

Multiple Pairwise Ranking with Implicit Feedback

Runlong Yu[§], Yunzhou Zhang[§], Yuyang Ye[§], Le Wu[†], Chao Wang[§], Qi Liu^{§*}, Enhong Chen[§]

[§]Anhui Province Key Laboratory of Big Data Analysis and Application, University of Science and Technology of China
{yrnl, zyunzhou, yeyuyang, wdyx2012}@mail.ustc.edu.cn, {qiliuq, cheneh}@ustc.edu.cn

[†]Hefei University of Technology, lewu@hfut.edu.cn

ABSTRACT

As users implicitly express their preferences to items on many real-world applications, the implicit feedback based collaborative filtering has attracted much attention in recent years. Pairwise methods have shown state-of-the-art solutions for dealing with the implicit feedback, with the assumption that users prefer the observed items to the unobserved items. However, for each user, the huge unobserved items are not equal to represent her preference. In this paper, we propose a Multiple Pairwise Ranking (MPR) approach, which relaxes the simple pairwise preference assumption in previous works by further tapping the connections among items with multiple pairwise ranking criteria. Specifically, we exploit the preference difference among multiple pairs of items by dividing the unobserved items into different parts. Empirical studies show that our algorithms outperform the state-of-the-art methods on real-world datasets.

CCS CONCEPTS

• Information systems → Collaborative filtering; Recommender systems

KEYWORDS

Collaborative filtering; Implicit feedback; Item recommendation

ACM Reference format:

Runlong Yu, Yunzhou Zhang, Yuyang Ye, Le Wu, Chao Wang, Qi Liu, and Enhong Chen. 2018. Multiple Pairwise Ranking with Implicit Feedback. In *2018 ACM Conference on Information and Knowledge Management (CIKM'18)*, October 22–26, 2018, Torino, Italy, 4 pages. <https://doi.org/10.1145/3269206.3269283>

1 INTRODUCTION

Collaborative filtering methods have become wildly used technologies in recommender systems [1, 2]. In many real-world scenarios, explicit feedback is not always available. On the contrary, there

are lots of data in one-class form, e.g., purchases in Tmall, likes in WeChat, watches in Netflix. Such data do not contain the scoring between users and items, which are usually called one-class or implicit feedback. Implicit feedback is different from explicit feedback: the latter explicitly express users' positive and negative preferences through the rating scores, while the former only contains positive feedback. Hence, the huge unobserved item feedbacks cannot be treated as negative preferences, as the most possible reason is that users have not seen them before.

Previous methods for dealing with the implicit feedback can be divided into two groups [3]: (1) pointwise regression methods, and (2) pairwise ranking methods. Pointwise methods take implicit feedback as absolute preference scores and minimize a pointwise square loss to approximate the absolute rating scores. Pairwise methods take pairs of items as basic units and try to maximize the likelihood of pairwise preferences over the observed items and the unobserved items. Bayesian Personalized Ranking (BPR) [4] is one of the most popular approaches that adopt such pairwise preference assumption. For an observed (user, item) interaction (u, i) and an unobserved (user, item) interaction (u, j) , BPR assumes that a user u has a higher preference on item i than on item j .

Many pairwise methods improve over BPR [5, 6]. These approaches have shown better performance compared to the pointwise based methods [7]. However, we argue that the huge unobserved item preferences are not well exploited in these pairwise based models. Specifically, all these pairwise based models inherited the assumption in BPR that users prefer the observed items to the unobserved items. Nevertheless, the unobserved item preference could be attributed to two aspects: the active user does not like the item, or simply the user does not observe or notice the item before. Thus, the core optimization idea of BPR that all the observed items should rank higher than all the unobserved items is too strict in the model design process.

In this paper, we propose a new pairwise ranking method, namely Multiple Pairwise Ranking (MPR), for attempting to solve the limitation. For each user, we further divide her unobserved item feedback into two parts: the possibly negative item feedback and the unknown preference for the items. Then, for each user, we have three parts of feedback: the positive feedback and two parts from the unobserved feedback. We take these three parts and further exploit the preference difference among multiple pairs of items, which can be thought of as a new multiple pairwise model. MPR is a novel basic algorithm and can be easily adopted by aforementioned works, like group preference, to further improve their performances.

* Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '18, October 22–26, 2018, Torino, Italy
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-6014-2/18/10.
<https://doi.org/10.1145/3269206.3269283>

2 THE PROPOSED METHOD

2.1 Preliminaries

Rendle et al. [4] proposed the BPR method, which takes pairs of items as basic units and tries to maximize the likelihood of pairwise preferences. The two fundamental assumptions adopted by BPR are as follows: (a) It assumes that user u prefers item i over all the unobserved items if item i has been observed by user u . (b) It assumes that the likelihood of pairwise preference of a user u is independent of the others. Mathematically, the likelihood of BPR among items can be formulated as

$$BPR = \prod_{u \in U} \prod_{i \in I_u^+} \prod_{j \in I - I_u^+} Pr(r_{ui} > r_{uj}) \times [1 - Pr(r_{uj} > r_{ui})] \quad (1)$$

where $U = \{u\}_{u=1}^n$ represents the set of users and $I = \{i\}_{i=1}^m$ represents the set of items, and each user $u \in U$ has expressed her positive feedbacks on items $I_u^+ \subset I$. r_{ui} denotes u 's preference on item i . The goal of implicit feedback problem is to recommend a personalized ranking list of items for each user u from the unobserved item set $I - I_u^+$.

Group Bayesian Personalized Ranking (GBPR) [5] relaxes assumption (b) among users in BPR by considering that users' preferences are influenced by other users with similar interests. The likelihood of GBPR among items can be given by

$$GBPR = \prod_{u \in U} \prod_{i \in I_u^+} \prod_{j \in I - I_u^+} Pr(r_{ugi} > r_{uj}) \times [1 - Pr(r_{uj} > r_{ugi})] \quad (2)$$

where $\mathcal{G} \subseteq U$ is a user group, $u \in \mathcal{G}$, $r_{ugi} = \frac{1}{|\mathcal{G}|} \sum_{w \in \mathcal{G}} r_{wi}$ is the group preference of users from the group \mathcal{G} on item i .

2.2 Multiple Pairwise Ranking

For each user u , besides the positive item set I_u^+ , the huge unobserved item feedback could be attributed to two reasons: she dislikes it or she has not seen it. Given this intuition, we relax the assumption (a) adopted by BPR with equal importance of the unobserved items. Specifically, we first consider three subsets contained in the set of items I .

- I_u^+ : The real positive items that user u has expressed her feedbacks on.
- I_u^- : The possibly negative items that user u has seen but not given her feedback on.
- I_u^* : The uncertainly negative items that user u has not seen.

It should be noted that $I_u^- \cup I_u^*$ makes up the unobserved items for u . Secondly, we define preference difference $r_{uij} = r_{ui} - r_{uj}$ as a difference value of user u 's preferences between item i and item j . Finally, we give an illustration of our preference assumption in Figure 1. Here we use "1" for "like", and "0" for "dislike", and "?" for "unclear". For a user u , we can suspect that the preferences $r_{ui} = r_{up} = r_{up'} = 1$, and $r_{uj} \approx r_{uq'} \approx 0$, and $r_{uq} = ?$ ($0 \leq ? \leq 1$) according to $i, p, p' \in I_u^+$, $j, q' \in I_u^-$, and $q \in I_u^*$. Since the user u has expressed her positive feedback on item i , and item j and item q' are possibly negative items for user u , we can believe r_{uij} is not less than $r_{uqq'}$ whether q is a positive or negative item for user u . Simultaneously, for p, p' , user u likes both of




Figure 1: Illustration of preference assumption.

them, so the difference value of user u 's preferences between them, namely $r_{upp'}$, can be small. Based on the assumption, we give a new criterion called Multiple Pairwise Ranking (MPR) showing overall likelihood for users and items,

$$MPR = \prod_{u \in U} \prod_{i, p, p' \in I_u^+} \prod_{j, q' \in I_u^-; q \in I_u^*} Pr(r_{uij} \geq r_{uqq'}, r_{uqq'} \geq r_{upp'}) \times [1 - Pr(r_{uij} < r_{uqq'}, r_{uqq'} < r_{upp'})] \quad (3)$$

