Neurocomputing 207 (2016) 501–510

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Group Bayesian personalized ranking with rich interactions for one-class collaborative filtering[☆,☆☆]

Weike Pan[a,b,*], Li Chen[b]

[a] College of Computer Science and Software Engineering, Shenzhen University, China
[b] Department of Computer Science, Hong Kong Baptist University, Hong Kong

A R T I C L E   I N F O

Article history:
Received 29 July 2015
Received in revised form
5 May 2016
Accepted 9 May 2016
Communicated by Peng Cui
Available online 16 May 2016

Keywords:
One-class collaborative filtering
Implicit feedback
Group pairwise preference
Item set

A B S T R A C T

Both researchers and practitioners in the field of collaborative filtering have shown keen interest to user behaviors of the "one-class" feedback form such as transactions in e-commerce and "likes" in social networks. This recommendation problem is termed as one-class collaborative filtering (OCCF). In most of the previous work, a pairwise preference assumption called Bayesian personalized ranking (BPR) was empirically proved to be able to exploit such one-class data well. In one of the most recent work, an upgraded model called group preference based BPR (GBPR) leverages the group preference and obtains better performance.

In this paper, we go one step beyond GBPR, and propose a new and generic assumption, i.e., group Bayesian personalized ranking with rich interactions (GBPR$^+$). In our GBPR$^+$, we adopt a set of items instead of one single item as used in GBPR, which is expected to introduce rich interactions. GBPR is a special case of our GPBR$^+$ when the item set contains only one single item. We study the empirical performance of our GBPR$^+$ with several state-of-the-art methods on four real-world datasets, and find that our GPBR$^+$ can generate more accurate recommendations.

1. Introduction

In industry, recommender system as a critical engine in various online entertainments [10] and shopping services [24] has caught much attention and contributed significantly in revenue growth in recent years. Lots of internet, electronic and telecom giants embed recommendation technologies in their existing systems, e.g., YouTube[1] and Amazon,[2] in order to increase user engagement and earn more revenues from sales of product or advertisement.

In academia, most research studies [2,6] on recommendation are biased towards the numerical rating prediction problem associated with the Netflix competition,[3] partially due to the public availability of the data. The Netflix contest can be categorized as a "multi-class" recommendation problem, where the inputs are categorical scores or classes. For example, the input data of Netflix are numerical ratings with "1" for bad, "2" for fair, "3" for good, "4" for excellent, and "5" for perfect, for which various algorithms have been proposed in order to predict the users' preference scores accurately. Such categorical data contain both positive feedback of "4" and "5" and negative feedback of "1" and "2". For such data, some regression or ranking loss functions were designed to fit the scores or to preserve the ordering. Matrix-factorization based models have been shown to be the most effective solutions [20,21,39,48], which makes use of these loss functions. There are also some work that exploits both multi-class ratings and additional information [33,55], such as social connections, interests and behaviors [16,25,38,53], content [1,3], taxonomy [29] and context [44,49], and adopts more than one recommendation techniques [14,17,26].

However, in most applications, user behavior data are in "one-class" form rather than in multi-class form, e.g., "like" in Facebook, "bought" in Amazon, and "click" in Google Advertisement. Such data are called implicit [28,41] or one-class [32] feedback. We illustrate the problem in Fig. 1, where some users have expressed positive feedback on some movies. This kind of one-class collaborative filtering problem is different from that of the 5-star rating prediction problem encountered in the Netflix competition, because the former only contains positive feedback compared to both positive feedback and negative feedback used in the

[☆] This work is an extension of our previous work [35]. Compared with our previous work [35], we have added the following new contents in this paper, (1) we have developed an extension of GBPR (i.e., GBPR$^+$) in Section 4; (2) we have included new experimental results (i.e., five additional state-of-the-art baselines in Table 4, more detailed results in Figs 5 and 6, and new results in Fig. 7) and associated analyses in Section 5; (3) we have added more related work and discussions in Sections 1 and 2; and (4) we have made many improvements on the Abstract, Introduction, illustrations in Figs 2 and 3, and result presentation.

[☆☆] Some of this work was done while Weike Pan was a post-doctoral research fellow in the Department of Computer Science, Hong Kong Baptist University.

* Corresponding author.
E-mail addresses: panweike@szu.edu.cn (W. Pan), lichen@comp.hkbu.edu.hk (L. Chen).

[1] http://www.youtube.com/.
[2] http://www.amazon.com/.
[3] http://www.netflixprize.com/.

http://dx.doi.org/10.1016/j.neucom.2016.05.019
0925-2312/

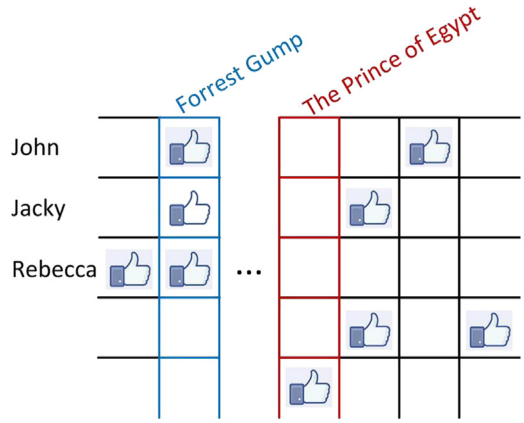

Fig. 1. An illustrative example of the one-class collaborative filtering (OCCF) problem. The observed data only contains positive feedback (e.g., "like") from users on items.

competition. Furthermore, the goal of OCCF is to rank items rather than predict ratings.

