieLens* is a commonly used movie recommendation dataset. *Kindle* contains Amazon product ratings collected from Kindle Store. *Yahoo* (Marlin and Zemel 2009) contains ratings for songs from Yahoo! Music. *CiteULike* is composed of users' collections of articles on CiteULike website. Following Wu, Hsieh, and Sharpnack(2018), we took two steps for data pre-processing. First, the original data of *MovieLens*, *Kindle* and *Yahoo* are in the form of 5-star ratings. We transformed them into implicit data, where each entry was marked as 1/0, depending on whether the ratings are greater than 3. Second, in order to make sure we have adequate positive feedback for better evaluating the recommendation algorithms, we filtered out users with less than 60, 20, 10, 10 positive items in *MovieLens*, *Kindle*, *Yahoo* and *CiteULike*, respectively. After data filtering, there are totally 3,937 users and 3,533 items with 923,473 positive entries in *MovieLens*, 4,379 users and 3,774 items with 102,545 positive entries in *Kindle*, 4,664 users and 921 items with 82,384 positive entries in *Yahoo*, 4,123 users and 7,849 items with 135,365 positive entries in *CiteULike*.

Evaluation protocols. We randomly sampled 50% of positive items for each user to construct the training set in each dataset, while the maximum number of item samples for each user was set as 10. Then, we sampled 1 positive item of each user as the validation set. Meanwhile, the rest data were

¹<https://grouplens.org/datasets/movielens/>

²<http://jmcauley.ucsd.edu/data/amazon/>

³<https://webscope.sandbox.yahoo.com/catalog.php?datatype=r>

⁴<http://www.citeulike.org>

Table 1: The recommendation performance of different approaches. (Methods with notation * are our proposed methods. We conducted the paired t-tests to verify that all improvements by SetRank are statistically significant for $p < 0.05$.)