In summary, our proposed model of three pairs of items further exploit user's preference difference between (I) an observed item and an unobserved item, and (II) two unobserved items, and (III) two observed items. The difference value of user's preferences between (I) is not less than (II), and (II) is not less than (III). Through the two inequalities, we relax the strict pairwise preference assumption in BPR and define a more realistic preference assumption. It could be considered that MPR digs implied positive feedbacks in the unobserved items by extracting the implicit partial ordering relation in the observed items, as well as in the unobserved items respectively in the optimizing process.

As the proposed MPR is a novel basic method, we can easily extend MPR to our new model Group Multiple Pairwise Ranking (GMPR) via introducing group preference of users. The likelihood of GMPR among items can be given by

$$GMPR = \prod_{u \in U} \prod_{i, p, p' \in I_u^+} \prod_{j, q' \in I_u^-; q \in I_u^*} Pr(r_{ugij} \geq r_{uqq'}, r_{uqq'} \geq r_{ugpp'}) \times [1 - Pr(r_{ugij} < r_{uqq'}, r_{uqq'} < r_{ugpp'})] \quad (4)$$

where r_{ugij} is the same as in Eq. (2), $r_{ugij} = r_{ugi} - r_{uj}$, $r_{ugpp'} = r_{ugp} - r_{ugp'}$.

2.3 Implicit Feedback Data Division

There are many ways to help us divide the unobserved items into I_u^- and I_u^* , and two of them are detailed. The first one is for online product recommendation problems. Many E-commerce platforms for online shopping record users' click data that can help us divide the unobserved data. An obvious idea is that: the products that the user has clicked and purchased are the real positive items, and the products that the user has clicked but not purchased are the possibly negative items, and the products that the user has not clicked are the uncertainly negative items.

However, additional click data are not always available. In another scenario like movie recommendation, "watches" are only feedbacks from customers. We give our second general method of dividing the unobserved items based on popularity. We rank items from I in descending order based on their observed counts by all users and choose the last part as the unpopular items I^e . We as-

sume that most people may not like the items in I^e . The previous researchers have found that the implied positive samples are comparatively few in the unobserved items [8], so we set $I_u^- = I^e - I_u^+$ in which have more chances to be the possibly negative items for u . As before, the uncertainly negative items $I_u^* = I - I_u^+ \cup I_u^-$.




Figure 2: Illustration of notations.

We illustrate this idea in Figure 2. User John likes movies “The Dark Knight” and “Alien”, which denote the items in I_u^+ . Several movies with few observed counts denote the unpopular items in I^e . We generally assume that movies in shadow part (items in I_u^-) may not interest John. It is hard to judge John’s preferences on items in I_u^* .

2.4 Learning the MPR

Based on the above explanation, here we represent $r_{uij} \geq r_{uqq'}$, $r_{uqq'} \geq r_{upp'}$ as $\lambda(r_{uij} - r_{uqq'}) + (1 - \lambda)(r_{uqq'} - r_{upp'})$, where λ is a trade-off parameter used to fuse their relation, and we abbreviate the above formula as $r_{>u}$. Following BPR, we use $\sigma(x) = \frac{1}{1+e^{-x}}$ to approximate the probability $Pr(\cdot)$. Then Eq. (3) can be simplified as $Pr(r_{>u})[1 - Pr(-r_{>u})] = \sigma(r_{>u})[1 - \sigma(-r_{>u})] = \sigma^2(r_{>u})$. Based on this trick, the objective function of MPR can be represented as follows,

$$\min_{\Theta} -\frac{1}{2} \ln MPR + \frac{1}{2} R(\Theta) \quad (5)$$

where $\Theta = \{U_u \in R^{1 \times d}, V_i \in R^{1 \times d}, b_i \in R, u \in U, i \in I\}$ is a set of model parameters to be learned, U_u is a latent feature vector describing user u , V_i is a latent feature vector describing item i , b_i is the bias of item i , and d is the number of latent factors.

