Popularity-Opportunity Bias in Collaborative Filtering

Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang, and James Caverlee

Texas A&M University

Recommenders – essential conduits

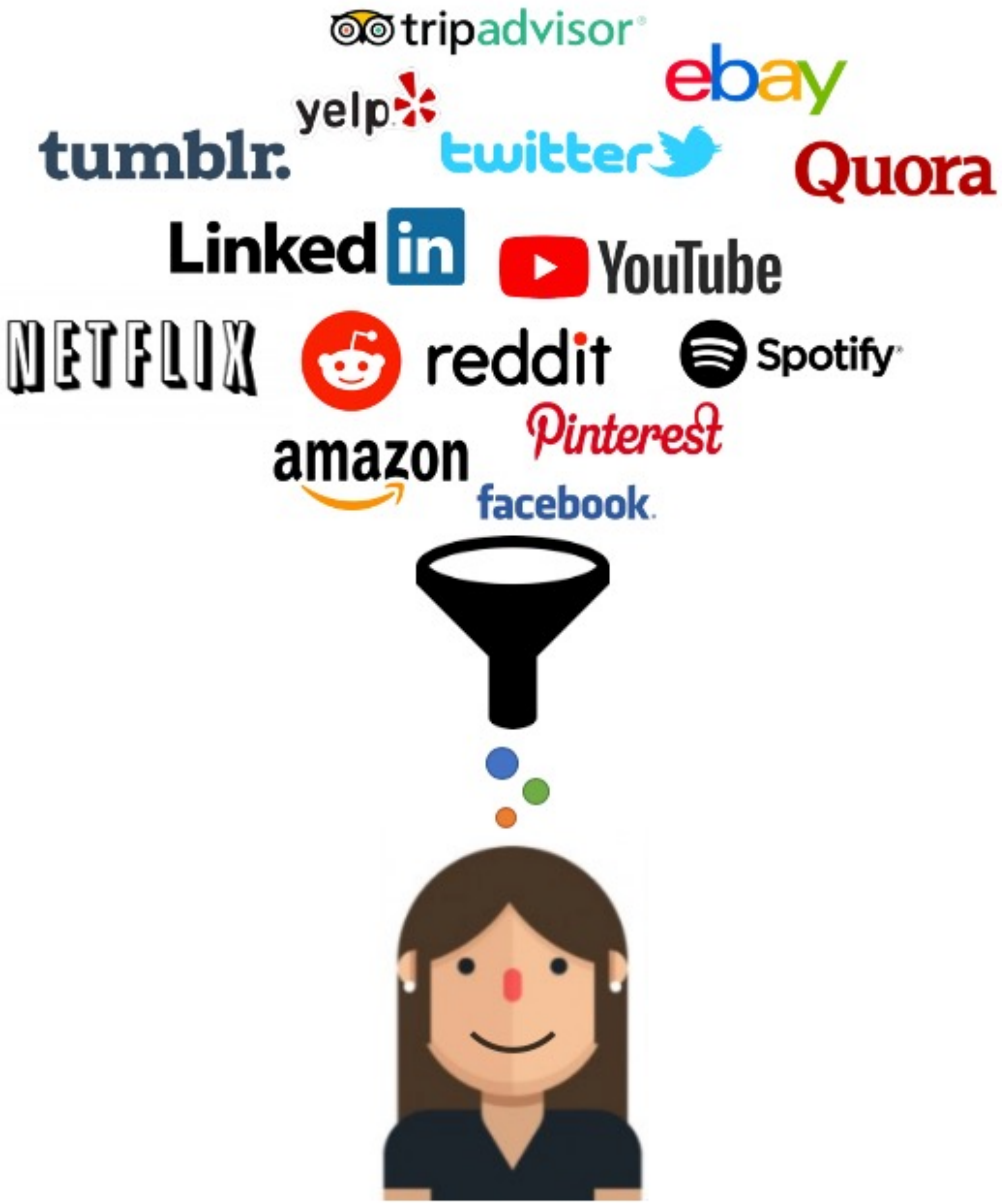

Recommenders – pursue higher and higher utility

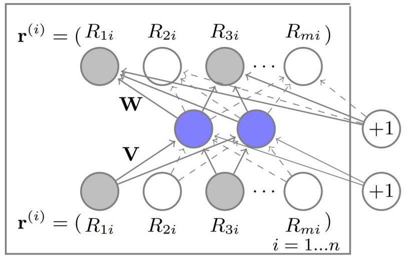

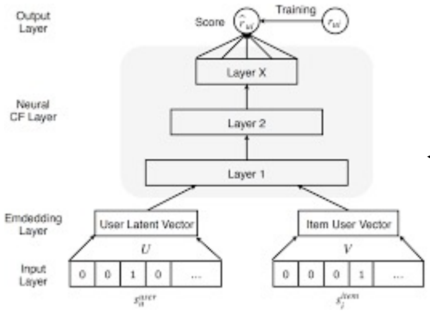

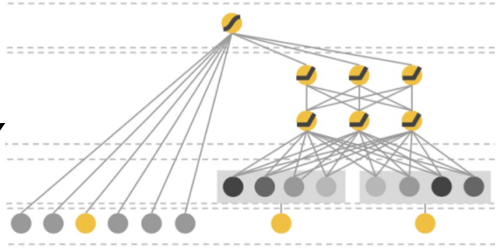

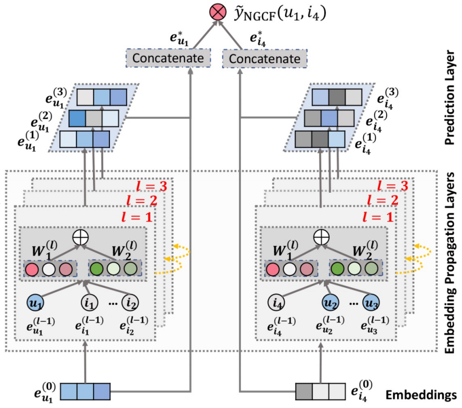

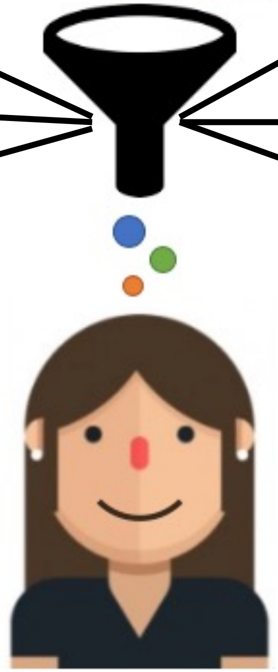

