RBPR: Role-based Bayesian Personalized Ranking for Heterogeneous One-Class Collaborative Filtering

Xiaogang Peng, Yaofeng Chen, Yuchao Duan, Weike Pan* and Zhong Ming*
College of Computer Science and Software Engineering, Shenzhen University
patrickpeng@126.com, {chenyaofeng,duanyuchao}@email.szu.edu.cn, {panweike,mingz}@szu.edu.cn
*: corresponding author.

ABSTRACT

Heterogeneous one-class collaborative filtering (HOCCF) is a recently studied important recommendation problem, which consists of different types of users' one-class feedback such as browses and purchases. In HOCCF, we aim to fully exploit the heterogenous feedback and learn users' preferences so as to make a personalized and ranking-oriented recommendation for each user. For HOCCF, we can apply existing solutions for OCCF with purchases only such as Bayesian personalized ranking (BPR) or make use of both browses and purchases such as transfer via joint similarity learning (TJSL). However, BPR may be not very accurate due to the ignorance of browses, and TJSL may be not very efficient due to the mechanism of joint similarity learning and base model aggregation. In this paper, we propose a novel perspective for the different types of one-class feedback via users' different *roles*, i.e., browser and purchaser. Specifically, we design a two-stage role-based preference learning framework, i.e., role-based Bayesian personalized ranking (RBPR). In RBPR, we first digest the combined one-class feedback *as a browser* to find the candidate items that a user will browse, and then we exploit the purchase feedback to refine the candidate list *as a purchaser*. Empirical results on five public datasets show that our RBPR is an efficient and accurate recommendation algorithm for HOCCF as compared with the state-of-the-art methods such as BPR and TJSL.

CCS Concepts

•Information systems → Personalization; •Human-centered computing → Collaborative filtering;

Keywords

Role-based Preference Learning; Bayesian Personalized Ranking; Heterogeneous One-Class Collaborative Filtering

1. INTRODUCTION

.

Intelligent recommendation systems and technology have played a more and more important role in various real-world applications, with a wide spectrum of entertainment, social and professional services. Some recent work show that one important line of research have gradually transferred from collaborative filtering (CF) with numerical ratings to one-class CF (OCCF) with homogeneous one-class feedback such as purchases [2] and heterogeneous OCCF (HOCCF) with more than one types of one-class feedback such as browses and purchases [3]. In this paper, we focus on the problem setting of HOCCF, which is very common in real industry scenarios.

The main challenge of HOCCF is the heterogeneity of the two different types of one-class feedback, since a user's preference behind a purchase action may be different from that of a browse action. In a very recent work [3], a similarity learning algorithm is proposed for this challenge, which aims to combine browses and purchases in a principled way. The improved performance in [3] shows the complementarity of browses to the well exploited feedback of purchases in OCCF models [1, 4]. However, the proposed algorithm, i.e., *transfer via joint similarity learning* (TJSL) [3], may be not efficient enough for large datasets due to the complex prediction rule and base model ensemble.

In this paper, we interpret the HOCCF problem from a novel view of users' roles, i.e., *a purchaser (as reflected in a purchase feedback) is converted from a browser in a sequential manner*. Based on this perspective, we propose a two-stage framework, including browser-based preference learning and purchaser-based preference learning. Those two preference learning tasks are connected via a candidate list of items that a user will browse, which is assumed to contain the potential items that a user will finally purchase. In each of the two tasks, we apply the seminal work for homogeneous one-class feedback, i.e., Bayesian personalized ranking (BPR) [4], and for this reason, we call our approach role-based BPR (RBPR).

In our empirical studies, we compare our RBPR with the state-of-the-art methods of BPR and TJSL using various ranking-oriented evaluation metrics on five public datasets. The studies show that our RBPR is able to produce competitive recommendations efficiently. We list our main contributions as follows: (i) we propose a novel and generic staged role-based preference learning framework, which is a frustratingly easy, scalable and effective solution for collaborative ranking with heterogeneous one-class feedback; and (ii) we conduct extensive empirical studies and obtain very promising results.