| Datasets | Methods | P@5 | P@10 | R@5 | R@10 | MAP@5 | MAP@10 |
|------------------|---------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|----------------------------|
| <i>MovieLens</i> | WMF | 0.5161 \pm 0.0050 | 0.4787 \pm 0.0039 | 0.0173 \pm 0.0002 | 0.0317 \pm 0.0003 | 0.4125 \pm 0.0066 | 0.3449 \pm 0.0051 |
| | BPR | 0.6330 \pm 0.0075 | 0.6058 \pm 0.0075 | 0.0213 \pm 0.0004 | 0.0401 \pm 0.0007 | 0.5403 \pm 0.0105 | 0.4835 \pm 0.0078 |
| | Cofiset | 0.6500 \pm 0.0050 | 0.6152 \pm 0.0033 | 0.0217 \pm 0.0003 | 0.0405 \pm 0.0002 | 0.5657 \pm 0.0060 | 0.5020 \pm 0.0042 |
| | SQL-Rank | 0.6609 \pm 0.0020 | 0.6227 \pm 0.0028 | 0.0227 \pm 0.0003 | 0.0421 \pm 0.0003 | 0.5749 \pm 0.0018 | 0.5081 \pm 0.0025 |
| | MF-SetRank* | 0.6762 \pm 0.0024 | 0.6398 \pm 0.0017 | 0.0231 \pm 0.0002 | 0.0427 \pm 0.0001 | 0.5940 \pm 0.0028 | 0.5296 \pm 0.0033 |
| | DeepMF | 0.6191 \pm 0.0136 | 0.5831 \pm 0.0120 | 0.0208 \pm 0.0006 | 0.0383 \pm 0.0009 | 0.5270 \pm 0.0159 | 0.4627 \pm 0.0143 |
| | Deep-BPR | 0.6624 \pm 0.0061 | 0.6242 \pm 0.0067 | 0.0225 \pm 0.0003 | 0.0416 \pm 0.0008 | 0.5798 \pm 0.0098 | 0.5141 \pm 0.0077 |
| | Deep-SQL | 0.6794 \pm 0.0074 | 0.6432 \pm 0.0024 | 0.0233 \pm 0.0004 | 0.0431 \pm 0.0003 | 0.5990 \pm 0.0076 | 0.5343 \pm 0.0036 |
| | Multi-VAE | 0.6806 \pm 0.0012 | 0.6429 \pm 0.0030 | 0.0235 \pm 0.0002 | 0.0434 \pm 0.0002 | 0.5996 \pm 0.0017 | 0.5338 \pm 0.0020 |
| | Deep-SetRank* | 0.6956 \pm 0.0030 | 0.6557 \pm 0.0013 | 0.0242 \pm 0.0003 | 0.0447 \pm 0.0003 | 0.6150 \pm 0.0039 | 0.5473 \pm 0.0019 |
| <i>Kindle</i> | WMF | 0.1281 \pm 0.0018 | 0.1105 \pm 0.0015 | 0.0577 \pm 0.0014 | 0.0957 \pm 0.0015 | 0.0830 \pm 0.0008 | 0.0580 \pm 0.0007 |
| | BPR | 0.1372 \pm 0.0030 | 0.1174 \pm 0.0024 | 0.0660 \pm 0.0023 | 0.1078 \pm 0.0030 | 0.0896 \pm 0.0019 | 0.0623 \pm 0.0012 |
| | Cofiset | 0.1438 \pm 0.0025 | 0.1223 \pm 0.0017 | 0.0690 \pm 0.0017 | 0.1127 \pm 0.0020 | 0.0940 \pm 0.0017 | 0.0653 \pm 0.0011 |
| | SQL-Rank | 0.1478 \pm 0.0024 | 0.1238 \pm 0.0008 | 0.0705 \pm 0.0010 | 0.1135 \pm 0.0009 | 0.0986 \pm 0.0017 | 0.0678 \pm 0.0009 |
| | MF-SetRank* | 0.1773 \pm 0.0027 | 0.1466 \pm 0.0019 | 0.0869 \pm 0.0010 | 0.1375 \pm 0.0016 | 0.1210 \pm 0.0025 | 0.0834 \pm 0.0015 |
| | DeepMF | 0.1484 \pm 0.0022 | 0.1236 \pm 0.0011 | 0.0679 \pm 0.0013 | 0.1085 \pm 0.0016 | 0.0992 \pm 0.0017 | 0.0684 \pm 0.0010 |
| | Deep-BPR | 0.1654 \pm 0.0025 | 0.1372 \pm 0.0013 | 0.0824 \pm 0.0010 | 0.1300 \pm 0.0017 | 0.1106 \pm 0.0024 | 0.0765 \pm 0.0014 |
| | Deep-SQL | 0.1754 \pm 0.0045 | 0.1435 \pm 0.0023 | 0.0842 \pm 0.0030 | 0.1320 \pm 0.0036 | 0.1201 \pm 0.0032 | 0.0823 \pm 0.0017 |
| | Multi-VAE | 0.1651 \pm 0.0019 | 0.1360 \pm 0.0010 | 0.0800 \pm 0.0013 | 0.1265 \pm 0.0011 | 0.1120 \pm 0.0015 | 0.0771 \pm 0.0009 |