NDCG; recall; precision; MRR; CTR; MAP…

Recommenders – raise higher bias at the same time

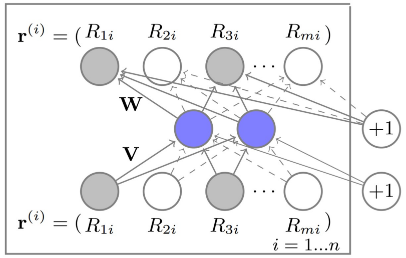

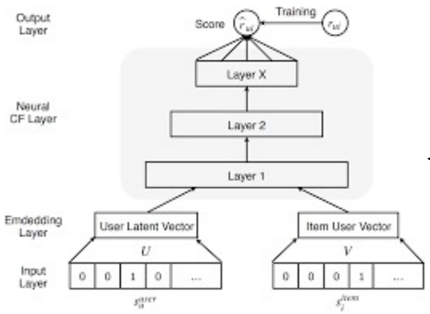

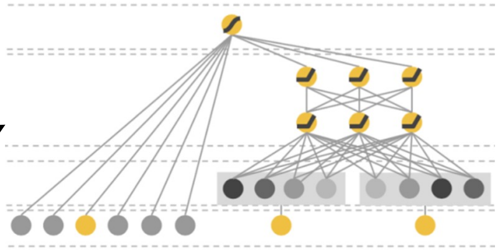

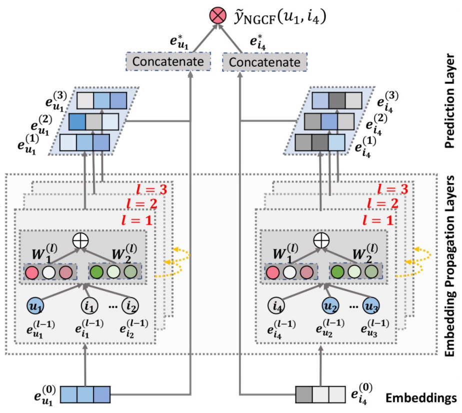

NDCG; recall; precision; MRR; CTR; MAP…

bias

Popularity Bias

Popular items are recommended more frequently than less popular items (a demographic parity based concept), leading to rich-get-richer problem.