$$\ln MPR = \prod_{u \in U} \prod_{i, p, p' \in I_u^+} \prod_{j, q' \in I_u^-, q \in I_u^*} 2 \ln \sigma(\lambda(r_{uij} - r_{uqq'}) + (1 - \lambda)(r_{uqq'} - r_{upp'})) \quad (6)$$

Eq. (6) is the log-likelihood of MPR. $R(\Theta) = \sum_{u \in U} \sum_{t \in S} [\alpha_u \|U_u\|^2 + \alpha_v \|V_t\|^2 + \beta_v \|b_t\|^2]$ is a regularization term to prevent overfitting in the learning process, and $S = \{i, j, p, p', q, q'\}$ is the sampled items. The individual preference is modeled by matrix factorization, for example $r_{ui} = U_u V_i^T + b_i$.

The optimization problem of the objective function in Eq. (5) can be solved by the widely used Stochastic Gradient Descent (SGD) algorithm. The process of SGD is to select a record, which includes a user u , six items containing i, j, p, p', q, q' , and iteratively update model parameters based on the sampled feedback records. The tentative objective function can be written as $f(u, S)$

$$\begin{aligned} &= -\ln \sigma(r_{>u}) + \frac{\alpha_u}{2} \|U_u\|^2 + \frac{\alpha_v}{2} \sum_{t \in S} \|V_t\|^2 + \frac{\beta_v}{2} \sum_{t \in S} \|b_t\|^2 = \\ &\ln[1 + \exp(-r_{>u})] + \frac{\alpha_u}{2} \|U_u\|^2 + \frac{\alpha_v}{2} \sum_{t \in S} \|V_t\|^2 + \frac{\beta_v}{2} \sum_{t \in S} \|b_t\|^2. \end{aligned}$$

With the above gradients, we can update the model parameters as $\Theta = \Theta - \gamma \frac{\partial f(u, S)}{\partial \Theta}$, where $\gamma > 0$ is the learning rate.

GMPR can be formulated in the same way. Here we directly give the objective function as $\min_{\Theta} -\frac{1}{2} \ln GMPR + R(\Theta_G)$, where $\ln GMPR = \sum_{u \in U} \sum_{i, p, p' \in I_u^+} \sum_{j, q' \in I_u^-, q \in I_u^*} 2 \ln \sigma(\lambda(r_{ugij} - r_{uqq'}) + (1 - \lambda)(r_{uqq'} - r_{upp'}))$ is the log-likelihood of GMPR, and $R(\Theta_G) = \sum_{u \in U} \sum_{t \in S} [\alpha_u \sum_{w \in G} \|U_w\|^2 + \alpha_v \|V_t\|^2 + \beta_v \|b_t\|^2]$ is the regularization term.

3 EXPERIMENTAL EVALUATION

3.1 Experimental Setup

3.1.1 Datasets. We employ five real-world datasets for experiments, including Tmall, MovieLens100K, MovieLens1M, UserTag and NF5K5K [8]. We take a pre-processing step mentioned in [5] to deal with the rating data. In summary, the statistics of all datasets are presented in Table 1.

Table 1: The statistics of datasets.

Statistics	Tmall	ML100K	ML1M	UserTag	NF5K5K
#users	28,059	943	6,040	3,000	5,000
#items	32,339	1,682	1,952	2,000	5,000
#feedbacks	464,426	100,000	1,000,209	246,436	282,474

Tmall dataset also has 1,024,575 click records. Here we conduct some experiments under various conditions of data sparseness to simulate realistic situations. For example, the case of “5%” denotes that we randomly select 5% of user-item pairs as training set and use the remainder for validation set and test set. For MPR with popularity, we rank items from the training set in descending order based on their observed counts and choose the last 20% as the unpopular set. The final experimental results are averaged over every evaluation metric on 20 copies of the test set.