| | Deep-SetRank* | 0.1837 \pm 0.0020 | 0.1502 \pm 0.0017 | 0.0894 \pm 0.0009 | 0.1402 \pm 0.0013 | 0.1266 \pm 0.0015 | 0.0867 \pm 0.0009 |
| <i>Yahoo</i> | WMF | 0.1813 \pm 0.0021 | 0.1465 \pm 0.0011 | 0.1297 \pm 0.0015 | 0.2048 \pm 0.0017 | 0.1182 \pm 0.0015 | 0.0772 \pm 0.0010 |
| | BPR | 0.2096 \pm 0.0032 | 0.1729 \pm 0.0020 | 0.1475 \pm 0.0031 | 0.2389 \pm 0.0039 | 0.1378 \pm 0.0029 | 0.0928 \pm 0.0010 |
| | Cofiset | 0.2196 \pm 0.0041 | 0.1791 \pm 0.0018 | 0.1554 \pm 0.0034 | 0.2478 \pm 0.0032 | 0.1457 \pm 0.0034 | 0.0979 \pm 0.0016 |
| | SQL-Rank | 0.2137 \pm 0.0031 | 0.1723 \pm 0.0011 | 0.1502 \pm 0.0025 | 0.2380 \pm 0.0017 | 0.1432 \pm 0.0026 | 0.0951 \pm 0.0012 |
| | MF-SetRank* | 0.2267 \pm 0.0012 | 0.1817 \pm 0.0007 | 0.1616 \pm 0.0008 | 0.2540 \pm 0.0018 | 0.1528 \pm 0.0016 | 0.1009 \pm 0.0006 |
| | DeepMF | 0.2167 \pm 0.0019 | 0.1764 \pm 0.0009 | 0.1529 \pm 0.0021 | 0.2437 \pm 0.0032 | 0.1445 \pm 0.0024 | 0.0964 \pm 0.0012 |
| | Deep-BPR | 0.2260 \pm 0.0019 | 0.1831 \pm 0.0010 | 0.1615 \pm 0.0008 | 0.2549 \pm 0.0020 | 0.1513 \pm 0.0018 | 0.1013 \pm 0.0009 |
| | Deep-SQL | 0.2278 \pm 0.0023 | 0.1817 \pm 0.0013 | 0.1613 \pm 0.0025 | 0.2517 \pm 0.0029 | 0.1542 \pm 0.0018 | 0.1021 \pm 0.0009 |
| | Multi-VAE | 0.2300 \pm 0.0018 | 0.1855 \pm 0.0016 | 0.1631 \pm 0.0011 | 0.2572 \pm 0.0014 | 0.1550 \pm 0.0025 | 0.1036 \pm 0.0020 |
| | Deep-SetRank* | 0.2372 \pm 0.0011 | 0.1894 \pm 0.0015 | 0.1694 \pm 0.0007 | 0.2637 \pm 0.0022 | 0.1608 \pm 0.0016 | 0.1067 \pm 0.0013 |
| <i>CiteULike</i> | WMF | 0.1714 \pm 0.0018 | 0.1447 \pm 0.0010 | 0.0539 \pm 0.0009 | 0.0866 \pm 0.0009 | 0.1195 \pm 0.0023 | 0.0859 \pm 0.0012 |
| | BPR | 0.1876 \pm 0.0027 | 0.1612 \pm 0.0016 | 0.0644 \pm 0.0009 | 0.1056 \pm 0.0012 | 0.1320 \pm 0.0027 | 0.0964 \pm 0.0019 |
| | Cofiset | 0.1881 \pm 0.0022 | 0.1590 \pm 0.0009 | 0.0664 \pm 0.0013 | 0.1058 \pm 0.0016 | 0.1317 \pm 0.0018 | 0.0953 \pm 0.0008 |
| | SQL-Rank | 0.1801 \pm 0.0024 | 0.1522 \pm 0.0015 | 0.0576 \pm 0.0008 | 0.0932 \pm 0.0010 | 0.1283 \pm 0.0019 | 0.0927 \pm 0.0013 |
| | MF-SetRank* | 0.2124 \pm 0.0016 | 0.1813 \pm 0.0014 | 0.0764 \pm 0.0007 | 0.1238 \pm 0.0015 | 0.1523 \pm 0.0018 | 0.1117 \pm 0.0010 |
| | DeepMF | 0.1853 \pm 0.0018 | 0.1555 \pm 0.0015 | 0.0600 \pm 0.0010 | 0.0962 \pm 0.0015 | 0.1296 \pm 0.0019 | 0.0935 \pm 0.0014 |
| | Deep-BPR | 0.2086 \pm 0.0009 | 0.1754 \pm 0.0015 | 0.0767 \pm 0.0012 | 0.1224 \pm 0.0014 | 0.1483 \pm 0.0015 | 0.1075 \pm 0.0004 |
| | Deep-SQL | 0.2077 \pm 0.0027 | 0.1733 \pm 0.0016 | 0.0730 \pm 0.0021 | 0.1152 \pm 0.0015 | 0.1478 \pm 0.0029 | 0.1064 \pm 0.0017 |
| | Multi-VAE | 0.2081 \pm 0.0027 | 0.1720 \pm 0.0016 | 0.0775 \pm 0.0014 | 0.1205 \pm 0.0022 | 0.1502 \pm 0.0031 | 0.1068 \pm 0.0015 |
| | Deep-SetRank* | 0.2233 \pm 0.0026 | 0.1856 \pm 0.0017 | 0.0830 \pm 0.0019 | 0.1300 \pm 0.0016 | 0.1606 \pm 0.0022 | 0.1159 \pm 0.0014 |