$$\frac{\#times\ of\ being\ ranked\ in\ top15}{\#users}$$

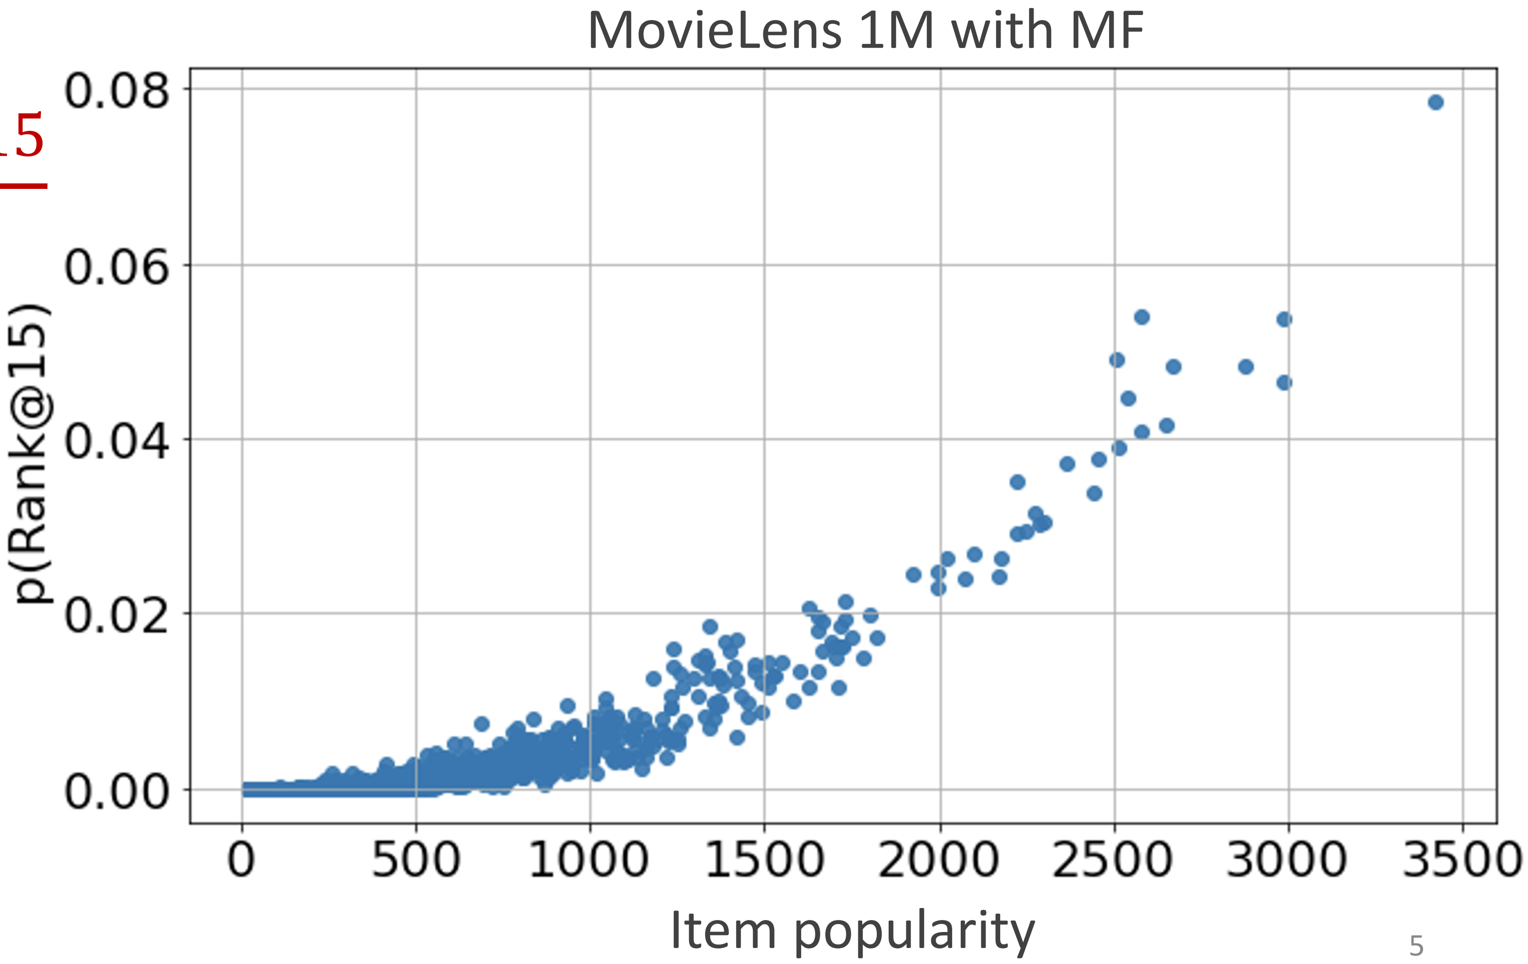

Drawback of Conventional Concept of Popularity Bias

Is the popularity bias really a problem?

a popular item A

recommended to 100 users

VS

a less popular item B

recommended to 10 users

Drawback of Conventional Concept of Popularity Bias