3.1.2 Baselines and Evaluation Metrics. We compare the performance of MPR with several state-of-the-art implicit feedback recommendation methods: Weighted Matrix Factorization (WMF) [9], which is a typical pointwise method based on matrix factorization, BPR [4], and Group Bayesian Personalized Ranking (GBPR) [5], which is a state-of-the-art extension of BPR. We use “MPR (GMPR) + click” to represent dividing the unobserved items with click data on Tmall dataset and the other situations use popularity. Prec@5, MAP, NDCG@5, and AUC are employed to evaluate the recommendation performance of models. We set iteration number $T = 100,000$, learning rate $\gamma = 0.01$, regularization terms $\alpha_u = \alpha_v = \beta_v = 0.01$, the number of latent dimensions $d = 20$, and the tradeoff parameter $\lambda = 0.7$ by default in both MPR and GMPR. For GMPR, we fix the size of user group $G = 3$, the parameter $\rho = 1.0$ introduced in [5]. For other model parameters, the optimal values are tuned according to NDCG@5 on the validation set.

Table 2: Performance comparisons of MPR, GMPR and baselines. Numbers in boldface are the best results, and marking * indicates MPR or GMPR is superior to other algorithms significantly.

Dataset	Method	Prec@5	MAP	NDCG@5	AUC	Dataset	Method	Prec@5	MAP	NDCG@5	AUC
Tmall (25%)	WMF	0.0050	0.0039	0.0062	0.1809	Tmall (50%)	WMF	0.0075	0.0052	0.0079	0.2420
	BPR	0.0184	0.0133	0.0196	0.7130		BPR	0.0203	0.0142	0.0215	0.7355
	MPR+click	0.0221*	0.0146*	0.0225*	0.7494*		MPR+click	0.0226	0.0155	0.0238	0.7592
	MPR	0.0210	0.0143	0.0217	0.7429		MPR	0.0230*	0.0161*	0.0254*	0.7621*
	GBPR	0.0222	0.0147	0.0229	0.7501		GBPR	0.0239	0.0168	0.0257	0.7622
	GMPR+click	0.0229*	0.0160*	0.0242*	0.7608*		GMPR+click	0.0246	0.0172	0.0266	0.7694
	GMPR	0.0224	0.0154	0.0231	0.7534	GMPR	0.0251*	0.0177*	0.0279*	0.7773*	
ML100K (5%)	WMF	0.1042	0.0756	0.1173	0.2537	ML100K (10%)	WMF	0.1498	0.0789	0.1592	0.3960
	BPR	0.2100	0.1300	0.2149	0.7449		BPR	0.2458	0.1456	0.2522	0.7720
	MPR	0.2190*	0.1530*	0.2258*	0.7849*		MPR	0.2857*	0.1884*	0.2923*	0.8258*
	GBPR	0.2211	0.1389	0.2277	0.7650		GBPR	0.2622	0.1580	0.2698	0.7965
	GMPR	0.2297*	0.1587*	0.2353*	0.7900*	GMPR	0.3223*	0.2022*	0.3288*	0.8326*	
ML1M (5%)	WMF	0.1869	0.0667	0.1972	0.3755	ML1M (10%)	WMF	0.2373	0.0880	0.2469	0.5257
	BPR	0.2529	0.1318	0.2590	0.8069		BPR	0.3223	0.1707	0.3306	0.8498
	MPR	0.3063*	0.1743*	0.3087*	0.8583*		MPR	0.3659*	0.1979*	0.3723*	0.8784*
	GBPR	0.2739	0.1436	0.2801	0.8278		GBPR	0.3570	0.1916	0.3656	0.8657
	GMPR	0.3381*	0.1857*	0.3406*	0.8630*	GMPR	0.3670*	0.1982*	0.3740*	0.8785*	
UserTag (5%)	WMF	0.1968	0.0884	0.2026	0.2999	UserTag (10%)	WMF	0.1938	0.0812	0.1995	0.4282
	BPR	0.2011	0.1038	0.2046	0.6279		BPR	0.2024	0.1108	0.2063	0.6493
	MPR	0.2074*	0.1239*	0.2100*	0.6464*		MPR	0.2214*	0.1374*	0.2239*	0.6697*
	GBPR	0.2239	0.1147	0.2281	0.6422		GBPR	0.2296	0.1260	0.2345	0.6670
	GMPR	0.2285*	0.1349*	0.2306*	0.6562*	GMPR	0.2476*	0.1497*	0.2499*	0.6812*	
NF5K5K (5%)	WMF	0.0631	0.0246	0.0636	0.1506	NF5K5K (10%)	WMF	0.0920	0.0363	0.0945	0.2685
	BPR	0.1509	0.1056	0.1553	0.8670		BPR	0.1656	0.1163	0.1707	0.8802
	MPR	0.1735*	0.1270*	0.1766*	0.8868*		MPR	0.1831*	0.1362*	0.1877*	0.9007*
	GBPR	0.1610	0.1125	0.1649	0.8770		GBPR	0.1812	0.1274	0.1865	0.8913
	GMPR	0.1844*	0.1322*	0.1887*	0.8880*	GMPR	0.1967*	0.1435*	0.2015*	0.9029*	