used for test. In this way, we randomly split each dataset five times and reported all the results by mean values. To evaluate the performance, we adopted three widely used evaluation metrics, i.e., $P@P$, $R@P$ and $MAP@P$ (Wu, Hsieh, and Sharpnack 2018; Wang, Wang, and Yeung 2015). For each user, P (Precision) $@P$ measures the ratio of correct prediction results among top- P items to P and R (Recall) $@P$ measures the ratio of correct prediction results among top- P items to all positive items. Furthermore, MAP (Mean Average Precision) $@P$ considers the ranking of correct prediction results among top- P items. The final results of three metrics are given in the average of all users.

Baselines. The recommendation methods for comparison are listed as follows:

- **WMF:** Weighted Matrix Factorization (Hu, Koren, and Volinsky 2008) is a popular rating prediction method for implicit data, which introduces the confidence levels into standard matrix factorization model.
- **BPR:** Bayesian Personalized Ranking (Rendle et al.

2009) is a widely used pairwise collaborative ranking approach, which transforms the original rating matrix into the form of independent pairs.

- **Cofiset:** Cofiset (Pan and Chen 2013a) defines the group preference as the mean value of each item in the group. Then the BPR loss function is used for optimization.
- **SQL-Rank:** Stochastic Queuing Listwise Ranking (Wu, Hsieh, and Sharpnack 2018) is a state-of-the-art listwise approach, which breaks ties randomly and generates multiple possible permutations.
- **DeepMF:** Deep Matrix Factorization (Xue et al. 2017) is a NN based matrix factorization model.
- **Multi-VAE:** Variational Autoencoders for Collaborative Filtering (Liang et al. 2018) is a state-of-the-art NN based method, which extends variational autoencoders to recommendations for implicit feedback.
- **Deep-BPR, Deep-SQL:** These two methods have the same network architecture with Deep-SetRank, but we replace the loss function in SetRank by those in BPR and




Figure 3: The performance of P@5 with different values of sampling ratio τ on the four datasets.




Figure 4: The performance of P@5 with different values of dimension r on the four datasets.

SQL-Rank, respectively. Hence, we obtain these two NN based pairwise and listwise approaches.

- **MF-SetRank, Deep-SetRank:** These two methods are our proposed setwise Bayesian approaches for collaborative ranking from implicit feedback. We release our code at <https://github.com/chadwang2012/SetRank>.

Please note that WMF, BPR, SQL-Rank, MF-SetRank are all implemented with a basic matrix factorization model and the four “Deep” methods are all implemented with the same neural network architecture, so it is a fair setting to compare the performances of different item ranking approaches.

Parameter settings. For the above baselines, we have carefully explored the corresponding parameters, i.e., the number of dimensions and regularization parameters. Besides, for SQL-Rank, we chose the ratio of subsampled unobserved items to positive items as 3 : 1 following the authors’ guidance. For MF-SetRank, we tuned the learning rate in $[0.1, 0.2, \dots, 1.0]$ and the decay rate in $[0.9, 0.93, 0.95, 0.97, 0.99]$. We also fixed the sampling ratio τ to 3. Then we tuned the number of dimensions r in $[50, 100, 150, 200, 250, 300]$ and the regularization parameter λ in $[0.2, 0.3, \dots, 1.9, 2.0]$.

For Multi-VAE, we set the encoder as 2-layer MLP with dimensions 600×200 and decoder with dimensions 200×600 . For the other four “Deep” methods, we fixed the user network as 2-layer MLP with dimensions 512×100 and item network with dimensions 1024×100 . Then we performed Adam (Kingma and Ba 2014) algorithm for optimization and tune the learning rate from 0.0001 to 0.01.

Overall Performance Comparison

We present the overall recommendation performance results of the nine methods in Table 1 under two types of

settings, i.e., $P = 5$ and $P = 10$, since the top recommended items are much more important in practical scenes. As shown in the results, Deep-SetRank achieves the best performance against all the baseline methods on every dataset. Specifically, Deep-SetRank outperforms the best baselines by an average relative boost of 4.28% for the metric P@5 on the four datasets. Besides, MF-SetRank achieves the best performance against all the other MF based baselines. Specifically, MF-SetRank outperforms the state-of-the-art MF based method, SQL-Rank, by an average relative boost of 11.57% for the metric P@5. We can also observe that NN based models have stronger embedding ability and can perform better than MF based models. Nevertheless, it is notable that MF-SetRank has achieved comparable performances with NN based methods, such as Multi-VAE and Deep-SQL. The outstanding performances clearly demonstrate the effectiveness of our setwise approaches. We can also observe that SetRank achieves the largest relative boost to the other baselines on the sparsest dataset, *CiteULike*, which shows its superior capacity for handling sparsity problem. Another notable thing is that listwise approaches seem to perform better in top-5 metrics than top-10 metrics. This is probably because they pay more attention to the top ranks in an item list. On the opposite, SetRank treats every positive item or every unobserved item fairly thus can perform well in both top-5 and top-10 metrics.

Hyper-parameter Investigations

Effectiveness of negative sampling. As mentioned in Section Implementation, it is unnecessary to utilize all the unobserved items for gradient calculations in SetRank. We can just randomly sample $\tau \cdot J_i$ negative items for each user i in each epoch. Since the number of positive items J_i is

usually far smaller than the number of total items, there are few unobserved item overlaps for each user among different epochs. In this subsection, we fix all the other parameters to be the same and evaluate the influence of sampling ratio τ on final recommendation results. The P@5 results are shown in Figure 3. We find that when $\tau = 3$, the performance is good enough. Even if we further enlarge the value of τ , the result would not increase significantly.