For solving the one-class collaborative filtering problem, previous matrix-factorization based algorithms can be roughly summarized into two categories, (1) pointwise regression methods and (2) pairwise ranking methods. The former methods learn a latent representation of both users and items by minimizing a pointwise square loss [13,32] in order to fit the absolute rating scores, while the latter methods take pairs of items as basic units and maximize the likelihood of pairwise preference over observed items and unobserved items [41]. Empirically, the pairwise ranking method [41] achieves much better performance than pointwise methods [13,32], and has been successfully adopted in many applications, e.g., tag recommendation [43], news recommendation [50], and shopping products recommendation [19]. Notice that there are also some work that generalize pairwise ranking to listwise ranking [45,46], though it is of low efficiency.

In this paper, we study the two fundamental assumptions made in the seminal work of pairwise ranking method [41], i.e., (1) *individual pairwise preference over two items* and (2) *independence between two users*, and point out that they may not always hold in real applications. As a response, we introduce *group pairwise preference* instead of *individual pairwise preference* in [41], in order to inject rich interactions among users and thus relax the *individual* and *independence* assumptions. We then propose a new and improved assumption called *group Bayesian personalized ranking with rich interactions* (GBPR^+) and design an efficient algorithm correspondingly. Notice that GBPR^+ is extended from our previous conference work GBPR [35]. Empirically, we find that our new assumption can make use of the one-class data more effectively and achieves better recommendation performance on all the four real-world datasets in our experiments.

We organize the rest of the paper as follows. We discuss some closely related work on one-class collaborative filtering in Section 2. We present the background and limitations of the existing pairwise preference learning methods in Section 3. We formally propose our assumption and algorithm in Section 4, and then study its empirical performance in Section 5. Finally, we give some concluding remarks in Section 6.

2. Related work

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

Pointwise methods with absolute preference assumptions: Pointwise methods take implicit feedback as *absolute* preference scores. For example, an observed user–item pair, (u,i), is interpreted as that user u likes item i with a high absolute score, e.g., weighted matrix factorization (WMF) [31,32], pure singular value decomposition (PureSVD) [9], sparse linear models (SLIM) [30], factored item similarity model with RMSE (FISM_{rmse}) [18] and implicit matrix factorization (IMF) [13] are typical approaches for solving this recommendation problem. WMF [32] uses different sampling strategies for unobserved user–item pairs and takes them as negative feedback to augment the observed positive feedback, so that existing matrix factorization methods can be applied. PureSVD [9] bypasses the sampling step in WMF [32] and takes all unobserved user–item pairs as negative feedback, which are then combined with observed pairs and fed to the mathematical tool of singular value decomposition (SVD) [5,11]. SLIM [30] and FISM_{rmse} [18] learn similarities between items instead of calculate them via some predefined measurement such as cosine similarity, where the former is associated with ℓ_1 regularization and the latter with factored latent features. IMF [13] introduces confidence weights on implicit feedback, which can then be approximated by two latent feature matrices. However, the limitation of WMF [32], PureSVD [9], SLIM [30] and FISM_{rmse} [18] is that taking unobserved user–item pairs as negative feedback may introduce errors. As for IMF [13], it requires auxiliary knowledge of confidence for each observed feedback, which may not be available in real applications.

There are also some non-negative matrix factorization (NMF) [23] based pointwise methods for one-class collaborative filtering, e.g., NMF with a low density assumption [47], NMF with side information of social trust and item content [51], etc.

Pairwise methods with relative preference assumptions: Pairwise methods take implicit feedback as *relative* preference rather than absolute ones, e.g., a user u is assumed to prefer an item i to an item j if the user–item pair (u,i) is observed, and (u,j) is not observed [41]. The proposed algorithm, Bayesian personalized ranking (BPR) [41], is the first method with such pairwise preference assumption for addressing the one-class collaborative filtering problem.

Due to the great success of pairwise methods in various one-class collaborative filtering problems, some new algorithms have been proposed to combine BPR with some auxiliary data, such as BPR with temporal information [42], BPR with user-side social connections [54], and BPR with item-side taxonomy [19], etc. There are also some work that (1) extend BPR from two dimensions to three dimensions [43], from one user–item matrix to multiple ones [22,37], (2) associate some sophisticated sampling strategies with BPR [40,52], or (3) revise the loss function [27,45,46] in BPR.

However, the limitation of pairwise methods can be attributed to the two fundamental assumptions made in BPR, namely *individual pairwise assumption over two items* and *independence assumption between two users*. Most follow-up work do not refine the fundamental assumptions, but just directly adopt the BPR criterion in their own applications. A recent algorithm [34] generalizes BPR via proposing a new assumption that an individual user is likely to prefer a set of observed items to a set of unobserved items.

Compared with the aforementioned work, our proposed *group Bayesian personalized ranking with rich interactions* (GBPR^+) is a novel algorithm in one-class collaborative filtering. In particular, GBPR^+ inherits the merit of pairwise methods, and improves the two fundamental assumptions in BPR via introducing *group pairwise preference*. Notice that GBPR^+ is a more generic version of our previous conference work GBPR [35], and it becomes equivalent to GBPR when the item set is of one single item. We summarize GBPR^+ and the aforementioned related work in Table 1.