3.2 Experimental Results

The experimental results are shown in Table 2, and Wilcoxon rank-sum statistical tests have been used to check whether the difference between our algorithms and other baselines are statistically significant (with a 0.05 significance level). From the table, we can see that MPR and GMPR significantly outperform other baselines on all evaluation metrics on all datasets, which shows that our proposed algorithm can recommend a more accurate rank biased list for users, especially for the top- k recommendation. MPR and GMPR shows significant improvement compared with other baselines on datasets for scenarios under 5% and 10%, which indicates that our algorithms can improve the quality of recommendation in situations of data sparsity. GMPR further improves GBPR and MPR on all datasets, which proves that our proposed algorithm can be used as a basic method and has extensive applicability. MPR (GMPR) with additional click data performs better than MPR (GMPR) using popularity on Tmall for scenarios under 25%, which indicates that click data may provide more information in situations of data sparsity.

4 CONCLUSION

In summary, we proposed a new pairwise ranking model, namely Multiple Pairwise Ranking (MPR), for item recommendation using implicit users' feedbacks. The main contribution of our approach is to dig implied positive feedbacks in the unobserved items based on the new multiple pairwise model. Empirical studies verified the effectiveness of our methods on real-world datasets. The multiple pairwise model is a new pairwise thinking that helps us under-

stand the preference difference among pairs of items and is not limited to the assumption in this paper. We encourage more application-oriented preference assumptions to be proposed based on our model.

ACKNOWLEDGMENTS

This research was partially supported by grants from the National Key Research and Development Program of China (No. 2016YFB1000904), and the National Natural Science Foundation of China (No.s 61672483, U1605251, and 61602147), and CCF-Tencent Open Fund, and the Young Elite Scientist Sponsorship Program of CAST, and the Youth Innovation Promotion Association of CAS (No. 2014299).

REFERENCES

- [1] Qi Liu, Enhong Chen, Hui Xiong, et al. Enhancing collaborative filtering by user interest expansion via personalized ranking. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 42.1 (2012): 218-233.
- [2] Qi Liu, Han Wu, Yuyang Ye, et al. Patent Litigation Prediction: A Convolutional Tensor Factorization Approach. In *IJCAI* 5052-5059.
- [3] Weike Pan and Li Chen. 2016. Group Bayesian personalized ranking with rich interactions for one-class collaborative filtering. *Neurocomputing* 207 (2016): 501-510.
- [4] Steffen Rendle, et al. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *UAI* 452-461.
- [5] Weike Pan and Li Chen. 2013. GBPR: Group Preference Based Bayesian Personalized Ranking for One-Class Collaborative Filtering. In *IJCAI* 2691-2697.
- [6] Xu Chen, Pengfei Wang, Zheng Qin, and Yongfeng Zhang. 2016. Hlbp: A hybrid local bayesian personal ranking method. In *WWW* 21-22.
- [7] Liang Du, Xuan Li, and Yi-Dong Shen. 2011. User graph regularized pairwise matrix factorization for item recommendation. In *ICADMA*. 372-385.
- [8] Rong Pan, Yunhong Zhou, Bin Cao, et al. 2008. One-class collaborative filtering. In *ICDM*. 502-511.
- [9] Yifan Hu, Yehuda Koren, and Chris Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *ICDM*. 263-272.