Popularity-Opportunity Bias in Collaborative Filtering

Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang and James Caverlee
Department of Computer Science and Engineering, Texas A&M University, USA

Introduction

- **Problem and goal:**

Popularity bias is a long-standing problem in recommender systems. However, the conventional concept of popularity bias is aligned with the concept of **statistical parity**, and so inherit its limitations. In this paper, we re-examine popularity bias from the perspective of the concept of **equal opportunity**, which evaluate the bias with ground truth of user-item matching into consideration.

- **Contributions:**

(i) We propose to study the **popularity-opportunity bias** and from the views of both **user-side** and **item-side**.

(ii) We conduct a comprehensive **data-driven study** over four datasets to investigate the presence of the popularity-opportunity bias.

(iii) We **theoretically analyze** the impact of item popularity on ranking by MF and BPR to confirm the existence of the bias in both methods.

(iv) We investigate the potential of a **post-processing approach** to reduce this bias. Through experiments on four datasets, we explore the trade-offs between debiasing effectiveness and recommendation utility, showing the more effective debiasing performance of the proposed method over existing debiasing baselines designed for conventional popularity bias.

Conventional Popularity Bias

- **Investigate whether popular items are recommended more frequently than less popular items (a statistical parity based concept), leading to rich-get-richer.**

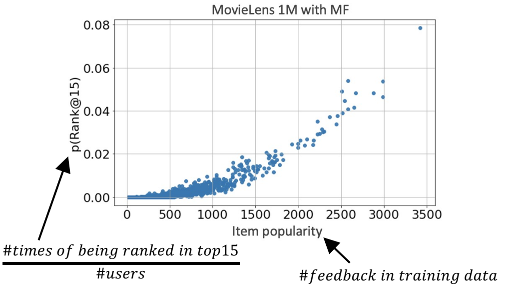

Drawback: conventional popularity bias is NOT always harmful

- **Recommendation showing conventional popularity bias is not necessarily problematic. Yet, enforcing zero conventional popularity bias may bring issues.**

Is conventional popularity bias harmful in this case?
Looks ok (rich-get-richer won't happen)

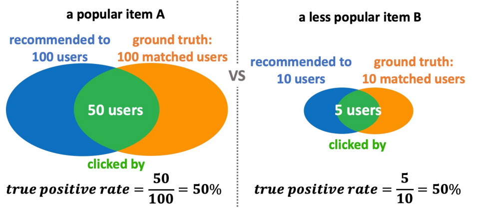

What if enforce no popularity bias following prior works?
Looks unfair!

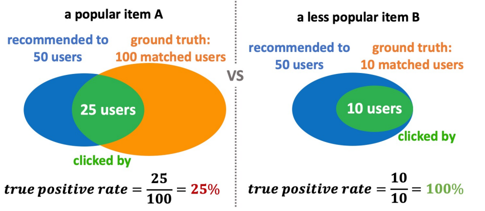

Popularity-Opportunity Bias

We propose to investigate the **popularity-opportunity bias**, which compares the **probability of being recommended to matched users (i.e., true positive rate)** for items of different popularity.

- **Conventional popularity bias** does not consider the ground truth of user-item matching.
- **Proposed popularity-opportunity bias** takes account of the ground truth of user-item matching.

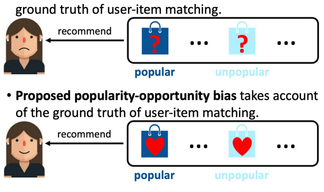

Two Views of the Bias

- **User-view popularity-opportunity bias (uPO bias)**

Given user u likes a popular item i and a less popular item j, whether i will be ranked higher than j?

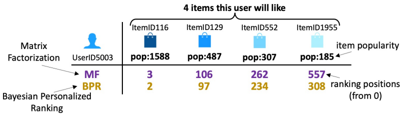

$$PRU = -\frac{1}{N}\sum_{u\in\mathcal{U}} \boxed{SRC}\left(\boxed{pop(\widetilde{O}_u^+)}, \boxed{rank_u(\widetilde{O}_u^+)}\right)$$

Spearman's rank correlation
popularity of items user u likes in testing set
popularity of items user u likes in testing set

- **Item-view popularity-opportunity bias (iPO bias)**

Whether popular items have higher expected ranking to matched users than less popular items?