Is the popularity bias really a problem?

a popular item A

recommended to 100 users

ground truth: 100 matched users

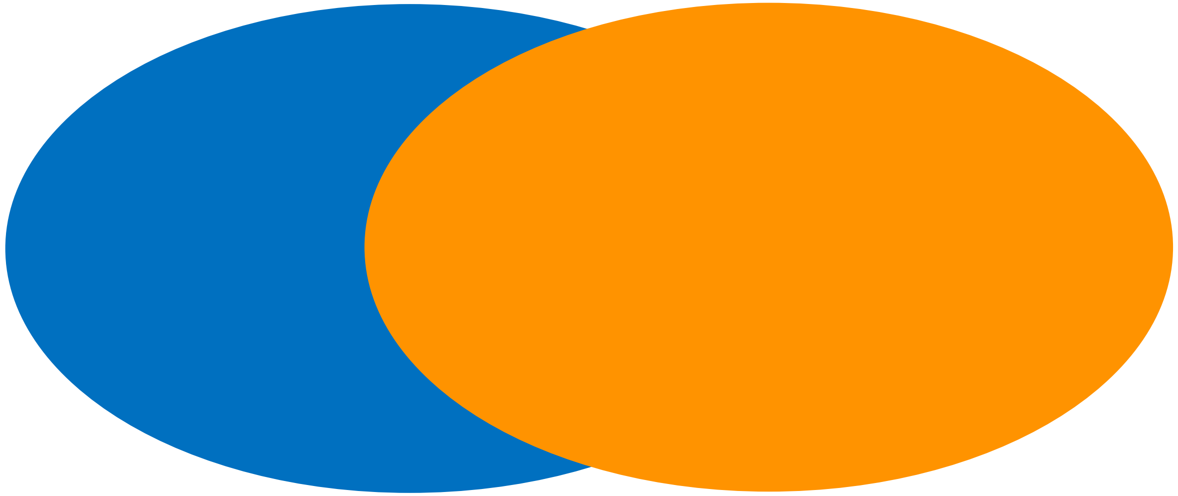

VS

a less popular item B

recommended to 10 users

ground truth: 10 matched users

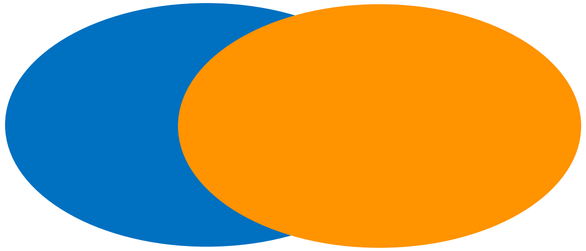

Drawback of Conventional Concept of Popularity Bias

Is the popularity bias really a problem?

