Asymmetric Bayesian Personalized Ranking for One-Class Collaborative Filtering

Shan Ouyang, Lin Li, Weike Pan* and Zhong Ming*

College of Computer Science and Software Engineering, Shenzhen University

{ouyangshan, lilin20171}@email.szu.edu.cn, {panweike, mingz}@szu.edu.cn

Presented at the 13th ACM Conference on Recommender Systems (RecSys 2019)

Abstract

In this paper, we propose a novel preference assumption for modeling users' one-class feedback such as "thumb up" in an important recommendation problem called one-class collaborative filtering (OCCF). Specifically, we address a fundamental limitation of a recent symmetric pairwise preference assumption and propose a novel and first asymmetric one, which is able to make the preferences of different users more comparable. With the proposed asymmetric pairwise preference assumption, we further design a novel recommendation algorithm called asymmetric Bayesian personalized ranking (ABPR). Extensive empirical studies on two large and public datasets show that our ABPR performs significantly better than several state-of-the-art recommendation methods with either pointwise preference assumption or pairwise preference assumption.

Problem Definition of One-Class Collaborative Filtering

- **Input:** A set of (user, item) pairs $\mathcal{R} = \{(u,i)\}$, where each (u,i) pair means that user u has a positive feedback to item i.
- **Goal:** recommend each user $u \in \mathcal{U}$ a personalized ranked list of items from the set of unobserved items, i.e., $\mathcal{I}\backslash\mathcal{I}_u$.

Asymmetric Pairwise Preference Assumption

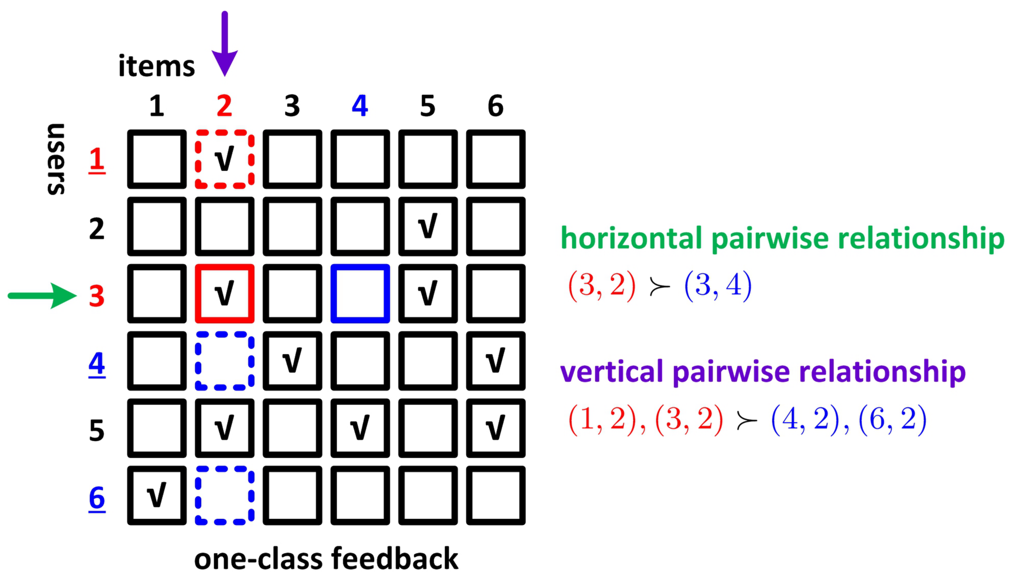

We keep the horizontal pairwise preference assumption in BPR [2, 3] and assume that an item is preferred by a group of interacted users to a group of un-interacted users in order to make the vertical one more reasonable and comparable,

$$\hat{r}_{ui} > \hat{r}_{uj}, \hat{r}_{\mathcal{P}i} > \hat{r}_{\mathcal{N}i}, \tag{1}$$

where $i \in \mathcal{I}_u$, $j \in \mathcal{I}\backslash\mathcal{I}_u$, $\mathcal{P} \subseteq \mathcal{U}_i$, $\mathcal{N} \subseteq \mathcal{U}\backslash\mathcal{U}_i$ and $u \in \mathcal{P}$.