Sensitivities of latent factors. In this paper, we factorize the score matrix into the product of user and item latent factors in a low-rank space. Therefore, the rank r of latent space is quite influential to the result. If the rank r is too small, the model could not fit the real-world data well while if r is too larger, it may cause the overfitting problem. We varied r to train our method and then presented the results in Figure 4. We can observe that the performance result of SetRank is not good when $r = 50$. With a larger value of r , the performance of MF-SetRank tends to be much better. Thus, we suggest adopting a large value for r to get the best performance in MF-SetRank. By comparison, $r = 100$ seems to be good enough for Deep-SetRank.

Conclusion

In this paper, we proposed a setwise Bayesian approach, namely SetRank, for collaborative ranking. SetRank has the ability in accommodating the characteristic of implicit feedback in recommender systems. Specifically, we first designed a novel setwise preference structure. Then, we maximized the posterior probability of the setwise preference structure to complete the Bayesian inference. In particular, we designed two implementations, MF-SetRank and Deep-SetRank. Moreover, we provided the theoretical analysis of SetRank to show that the bound of excess risk can be proportional to $\sqrt{M/N}$. Finally, extensive experiments on four real-world datasets clearly validated the advantages of SetRank over various state-of-the-art baselines.

Acknowledgments

This work was supported by grants from the National Natural Science Foundation of China (No.91746301, 61836013).

Appendix

Proof of Theorem 1

Proof. Owing to the independence assumption, the following equation holds:

$$\begin{aligned} p(>_{total} | X) &= \prod_{i=1}^N p(>_i | X_i) \\ &= \prod_{i=1}^N \prod_{j \in P_i} p(Y_{ij} \leq \min_{k \in O_i} \{Y_{ik}\} | X_i), \end{aligned} \quad (12)$$

where $\min_{k \in O_i} \{Y_{ik}\}$ obeys an exponential distribution with rate $\sum_{k \in O_i} \phi(X_{ik})$. Then we have

$$\begin{aligned} p(Y_{ij} \leq \min_{k \in O_i} \{Y_{ik}\} | X_i) &= \int_0^\infty \phi(X_{ij}) e^{-u\phi(X_{ij})} e^{-\sum_{k \in O_i} u\phi(X_{ik})} du \\ &= \frac{\phi(X_{ij})}{\phi(X_{ij}) + \sum_{k \in O_i} \phi(X_{ik})}. \end{aligned} \quad (13)$$

With Equation 12 and 13, we obtain the conclusion. \square

Proof of Theorem 2 and Theorem 3

To prove Theorem 2, we first follow (Wu, Hsieh, and Sharpnack 2018) to propose an important lemma to bound the excess risk by an empirical process term.

Lemma 1 *Supposing $\hat{X} := \arg \min_X -\log p(>_{total} | X)$ such that $X \in \mathcal{X}$ and there is a $X^* \in \mathcal{X}$ such that $>_{total}$ is generated from $p(>_{total} | X^*)$. Then we have the following inequality, where \mathbb{E} is for the draw of $>_i$:*

$$D(X^*, \hat{X}) \leq -\frac{1}{N} \sum_{i=1}^N \left(\log \frac{p(>_i | X_i^*)}{p(>_i | \hat{X}_i)} - \mathbb{E} \log \frac{p(>_i | X_i^*)}{p(>_i | \hat{X}_i)} \right). \quad (14)$$

Proof. Due to the optimality condition, we have

$$\sum_{i=1}^N -\log p(>_i | \hat{X}_i) \leq \sum_{i=1}^N -\log p(>_i | X_i^*). \quad (15)$$

Actually, Equation 15 is equivalent to

$$\frac{1}{N} \sum_{i=1}^N \log \frac{p(>_i | X_i^*)}{p(>_i | \hat{X}_i)} \leq 0.$$

Thus, it is easy to obtain the conclusion. \square

As we can see from Lemma 1, if we fix \hat{X} , the empirical process term (the RHS of Equation 14) is a random function of the preference structure $>_{total}$ with mean zero. However, \hat{X} is also random so that we have to uniformly bound the empirical process term over $\hat{X} \in \mathcal{X}$. To apply Dudley's chaining (Talagrand 2006), we first bound the variations between two preference scores X_i and X'_i with Lemma 2:

Lemma 2 *Define the difference function $\Delta(>_i | X_i, X'_i) := \log \frac{p(>_i | X_i)}{p(>_i | X'_i)}$. If a single entry Y_{il} changes, it would cause the transformation of setwise preference structure, i.e., $>_i$ would be converted into $>'_i$. We can bound the variations of the difference function in the form of:*

$$\begin{aligned} |\Delta(>_i | X_i, X'_i) - \Delta(>'_i | X_i, X'_i)| &\leq C \|\log \phi(X_i) - \log \phi(X'_i)\|_\infty, \end{aligned} \quad (16)$$

where $C = 2 + 2e^{2\alpha} J/K$.