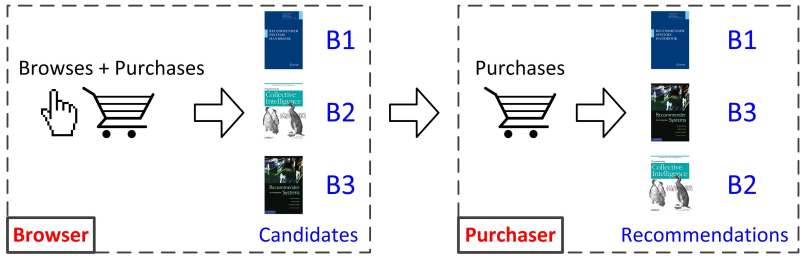

Figure 1: An illustration of role-based preference learning for heterogeneous one-class collaborative filtering (HOCCF), including browser-based preference learning and purchaser-based preference learning.

2. ROLE-BASED BAYESIAN PERSONALIZED RANKING

2.1 Problem Definition

In HOCCF, we have a set of n users ($\mathcal{U}$), a set of m items ($\mathcal{I}$), and two different sets of user feedback, e.g., browses $\mathcal{B}$ and purchases $\mathcal{P}$. Our goal is to find some likely-to-purchase items from unpurchased items for each user.

In order to fully exploit heterogeneous feedback in HOCCF such as browses and purchases, we propose not to model those different feedback jointly as a whole as done in a recent work [3], but separately in a staged manner. Specifically, we model different feedback of a typical user via different roles such as browser and purchaser. From the perspective of browser and purchaser, in our role-based Bayesian personalized ranking (RBPR), we have two tasks of preference learning, including browser-based preference learning and purchaser-based preference learning. We illustrate the main procedure of our proposed solution in Figure 1.

2.2 Browser-based Preference Learning

In the first step, we assume that a typical user is first a browser before he/she is converted to a purchaser. And thus, in our first task, we focus on answering the question of *"whether a user will browse an item"*.

In order to address this task, we propose to combine the two types of one-class feedback, i.e., browses and purchases, together, and then apply an algorithm for homogeneous one-class feedback such as BPR [4], i.e., BPR($\mathcal{B} \cup \mathcal{P}$). Mathematically, we will solve the following optimization problem,

$$\min_{\Theta_{\mathcal{B}\cup\mathcal{P}}} \sum_{u\in\mathcal{U}} \sum_{i\in(\mathcal{B}_u\cup\mathcal{P}_u)} \sum_{j\in\mathcal{I}\setminus(\mathcal{B}_u\cup\mathcal{P}_u)} f_{uij}, \quad (1)$$

where $\mathcal{B}_u$ and $\mathcal{P}_u$ are item sets browsed and purchased by user u, respectively, f_{uij} is the tentative objective function for a randomly sampled triple (u, i, j), and $\Theta_{\mathcal{B}\cup\mathcal{P}}$ denotes the set of model parameters to be learned [4].

Once we have learned the model parameters, we can generate a candidate list of items that a user is likely to browse. Specifically, for a top-K recommended problem, we will generate $3K$ items in this step, so that the refinement in next step may have more room for improvement.

2.3 Purchaser-based Preference Learning

In the second step, we assume that a user will most likely choose an item from the candidate list that he/she has browsed. For this reason, in our second task, we mainly answer the question of *"whether a user will purchase an item"*.

Input: Users' browses $\mathcal{B}$ and purchases $\mathcal{P}$.

Output: Top-K recommended items for each user.

Step 1. Conduct browser-based preference learning via BPR($\mathcal{B}\cup\mathcal{P}$) as shown in Eq.(1) and obtain $3K$ candidate items with highest predicted scores.
Step 2. Conduct purchaser-based preference learning via BPR($\mathcal{P}$) as shown in Eq.(2); predict the scores on the $3K$ candidate items and refine the list.

Figure 2: The algorithm of role-based Bayesian personalized ranking (RBPR).

In order to solve this task, we propose to use the purchase data only to refine the candidate list from the first step. The reason is that the purchase feedback is more helpful in answering whether a certain item will be bought by a user. Due to the fact that a user's purchase feedback are few, we may not get good results if we only apply the second step, i.e., only use the purchase feedback to find items that will be bought by a user. This phenomenon is also observed in our empirical studies.

Similarly, we again adopt BPR [4] for model training, but use purchase feedback $\mathcal{P}$ only. Mathematically, we learn the model parameters as follows,

$$\min_{\Theta_{\mathcal{P}}} \sum_{u\in\mathcal{U}} \sum_{i\in\mathcal{P}_u} \sum_{j\in\mathcal{I}\setminus\mathcal{P}_u} f_{uij}, \quad (2)$$

where $\Theta_{\mathcal{P}}$ denotes the model parameters to be learned from the purchase data only.

