An Improved Sampler for Bayesian Personalized Ranking by Leveraging View Data*

Jingtao Ding¹, Fuli Feng², Xiangnan He², Guanghui Yu¹, Yong Li¹, Depeng Jin¹

¹Department of Electronic Engineering, Tsinghua University

²School of Computing, National University of Singapore

liyong07@tsinghua.edu.cn;xiangnanhe@gmail.com

ABSTRACT

Bayesian Personalized Ranking (BPR) is a representative pairwise learning method for optimizing recommendation models. It is widely known that the performance of BPR depends largely on the quality of the negative sampler. In this short paper, we make two contributions with respect to BPR. First, we find that sampling negative items from the whole space is unnecessary and may even degrade the performance. Second, focusing on the purchase feedback of the E-commerce domain, we propose a simple yet effective sampler for BPR by leveraging the additional view data. Compared to the vanilla BPR that applies a uniform sampler on all candidates, our view-aware sampler enhances BPR with a relative improvement of 27.36% and 69.54% on two real-world datasets respectively.

KEYWORDS

BPR; recommendation; sampler; view data.

ACM Reference Format:

Jingtao Ding, Fuli Feng, Xiangnan He, Guanghui Yu, Yong Li, Depeng Jin. 2018. An Improved Sampler for Bayesian Personalized Ranking by Leveraging View Data. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3184558.3186905>

1 INTRODUCTION

Due to the prevalence of user implicit feedback in online information systems, recent research on recommendation has shifted from explicit ratings to implicit feedback, such as purchases, clicks, watches and so on [1, 3]. To learn recommender models from binary implicit feedback, Rendle et al. [5] proposed the BPR method, which assumes that an observed interaction should be predicted with a higher score than its unobserved counterparts (i.e., the missing interactions). Mathematically, the objective function for BPR can be formulated as

$$\arg \min_{\Theta} \sum_{(u, i, j) \in \mathcal{D}} -\ln \sigma(\hat{y}_{ui}(\Theta) - \hat{y}_{uj}(\Theta)), \quad (1)$$

where $\hat{y}(\Theta)$ is the predictive model, Θ denotes the model parameters, $\sigma(x) = \frac{1}{1+\exp(-x)}$ is the sigmoid function to convert the margin to

*This work was supported in part by the National Nature Science Foundation of China under 61621091 and 61673237, and research fund of Tsinghua University - Tencent Joint Laboratory for Internet Innovation Technology.

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3186905>

a probability, and $\mathcal{D}$ denotes the set of pairwise training examples: $\{(u, i, j) | i \in \mathcal{R}_u^+ \wedge j \notin \mathcal{R}_u^+\}$, where $\mathcal{R}_u^+$ denotes the set of items that u has interacted with before. Note that we have omitted the L_2 regularization terms for clarity. The optimization of BPR is usually achieved by the stochastic gradient descent (SGD). In each step, it first randomly draws an observed interaction (u, i) , and then selects an item j that u has not interacted with before to constitute (u, i, j) . Such a process of selecting j is also known as *negative sampling*.

In the original paper of BPR [5], Rendle et al. applied a uniform negative sampler, i.e., sampling j from **all items** that u has not consumed before with an **equal probability**. Later on, it was reported that such a uniform negative sampler is highly ineffective and slows down the convergence of BPR [4, 6], especially for datasets that have a large number of items. To this end, [4, 6] proposed dynamic negative sampling (DNS) strategies, aiming to maximize the utility of a gradient step by choosing “difficult” negative examples — i.e., the negative examples that lead to a large prediction loss by the current model. Despite that significant improvements have been observed, existing DNS strategies sample negative items from the whole item space, which arguably may still suffer from low efficiency when the number of items is large.

In this work, we aim to answer the following two research questions: 1) Is it necessary to sample negative items from the whole space? and 2) Can we design a better sampler for BPR?

2 EXPERIMENTAL SETTINGS

Datasets. We perform experiments on two real-world datasets.

Beibei¹: Beibei is the largest E-commerce platform for maternal and infant products in China. We sampled a subset of user interactions that contain views and purchases from Beibei within the time period from 2017/05/25 to 2017/06/28.

Tmall²: Tmall is the largest business-to-consumer E-commerce platform in China. To allow our results to be reproducible, we used a public benchmark released by the ICJAI-2015³.

We took three steps for data preprocessing. We first merged the repetitive purchases into one purchase with the earliest timestamp, as we aim to recommend novel items. Next we filtered out users' views on their purchased items to avoid information leaking. Finally, we filtered out users and items with less than 12 and 16 purchases, respectively, to overcome the high sparsity of the raw datasets. Table 1 summarizes the statistics of our experiment datasets.

Evaluation Methodology. We adopted the *leave-one-out* protocol [3, 5], where the latest purchase interaction of each user is held out for testing. For hyperparameter tuning, we randomly sampled

¹<http://www.beibei.com/>

²<https://www.tmall.com/>

³The dataset is downloaded from <https://tianchi.aliyun.com/datalab/dataSet.htm?id=5>

Table 1: Statistics of the evaluation datasets.

Dataset	Purchase#	View#	User#	Item#	Sparsity
Beibei	2,654,467	46,912,880	158,907	119,012	99.99%/99.75%
Tmall	464,426	1,585,225	28,059	32,339	99.95%/99.83%

one purchase interaction for each user as the validation set. The training process was stopped once we observed increasing in the validation loss. We employed *Hit Ratio* (HR) [3] for each user by truncating the ranked list of non-purchased items at the position of 100 and reported the average score of all users. We used the standard matrix factorization [5] as the predictive model, where the number of latent factors equals to 32.