For instantiation of the relationship between the group preferences $\hat{r}_{\mathcal{P}i}$ and $\hat{r}_{\mathcal{N}i}$, we propose "Many 'Group vs. One' (MGO)" inspired by "Many 'Set vs. One' (MSO)" [1],

$$\hat{r}_{\mathcal{P}i} > \hat{r}_{wi}, w \in \mathcal{N}, \tag{2}$$

where $\hat{r}_{\mathcal{P}i} = \frac{1}{|\mathcal{P}|}\sum_{u' \in \mathcal{P}} \hat{r}_{u'i}$ is the overall preference of user-group $\mathcal{P}$ to item i.

Notice that $\hat{r}_{u'i} = U_{u'\cdot}V_{i\cdot}^T + b_{u'} + b_i$ is the prediction rule for the preference of user u' to item i, where $U_{u'\cdot} \in \mathbb{R}^{1\times d}$ and $V_{i\cdot} \in \mathbb{R}^{1\times d}$ are latent feature vectors of user u' and item i, respectively, and $b_{u'} \in \mathbb{R}$ and $b_i \in \mathbb{R}$ are the bias of user u' and item i, respectively.

Asymmetric Bayesian Personalized Ranking (ABPR)

Based on the asymmetric pairwise preference assumption in Eqs.(1-2), we reach an objective function in our asymmetric Bayesian personalized ranking (ABPR) for each quintuple $(u, i, j, \mathcal{P}, \mathcal{N})$,

$$\min_{\Theta} -\ln\sigma(\hat{r}_{uij}) - \frac{1}{|\mathcal{N}|}\sum_{w\in\mathcal{N}} \ln\sigma(\hat{r}_{i\mathcal{P}w}) + \text{reg}(u,i,j,\mathcal{P},\mathcal{N}), \tag{3}$$

where $\Theta = \{U_{u\cdot}, V_{i\cdot}, b_u, b_i, u \in \mathcal{U}, i \in \mathcal{I}\}$ are the model parameters to be learned, $\hat{r}_{uij} = \hat{r}_{ui} - \hat{r}_{uj}$ and $\hat{r}_{i\mathcal{P}w} = \hat{r}_{\mathcal{P}i} - \hat{r}_{wi}$ denote the corresponding preference differences, and $\text{reg}(u,i,j,\mathcal{P},\mathcal{N}) = \frac{\alpha}{2}||V_{i\cdot}||^2 + \frac{\alpha}{2}||V_{j\cdot}||^2 + \frac{\alpha}{2}||b_i||^2 + \frac{\alpha}{2}||b_j||^2 + \sum_{u'\in\mathcal{P}}[\frac{\alpha}{2}||U_{u'\cdot}||^2 + \frac{\alpha}{2}||b_{u'}||^2] + \sum_{w\in\mathcal{N}}[\frac{\alpha}{2}||U_{w\cdot}||^2 + \frac{\alpha}{2}||b_w||^2]$ is the regularization term used to avoid overfitting.