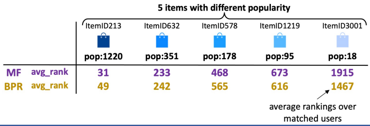

$$PRI = -SRC\left(\boxed{pop(\mathcal{I})}, \boxed{avg\_rank(\mathcal{I})}\right)$$

average rankings of items for matched users
popularity of all items

Prevalence of the Bias

Table: Measuring uPO bias (PRU) and iPO bias (PRI) for MF and BPR on four datasets.

	ML1M		Ciao		Epinions		App	
	MF	BPR	MF	BPR	MF	BPR	MF	BPR
PRU	0.835	0.779	0.542	0.591	0.684	0.708	0.567	0.636
PRI	0.980	0.969	0.363	0.433	0.535	0.573	0.609	0.692

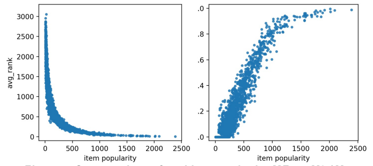

Figures: Scatter plots of ranking results by MF on ML1M.

Conclusion

- Propose the study of **popularity-opportunity bias**;
- **Empirically** show the prevalence of the bias;
- **Theoretically** show how two models inherently produce bias on both user and item sides (**refer to the paper**);
- Propose the **Popularity Compensation** debiasing method, and empirically show the effectiveness of the proposed method.

Debiasing: Popularity Compensation (PC)

- **Promote less popular items by adding compensation to predicted scores**

1.Calculate compensation based on popularity: $C_{u,i} = \boxed{\frac{1}{pop(i)}} \cdot (\beta \cdot \boxed{\widehat{\mathbf{R}}_{u,i}})$

less popular items are compensated more
items more likely to be liked are compensated more

2.Add the compensation to predicted score: $\widehat{\mathbf{R}}^*_{u,i} = \widehat{\mathbf{R}}_{u,i} + \boxed{\alpha} \cdot C_{u,i} \cdot \frac{\|\mathbf{C}_u\|}{\|\mathbf{R}_u\|}$

weight to control the strength of debiasing

Debiasing Experiments

		NDCG@k @20	*NDCG@k* @50	*PRU*	*PRI*
ML1M	MF	0.2726	0.2930	0.8350	0.9799
	MF-weight	0.1484	0.1793	0.4845	0.6407
	MF-rescale	0.1361	0.1658	0.4365	0.6936
	MF-PC	0.1435	0.1980	0.4552	0.5594
Ciao	MF	0.0717	0.0934	0.5420	0.3625
	MF-weight	0.0447	0.0675	0.3174	0.3293
	MF-rescale	0.0425	0.0608	0.3219	0.2526
	MF-PC	0.0647	0.0845	0.3073	−0.0150
Epinions	MF	0.0693	0.0938	0.6840	0.5351
	MF-weight	0.0349	0.0526	0.3453	0.2341
	MF-rescale	0.0343	0.0509	0.3678	0.2182
	MF-PC	0.0605	0.0848	0.3549	−0.0415
App	MF	0.1026	0.1359	0.5667	0.6089
	MF-weight	0.0388	0.0596	0.3552	0.2334
	MF-rescale	0.0384	0.0583	0.3350	0.2147
	MF-PC	0.0965	0.1280	0.3527	−0.0487

Our proposed method
Baselines to reduce the conventional popularity bias

- All of the proposed method and two baselines can **reduce the bias** from both user-view and item-view, but also **sacrifice recommendation utility** at the same time;
- Proposed PC reduce the bias to similar degree as the two baselines;
- Proposed PC **preserve the recommendation utility** better than the two baselines.