With the learned model parameters $\Theta_{\mathcal{P}}$, we can predict the preference of each item i in the candidate list of each user u, and then re-rank the items in the list. The refined list is expected to better represent the purchase likelihood of a certain user, i.e., the recommendation may be more accurate, which is also verified in our empirical studies. We illustrate the effect of the difference between those two ranked lists in Figure 1.

For the optimization problems in the aforementioned two learning tasks, we can apply stochastic gradient descent to learn the model parameters [4]. We put the two preference learning tasks in one single algorithm in Figure 2 in order to get a complete picture.

Table 1: Description of the datasets used in the experiments, including the numbers of users, items, purchases and browses, and the number of purchases in validation and test data. Note that the data of ML100K, ML1M and Alibaba are from [3], and the statistics of ML10M and Netflix are for the first copy of the three generated copies of each dataset.

Dataset	# user	# item	# purchase	# browse	# purchase (validation)	# purchase (test)
ML100K	943	1682	9438	45285	–	2153
ML1M	6040	3952	90848	400083	–	45075
Alibaba2015	7475	5257	9290	60659	–	2322
ML10M	71567	10681	309317	4000024	308673	308702
Netflix	480189	17770	4554888	39628846	4556347	4558506

3. EXPERIMENTAL RESULTS

3.1 Datasets and Evaluation Metrics

In our empirical studies, in order to directly compare our RBPR with the very recent algorithm for HOCCF, i.e., TJSL [3]. We first use the three public datasets in [3][1], including MovieLens 100K (ML100K), MovieLens 1M (ML1M) and Alibaba2015. The detailed description of those three data can be found in [3]. We also study the performance of our RBPR on two large datasets, including MovieLens 10M (ML10M)[2] and Netflix.

ML10M is a public data with about 10 million numerical ratings in $\{0.5, 1, 1.5, ..., 4.5, 5\}$, and Netflix is the dataset used in the famous $100 Million competition with about 0.1 billion scores in $\{1, 2, 3, 4, 5\}$. For both ML10M and Netflix, we first divide the data into five parts with equal numbers of (u, i, r_{ui}) triples, we then take one part and keep the (u, i) pairs with $r_{ui} = 5$ as purchases for training, take one part and keep the (u, i) pairs with $r_{ui} = 5$ as purchases for validation, and take one part and keep the (u, i) pairs with $r_{ui} = 5$ as purchases for test, and finally take the remaining two parts and keep all the (u, i) pairs as browses. We repeat this procedure for three times in order to obtain three copies of data.

We put the statistics of the datasets in Table 1.

For evaluation, we use five ranking-oriented metrics, including Precision@5, Recall@5, F1@5, NDCG@5 and 1-call@5.

3.2 Baselines and Parameter Settings

Because HOCCF is a relatively new recommendation problem, very few solutions have been proposed. In our empirical studies, we thus include the very recent algorithm TJSL [3] for HOCCF and also the seminal work BPR [4] for OCCF.

- BPR (Bayesian personalized ranking) is an efficient and accurate recommendation algorithm for homogeneous one-class feedback such as purchases, which mines users' preferences by assuming that a user prefers a purchased item to an unpurchased item.

- TJSL (transfer via joint similarity learning) is the state-of-the-art method for heterogeneous one-class feedback such as browses and purchases, which jointly learns the similarity between a candidate item and a purchased item, and the similarity between a candidate item and a likely-to-purchase item.

[1] http://www.cse.ust.hk/∼weikep/TL4HOCCF/
[2] http://grouplens.org/datasets/movielens/10m/

For BPR, TJSL and RBPR, we fix the dimension as $d = 20$ and the learning rate as $\gamma = 0.01$. For BPR and TJSL on ML100K, ML1M and Alibaba2015, we directly use the results from [3]. For RBPR on all the datasets and BPR on ML10M and Netflix, we search the best tradeoff parameter from $\{0.001, 0.01, 0.1\}$ and iteration number from $\{100, 500, 1000\}$ via NDCG@15. In order to make the results consistent and comparable with [3], we run five times of RBPR on ML100K, ML1M and Alibaba2015, and report the averaged performance. For ML10M and Netflix, we report the averaged results on three copies of data.