Proof. If the change of Y_{il} does not lead to the change of P_i and O_i , there is no influence for preference structure, i.e., $>_i = >'_i$ and $|\Delta(>_i | X_i, X'_i) - \Delta(>'_i | X_i, X'_i)| = 0$.

Otherwise, we assume that item $j' \in P_i$ and $k' \in O_i$ exchange their status with each other so that in the new preference structure $>'_i$, we have $j' \in O'_i$ and $k' \in P'_i$. In the following part of the proof, for ease of the statement, we denote $\lambda_l = \phi(X_{il})$ and $\Lambda_j = \lambda_j + \sum_{k \in O_i \setminus \{k'\}} \lambda_k$. λ'_l and Λ'_j are defined analogously with X' .

So, we have

$$\Delta(>_i | X_i, X'_i) = \sum_{j \in P_i} \left(\log \frac{\lambda_j}{\lambda'_j} - \log \frac{\Lambda_j + \lambda_{k'}}{\Lambda'_j + \lambda'_{k'}} \right),$$

$$\begin{aligned} \Delta(>'_i | X_i, X'_i) &= \sum_{j \in P_i \setminus \{j'\}} \left(\log \frac{\lambda_j}{\lambda'_j} - \log \frac{\Lambda_j + \lambda_{j'}}{\Lambda'_j + \lambda'_{j'}} \right) \\ &\quad + \log \frac{\lambda_{k'}}{\lambda'_{k'}} - \log \frac{\Lambda_{j'} + \lambda_{k'}}{\Lambda'_{j'} + \lambda'_{k'}}, \end{aligned}$$

and thus,

$$\begin{aligned} &|\Delta(>_i | X_i, X'_i) - \Delta(>'_i | X_i, X'_i)| \\ &= \left| \log \frac{\lambda_{j'}}{\lambda'_{j'}} - \log \frac{\lambda_{k'}}{\lambda'_{k'}} \right. \\ &\quad \left. + \sum_{j \in P_i \setminus \{j'\}} \left(\log \frac{\Lambda_j + \lambda_{j'}}{\Lambda'_j + \lambda'_{j'}} - \log \frac{\Lambda_j + \lambda_{k'}}{\Lambda'_j + \lambda'_{k'}} \right) \right| \\ &\leq \left| \log \frac{\lambda_{j'}}{\lambda'_{j'}} - \log \frac{\lambda_{k'}}{\lambda'_{k'}} \right| \\ &\quad + \sum_{j \in P_i \setminus \{j'\}} \left| \log \frac{\Lambda_j + \lambda_{j'}}{\Lambda'_j + \lambda'_{j'}} - \log \frac{\Lambda_j + \lambda_{k'}}{\Lambda'_j + \lambda'_{k'}} \right|. \end{aligned}$$

Notice that

$$\begin{aligned} &\left| \log \frac{\Lambda_j + \lambda_{j'}}{\Lambda'_j + \lambda'_{j'}} - \log \frac{\Lambda_j + \lambda_{k'}}{\Lambda'_j + \lambda'_{k'}} \right| \\ &\leq \left| \log \frac{\Lambda_j + \lambda_{j'}}{\Lambda'_j + \lambda'_{j'}} - \log \frac{\Lambda_j}{\Lambda'_j} \right| + \left| \log \frac{\Lambda_j}{\Lambda'_j} - \log \frac{\Lambda_j + \lambda_{k'}}{\Lambda'_j + \lambda'_{k'}} \right| \\ &= \left| \log \frac{\Lambda_j + \lambda_{j'}}{\Lambda_j} - \log \frac{\Lambda'_j + \lambda'_{j'}}{\Lambda'_j} \right| \\ &\quad + \left| \log \frac{\Lambda_j + \lambda_{k'}}{\Lambda_j} - \log \frac{\Lambda'_j + \lambda'_{k'}}{\Lambda'_j} \right|. \end{aligned}$$