Experiments

Dataset	Method		Precision@5	Recall@5	F1@5	NDCG@5	1-call@5
ML20M	Pointwise	MF	$0.1249_{\pm 0.0014}$	$0.0755_{\pm 0.0015}$	$0.0773_{\pm 0.0012}$	$0.1378_{\pm 0.0014}$	$0.4479_{\pm 0.0028}$
		LogMF	$0.1622_{\pm 0.0002}$	$0.0918_{\pm 0.0002}$	$0.0951_{\pm 0.0001}$	$0.1805_{\pm 0.0001}$	$0.5269_{\pm 0.0007}$
		FISM	$0.1351_{\pm 0.0014}$	$0.0821_{\pm 0.0010}$	$0.0836_{\pm 0.0009}$	$0.1505_{\pm 0.0019}$	$0.4755_{\pm 0.0028}$
	Pairwise	BPR	$0.1645_{\pm 0.0009}$	$0.0862_{\pm 0.0007}$	$0.0921_{\pm 0.0007}$	$0.1810_{\pm 0.0013}$	$0.5228_{\pm 0.0016}$
		BPR^T	$0.0632_{\pm 0.0024}$	$0.0364_{\pm 0.0015}$	$0.0373_{\pm 0.0015}$	$0.0693_{\pm 0.0025}$	$0.2536_{\pm 0.0078}$
		MBPR	$0.1609_{\pm 0.0004}$	$0.0893_{\pm 0.0007}$	$0.0931_{\pm 0.0006}$	$0.1797_{\pm 0.0004}$	$0.5262_{\pm 0.0025}$
		ABPR	**$0.1709_{\pm 0.0005}$**	**$0.0940_{\pm 0.0005}$**	**$0.0985_{\pm 0.0004}$**	**$0.1907_{\pm 0.0007}$**	**$0.5482_{\pm 0.0016}$**
NF50KU	Pointwise	MF	$0.1385_{\pm 0.0002}$	$0.0497_{\pm 0.0009}$	$0.0581_{\pm 0.0005}$	$0.1464_{\pm 0.0008}$	$0.4685_{\pm 0.0019}$
		LogMF	$0.1665_{\pm 0.0012}$	$0.0580_{\pm 0.0005}$	$0.0671_{\pm 0.0004}$	$0.1773_{\pm 0.0007}$	$0.5185_{\pm 0.0016}$
		FISM	$0.1447_{\pm 0.0017}$	$0.0520_{\pm 0.0010}$	$0.0606_{\pm 0.0008}$	$0.1536_{\pm 0.0024}$	$0.4833_{\pm 0.0028}$
	Pairwise	BPR	$0.1664_{\pm 0.0015}$	$0.0560_{\pm 0.0006}$	$0.0654_{\pm 0.0006}$	$0.1765_{\pm 0.0011}$	$0.5112_{\pm 0.0025}$
		BPR^T	$0.0912_{\pm 0.0048}$	$0.0302_{\pm 0.0014}$	$0.0349_{\pm 0.0017}$	$0.0971_{\pm 0.0057}$	$0.3275_{\pm 0.0139}$
		MBPR	$0.1794_{\pm 0.0012}$	$0.0659_{\pm 0.0009}$	$0.0748_{\pm 0.0008}$	$0.1928_{\pm 0.0013}$	$0.5489_{\pm 0.0033}$
		ABPR	**$0.1854_{\pm 0.0008}$**	**$0.0685_{\pm 0.0004}$**	**$0.0777_{\pm 0.0004}$**	**$0.1995_{\pm 0.0004}$**	**$0.5601_{\pm 0.0010}$**

Table 1: Recommendation performance of our ABPR, three pointwise preference learning methods, i.e., MF, LogMF and FISM, and three pairwise preference learning methods, i.e., BPR, BPR^T and MBPR, on ML20M and NF50KU w.r.t. five commonly used ranking-oriented evaluation metrics. The significantly best results are marked in bold (p-value < 0.015).

Observations:

- (i) our ABPR performs significantly better (p-value is smaller than 0.015) than all the baseline methods in all cases, which clearly shows the effectiveness of our proposed asymmetric pairwise preference assumption in modeling one-class feedback;
- (ii) MBPR performs similar to BPR on ML20M and better than BPR on NF50KU, which shows the sensitivity of the mutual pairwise preference assumption in MBPR w.r.t. different datasets (notice that our ABPR performs better than BPR on both datasets showcasing the superiority of our relaxed vertical pairwise relationship in our asymmetric assumption);
- (iii) the performance of BPR is much better than that of BPR^T, which is expected as the horizontal preference relationship, i.e., $(u,i) \succ (u,j)$, is more reasonable than the transposed one, i.e., $(u,i) \succ (w,i)$, considering the probable incomparability of preferences between different users (notice that our ABPR with asymmetric assumption, i.e., $(u,i) \succ (u,j)$ and $(\mathcal{P},i) \succ (\mathcal{N},i)$, performs the best); and
- (iv) for the methods with pointwise preference assumption, LogMF performs similar to BPR and better than MF and FISM, which shows the importance of an appropriate loss function in modeling one-class feedback.

Conclusions

- We study an important recommendation problem called one-class collaborative filtering (OCCF).
- We propose a novel preference assumption called asymmetric pairwise preference assumption.
- We design a novel recommendation algorithm called asymmetric Bayesian personalized ranking (ABPR).

References

[1] Weike Pan, Li Chen, and Zhong Ming. Personalized recommendation with implicit feedback via learning pairwise preferences over item-sets. *Knowledge and Information Systems*, 58(2):295–318, 2019.

[2] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars 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.

[3] Lu Yu, Ge Zhou, Chu-Xu Zhang, Junming Huang, Chuang Liu, and Zi-Ke Zhang. Rankmbpr: Rank-aware mutual bayesian personalized ranking for item recommendation. In *Proceedings of the 17th International Conference Web-Age Information Management*, WAIM'16, pages 244–256, 2016.

Acknowledgements

We thank the support of National Natural Science Foundation of China Nos. 61872249, 61836005 and 61672358.