3.3 Results

We report the recommendation performance in Table 2. We can have the following observations:

- RBPR and TJSL are better than BPR in all cases including five evaluation metrics and five datasets, which clearly shows that the feedback browses are useful for learning and mining users' hidden preferences, and RBPR and TJSL are able to make use of users' heterogeneous feedback well.

- RBPR and TJSL are comparable on three small datasets, e.g., TJSL is the best on ML100K, RBPR is the best on ML1M, and TJSL and RBPR are comparable on Alibaba2015.

- TJSL is too slow to generate recommendations on two large datasets within 24 hours, while RBPR can produce significantly better results than BPR, which shows that our RBPR is a more practical solution regarding the efficiency.

The overall performance in Table 2 shows that our RBPR performs the best in making use of the heterogeneous one-class feedback.

In order to check the performance improvement of our two-stage role-based preference learning solution, we also check the performance of the generated candidate items as shown in Figure 1. Specifically, we denote the method for generating the candidates as RBPR(Browser) since it is based on the role of browser only, and the final recommendation as RBPR(Browser,Purchaser). We report the performance on Precision and NDCG in Figure 3 (other metrics are similar), from which we can see that the second stage of candidate refinement using the purchase data can significantly improve the performance. The improvement also verifies our main assumption that there are usually two separate stages for a user's shopping action, i.e., browse and purchase.

Table 2: Recommendation performance of RBPR, BPR and TJSL on ML100K, ML1M, Alibaba2015, ML10M and Netflix using Prec@5, Rec@5, F1@5, NDCG@5 and 1-call@5. The significantly best results are marked in bold (p value < 0.01). Note that the results of BPR and TJSL on three small datasets are from [3]. We use "–" to denote the case that the training procedure does not finish within 24 hours.

Dataset	Method	Prec@5	Rec@5	F1@5	NDCG@5	1-call@5
ML100K	BPR	$0.0552_{\pm 0.0006}$	$0.1032_{\pm 0.0019}$	$0.0673_{\pm 0.0007}$	$0.0874_{\pm 0.0020}$	$0.2425_{\pm 0.0034}$
	TJSL	$\mathbf{0.0697}_{\pm \mathbf{0.0016}}$	$\mathbf{0.1393}_{\pm \mathbf{0.0028}}$	$\mathbf{0.0864}_{\pm \mathbf{0.0019}}$	$\mathbf{0.1133}_{\pm \mathbf{0.0047}}$	$\mathbf{0.3033}_{\pm \mathbf{0.0071}}$
	RBPR	$0.0654_{\pm 0.0013}$	$0.1275_{\pm 0.0048}$	$0.0803_{\pm 0.0021}$	$0.1058_{\pm 0.0047}$	$0.2890_{\pm 0.0047}$
ML1M	BPR	$0.0928_{\pm 0.0008}$	$0.0829_{\pm 0.0002}$	$0.0717_{\pm 0.0003}$	$0.1121_{\pm 0.0010}$	$0.3609_{\pm 0.0018}$
	TJSL	$0.1012_{\pm 0.0011}$	$0.0968_{\pm 0.0012}$	$0.0821_{\pm 0.0009}$	$0.1248_{\pm 0.0010}$	$0.3961_{\pm 0.0022}$
	RBPR	$\mathbf{0.1086}_{\pm \mathbf{0.0009}}$	$\mathbf{0.1017}_{\pm \mathbf{0.0015}}$	$\mathbf{0.0858}_{\pm \mathbf{0.0009}}$	$\mathbf{0.1327}_{\pm \mathbf{0.0016}}$	$\mathbf{0.4151}_{\pm \mathbf{0.0055}}$
Alibaba2015	BPR	$0.0050_{\pm 0.0006}$	$0.0193_{\pm 0.0026}$	$0.0077_{\pm 0.0009}$	$0.0138_{\pm 0.0017}$	$0.0246_{\pm 0.0031}$
	TJSL	$0.0071_{\pm 0.0004}$	$0.0283_{\pm 0.0016}$	$0.0110_{\pm 0.0006}$	$0.0200_{\pm 0.0008}$	$0.0347_{\pm 0.0017}$
	RBPR	$\mathbf{0.0076}_{\pm \mathbf{0.0005}}$	$0.0304_{\pm 0.0023}$	$0.0118_{\pm 0.0008}$	$\mathbf{0.0220}_{\pm \mathbf{0.0013}}$	$0.0367_{\pm 0.0024}$
ML10M	BPR	$0.0629_{\pm 0.0002}$	$0.0855_{\pm 0.0006}$	$0.0603_{\pm 0.0003}$	$0.0861_{\pm 0.0004}$	$0.2648_{\pm 0.0017}$
	TJSL	–	–	–	–	–
	RBPR	$\mathbf{0.0719}_{\pm \mathbf{0.0013}}$	$\mathbf{0.0977}_{\pm \mathbf{0.0017}}$	$\mathbf{0.0690}_{\pm \mathbf{0.0014}}$	$\mathbf{0.0994}_{\pm \mathbf{0.0020}}$	$\mathbf{0.2990}_{\pm \mathbf{0.0050}}$
Netflix	BPR	$0.0716_{\pm 0.0007}$	$0.0480_{\pm 0.0005}$	$0.0446_{\pm 0.0005}$	$0.0818_{\pm 0.0011}$	$0.2846_{\pm 0.0022}$
	TJSL	–	–	–	–	–
	RBPR	$\mathbf{0.0797}_{\pm \mathbf{0.0002}}$	$\mathbf{0.0595}_{\pm \mathbf{0.0004}}$	$\mathbf{0.0527}_{\pm \mathbf{0.0003}}$	$\mathbf{0.0939}_{\pm \mathbf{0.0003}}$	$\mathbf{0.3174}_{\pm \mathbf{0.0011}}$