Hence, we let $\delta = \|\log \phi(X_i) - \log \phi(X'_i)\|_\infty$, and have

$$\left| \log \frac{\Lambda_j}{\Lambda'_j} \right| \leq \max_l \left| \log \frac{\lambda_l}{\lambda'_l} \right| \leq \delta.$$

Further, we assume $\beta_j = \max\{\frac{\lambda_{j'}}{\Lambda_j}, \frac{\lambda'_{j'}}{\Lambda'_j}\}$, and then

$$\begin{aligned} &\left| \log \left(1 + \frac{\lambda_{j'}}{\Lambda_j} \right) - \log \left(1 + \frac{\lambda'_{j'}}{\Lambda'_j} \right) \right| \\ &\leq |\log(1 + \beta_j) - \log(1 + e^{-2\delta}\beta_j)| \leq |1 - e^{-2\delta}| |\beta_j|. \end{aligned}$$

Considering that we have $\Lambda_j \geq Ke^{-\alpha}$, we can derive $\beta_j \leq e^{2\alpha}/K$. Similar conclusion can be obtained for $\left| \log \left(1 + \frac{\lambda_{k'}}{\Lambda_j} \right) - \log \left(1 + \frac{\lambda'_{k'}}{\Lambda'_j} \right) \right|$.

Synthesize the analysis above, we thus have

$$\begin{aligned} &|\Delta(>_i | X_i, X'_i) - \Delta(>'_i | X_i, X'_i)| \\ &\leq 2\delta + 2 \sum_{j \in P_i \setminus \{j'\}} |1 - e^{-2\delta}| e^{2\alpha}/K \\ &\leq \delta(2 + 2e^{2\alpha}J/K). \end{aligned}$$

Consequently, let $C = 2 + 2e^{2\alpha}J/K$ and we can come to the conclusion. \square

Proof of Theorem 2. The empirical process function is defined as

$$\rho_N(x) := -\frac{1}{N} \sum_{i=1}^N \left(\log \frac{p(>_i | X_i^*)}{p(>_i | \hat{X}_i)} - \mathbb{E} \log \frac{p(>_i | X_i^*)}{p(>_i | \hat{X}_i)} \right).$$

From Theorem 1, we know that $\rho_N(x)$ is a function of $N \times M$ independent exponential random variables. And from Lemma 2, we know that the change of preference structure caused by the change of a single entry Y_{il} is bounded. Specifically, the accumulative squares of bounds are

$$\sum_{i=1}^N \sum_{l=1}^M C^2 \|\log \phi(X_i) - \log \phi(X'_i)\|_\infty^2 = MC^2 \|Z - Z'\|_{\infty,2}^2.$$

Then according to McDiarmid's inequality (McDiarmid 1989), we have

$$p\{N(\rho_N(x) - \rho_N(x')) > \epsilon\} \leq \exp\left(\frac{-2\epsilon^2}{MC^2 \|Z - Z'\|_{\infty,2}^2}\right).$$

As a result, the stochastic process $\{N\rho_N(X) | X \in \mathcal{X}\}$ is a subGaussian field with canonical distance $d(X, X') = \sqrt{MC} \|Z - Z'\|_{\infty,2}$. Following Dudley's chaining (Talagrand 2006), we can get the conclusion. \square

Proof of Theorem 3. Wu, Hsieh, and Sharpnack(2018) have proved that $g(\mathcal{Z}) \leq c'\sqrt{rN}$ in the personalized collaborative setting, where c' is an absolute constant. Thus we can conclude the proof immediately. \square

References

Cao, Z.; Qin, T.; Liu, T.-Y.; Tsai, M.-F.; and Li, H. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*, 129–136. ACM.