3 UNNECESSARY TO SAMPLE FROM ALL ITEMS

We controlled the sampling space of BPR on a user basis. For each user, the negative items are only sampled from a fraction of items, i.e., a randomly reduced item space. We varied the size and summarized the performance in Table 2. For each setting, we repeated the experiment five times and reported the average score. The first row indicates the performance of the original BPR that samples negative items from the whole space.

Table 2: Performance of BPR with different settings on the ratio of the reduced sampling space. “Size” means the number of items in sampling space for each user (Ratio $\times$ Item#).

Beibei				Tmall			
Ratio	Size	HR	Δ HR	Ratio	Size	HR	Δ HR
2^0	119,012	0.1094	0	2^0	32,339	0.0301	0
2^{-6}	1,859	0.1112	+2.03%	2^{-3}	4,042	0.0300	-0.27%
2^{-7}	930	0.1103	+1.22%	2^{-4}	2,021	0.0300	-0.33%
2^{-8}	465	0.1106	+1.50%	2^{-5}	1010	0.0297	-1.33%
2^{-9}	232	0.1104	+1.26%	2^{-6}	505	0.0299	-0.60%

Surprisingly on the Beibei dataset, the performance is not decreased but increased after reducing the sampling space. For example, when the sampling space is $1/2^9$ smaller, we obtained a relative improvement of 1.26% over the original BPR. This finding is novel and encouraging, meaning that sampling from the whole item space is not only unnecessary for BPR, but may even hurt the performance. On the Tmall dataset, as the original item space is not that large (which is one magnitude smaller), we did not observe improvements by reducing the sampling space. But still, we can see that with a much smaller sampling space, the performance remains the same level as the original BPR. This provides further evidence on the inefficiency of the uniform sampler for BPR.

4 A VIEW-ENHANCED SAMPLER

In E-commerce recommender systems, besides the purchase feedback that is directly related to optimizing the conversion rate, the view logs of users can be intuitively treated as an intermediate feedback between the purchased and missing interactions. Since a training example $(u, i, j) \in \mathcal{D}$ in BPR assumes that u prefers i over j , we can integrate the view signal by augmenting the training data. In our proposed view-enhanced sampler, we split the item space into three sets for each user u , namely $\mathcal{P}_u$, $\mathcal{V}_u$, and $\mathcal{R}_u$, which indicate the purchased items, viewed (but not purchased) items, and remaining items, respectively. Then, we sample an item pair (i, j) from three candidate sets, $\{(i, j) | i \in \mathcal{P}_u, j \in \mathcal{V}_u\}$, $\{(i, j) | i \in \mathcal{P}_u, j \in \mathcal{R}_u\}$, and $\{(i, j) | i \in \mathcal{V}_u, j \in \mathcal{R}_u\}$, with predefined probabilities $[\omega_1, \omega_2, \omega_3]$

respectively, where $\omega_1 + \omega_2 + \omega_3 = 1$. The generated training example (u, i, j) is finally used to update the model parameters in Eq. (1). We term the BPR method with this view-enhanced sampler as *BPR+view*.

Our proposed BPR+view achieved the best performance when $[\omega_1, \omega_2, \omega_3]$ are set as $[0.3, 0.3, 0.4]$ and $[0.01, 0.09, 0.9]$ on the Beibei and Tmall datasets, respectively. To demonstrate its effectiveness, we compare it with 1) the vanilla BPR [5], and 2) BPR-DNS [6], which selects the item with the highest prediction score among X randomly sampled negatives. For BPR-DNS, we tuned the X in the same way as the original paper. To our knowledge, DNS is the most effective sampler to date for BPR based on the interaction data only, and empirically outperforms [4]. In addition, we evaluated a common baseline Popularity [3], which simply recommends items based on their popularity evidenced by the number of purchases.

Figure 1: Performance comparison in each iteration.

Figure 1 shows the testing HR of the compared methods in each training iteration. As can be seen, upon convergence, BPR+view significantly outperforms all other methods, and the improvements are more significant on the Tmall dataset (60%+ relative improvements over BPR and BPR-DNS). This justifies the efficacy of accounting for the preference signal in the view data using our proposed sampler. Besides, we observed that Popularity performs as well as BPR on the Beibei dataset, which is unexpected since BPR is a personalized recommendation method. Our further investigation finds that the reason is because the Beibei dataset is highly popularity-skewed — the top-1% items contributed almost 50% of purchases.

5 FUTURE WORK

We have demonstrated that sampling negative items from the whole space is unnecessary for BPR, and proposed an enhanced sampler based on the view data. In future, we will design an adaptive sampler to leverage view data and other implicit feedback more sufficiently. Furthermore, we plan to explore more generic feature-based recommender models, such as the state-of-the-art neural factorization machine [2].

REFERENCES

- [1] I. Bayer, X. He, B. Kanagal, and S. Rendle. A generic coordinate descent framework for learning from implicit feedback. In *WWW*, pages 1341–1350, 2017.
- [2] X. He and T.-S. Chua. Neural factorization machines for sparse predictive analytics. In *SIGIR*, pages 355–364, 2017.
- [3] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua. Neural collaborative filtering. In *WWW*, pages 173–182, 2017.
- [4] S. Rendle and C. Freudenthaler. Improving pairwise learning for item recommendation from implicit feedback. In *WSDM*, pages 273–282, 2014.
- [5] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *UAI*, pages 452–461, 2009.
- [6] W. Zhang, T. Chen, J. Wang, and Y. Yu. Optimizing top-n collaborative filtering via dynamic negative item sampling. In *SIGIR*, pages 785–788, 2013.