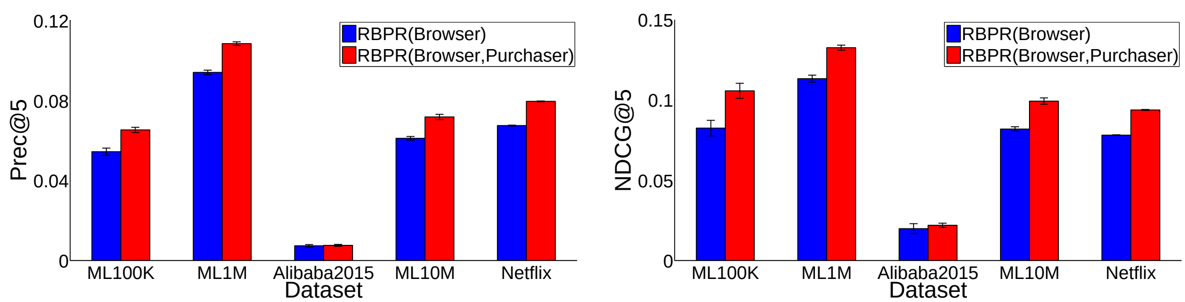

Figure 3: Recommendation performance of RBPR with different configurations, i.e., RBPR(Browser) for browser only, and BPR(Browser,Purchaser) for both browser and purchaser.

4. CONCLUSIONS AND FUTURE WORK

In this paper, we study an important recommendation problem called heterogeneous one-class collaborative filtering (HOCCF) from a novel perspective of users' roles. Specifically, we propose a novel role-based preference learning framework, i.e., role-based Bayesian personalized ranking (RBPR), based on a seminal work [4]. Extensive empirical studies show that our RBPR is more accurate than the seminal work for OCCF, i.e., BPR [4], and a very recent similarity learning algorithm for HOCCF, i.e., TJSL [3]. Furthermore, our RBPR is very efficient with the inherited merits of BPR [4], while TJSL [3] is difficult to produce recommendations on two large datasets.

For future work, we are interested in extending and applying our role-based preference learning framework to other recommendation settings with more types of users' roles such as searcher, browser, purchaser, rater and friends.

5. ACKNOWLEDGMENT

We thank the support of Natural Science Foundation of China (NSFC) No. 61502307 and Natural Science Foundation of Guangdong Province No. 2014A030310268 and No. 2016A030313038.

6. REFERENCES

[1] S. Kabbur, X. Ning, and G. Karypis. Fism: Factored item similarity models for top-n recommender systems. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '13, pages 659–667, 2013.

[2] R. Pan, Y. Zhou, B. Cao, N. N. Liu, R. Lukose, M. Scholz, and Q. Yang. One-class collaborative filtering. In *Proceedings of the 8th IEEE International Conference on Data Mining*, ICDM '08, pages 502–511, 2008.

[3] W. Pan, M. Liu, and Z. Ming. Transfer learning for heterogeneous one-class collaborative filtering. *IEEE Intelligent Systems*, 2016.

[4] S. Rendle, C. Freudenthaler, Z. Gantner, and L. 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, 2009.