- Chapelle, O., and Keerthi, S. S. 2010. Efficient algorithms for ranking with svms. *Information retrieval* 13(3):201–215.
- Chen, W.; Liu, T.-Y.; Lan, Y.; Ma, Z.-M.; and Li, H. 2009. Ranking measures and loss functions in learning to rank. In *NIPS*, 315–323.
- Freund, Y.; Iyer, R.; Schapire, R. E.; and Singer, Y. 2003. An efficient boosting algorithm for combining preferences. *Journal of machine learning research* 4(Nov):933–969.
- He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *WWW*, 173–182. International World Wide Web Conferences Steering Committee.
- Hsieh, C.-J.; Natarajan, N.; and Dhillon, I. S. 2015. Pu learning for matrix completion. In *ICML*, 2445–2453.
- Hu, Y.; Koren, Y.; and Volinsky, C. 2008. Collaborative filtering for implicit feedback datasets. In *ICDM*, volume 8, 263–272. Citeseer.
- Huang, S.; Wang, S.; Liu, T.-Y.; Ma, J.; Chen, Z.; and Vejalainen, J. 2015. Listwise collaborative filtering. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 343–352. ACM.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krohn-Grimberghe, A.; Drumond, L.; Freudenthaler, C.; and Schmidt-Thieme, L. 2012. Multi-relational matrix factorization using bayesian personalized ranking for social network data. In *Proceedings of the fifth ACM international conference on Web search and data mining*, 173–182. ACM.
- Liang, D.; Krishnan, R. G.; Hoffman, M. D.; and Jebara, T. 2018. Variational autoencoders for collaborative filtering. In *WWW*, 689–698. International World Wide Web Conferences Steering Committee.
- Liu, H.; Li, T.; Hu, R.; Fu, Y.; Gu, J.; and Xiong, H. 2019a. Joint representation learning for multi-modal transportation recommendation. In *2019 AAAI Conference on Artificial Intelligence (AAAI'19)*, 1036–1043.
- Liu, H.; Tong, Y.; Zhang, P.; Lu, X.; Duan, J.; and Xiong, H. 2019b. Hydra: A personalized and context-aware multi-modal transportation recommendation system. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2314–2324.
- Marlin, B. M., and Zemel, R. S. 2009. Collaborative prediction and ranking with non-random missing data. In *Proceedings of the third ACM conference on Recommender systems*, 5–12. ACM.
- McDiarmid, C. 1989. On the method of bounded differences. *Surveys in combinatorics* 141(1):148–188.
- Mnih, A., and Salakhutdinov, R. R. 2008. Probabilistic matrix factorization. In *Advances in neural information processing systems*, 1257–1264.
- Pan, W., and Chen, L. 2013a. Cofiset: Collaborative filtering via learning pairwise preferences over item-sets. In *Proceedings of the 2013 SIAM international conference on data mining*, 180–188. SIAM.
- Pan, W., and Chen, L. 2013b. Gbpr: group preference based bayesian personalized ranking for one-class collaborative filtering. In *Twenty-Third International Joint Conference on Artificial Intelligence*.
- Qin, C.; Zhu, H.; Zhu, C.; Xu, T.; Zhuang, F.; Ma, C.; Zhang, J.; and Xiong, H. 2019. Duerquiz: A personalized question recommender system for intelligent job interview. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2165–2173. ACM.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, 452–461. AUAI Press.
- Shi, Y.; Larson, M.; and Hanjalic, A. 2010. List-wise learning to rank with matrix factorization for collaborative filtering. In *Proceedings of the fourth ACM conference on Recommender systems*, 269–272. ACM.
- Talagrand, M. 2006. *The generic chaining: upper and lower bounds of stochastic processes*. Springer Science & Business Media.
- Wang, S.; Huang, S.; Liu, T.-Y.; Ma, J.; Chen, Z.; and Vejalainen, J. 2016. Ranking-oriented collaborative filtering: A listwise approach. *ACM Transactions on Information Systems (TOIS)* 35(2):10.
- Wang, H.; Wang, N.; and Yeung, D.-Y. 2015. Collaborative deep learning for recommender systems. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1235–1244. ACM.
- Weimer, M.; Karatzoglou, A.; Le, Q. V.; and Smola, A. J. 2008. Cofi rank-maximum margin matrix factorization for collaborative ranking. In *Advances in neural information processing systems*, 1593–1600.
- Wu, L.; Hsieh, C.-J.; and Sharpnack, J. 2017. Large-scale collaborative ranking in near-linear time. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 515–524. ACM.
- Wu, L.; Hsieh, C.-J.; and Sharpnack, J. 2018. Sql-rank: A listwise approach to collaborative ranking. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, 5315–5324.
- Xia, F.; Liu, T.-Y.; Wang, J.; Zhang, W.; and Li, H. 2008. Listwise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th international conference on Machine learning*, 1192–1199. ACM.
- Xia, L. 2019. Learning and decision-making from rank data. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 13(1):1–159.
- Xue, H.-J.; Dai, X.; Zhang, J.; Huang, S.; and Chen, J. 2017. Deep matrix factorization models for recommender systems. In *IJCAI*, 3203–3209.
- Zhu, C.; Zhu, H.; Xiong, H.; Ma, C.; Xie, F.; Ding, P.; and Li, P. 2018. Person-job fit: Adapting the right talent for the right job with joint representation learning. *ACM Transactions on Management Information Systems (TMIS)* 9(3):12.