Looks ok (rich-get-richer won't happen)

a popular item A

recommended to 100 users

ground truth: 100 matched users

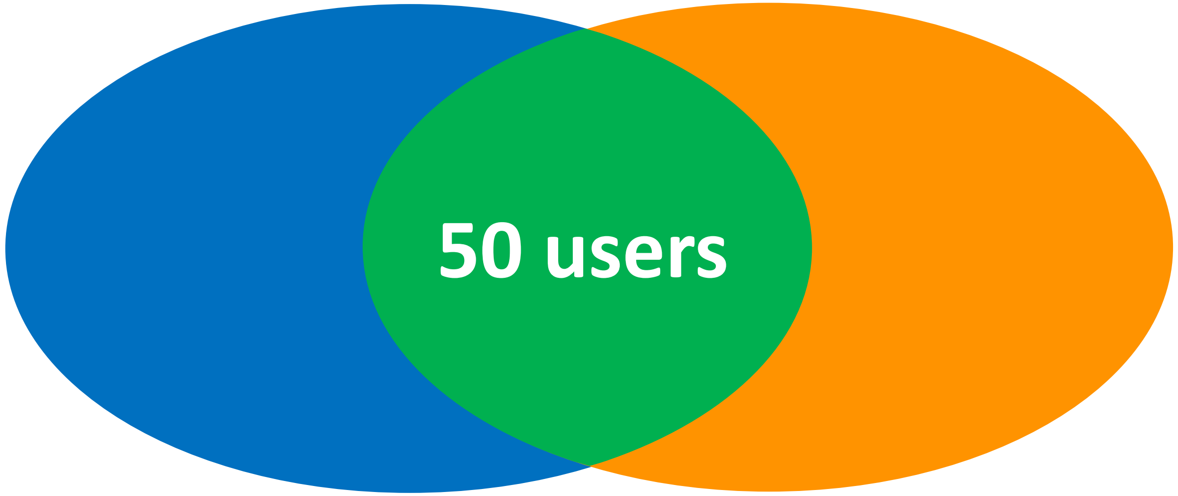

clicked by

$$true\ positive\ rate = \frac{50}{100} = 50\%$$

VS

a less popular item B

recommended to 10 users

ground truth: 10 matched users

clicked by

$$true\ positive\ rate = \frac{5}{10} = 50\%$$

Drawback of Conventional Concept of Popularity Bias

Enforce no popularity bias following existing works.

a popular item A

recommended to 50 users

VS

a less popular item B

recommended to 50 users

Drawback of Conventional Concept of Popularity Bias

Enforce no popularity bias following existing works.

Looks unfair

a popular item A

recommended to 50 users

ground truth: 100 matched users

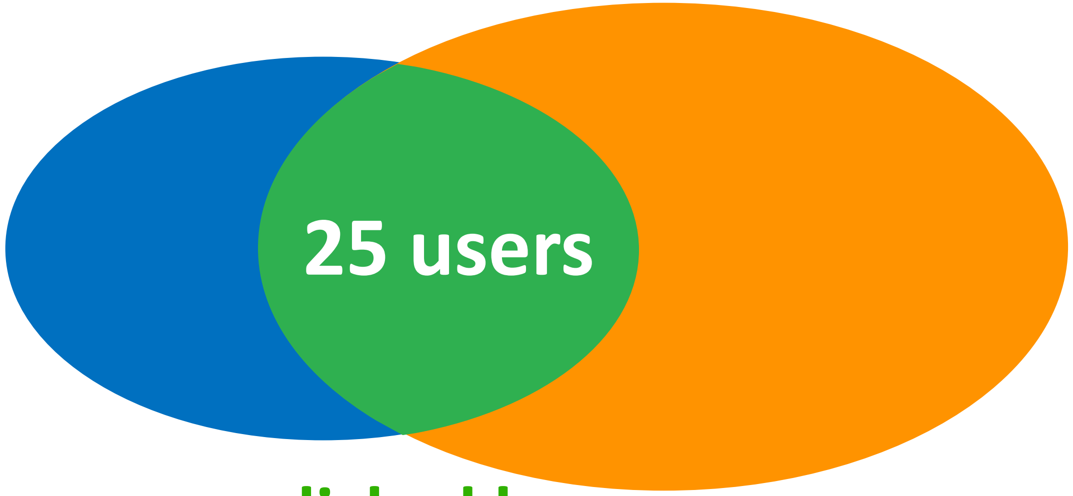

clicked by

$$\boldsymbol{true\ positive\ rate} = \frac{\mathbf{25}}{\mathbf{100}} = \mathbf{25\%}$$

VS

a less popular item B

recommended to 50 users

ground truth: 10 matched users

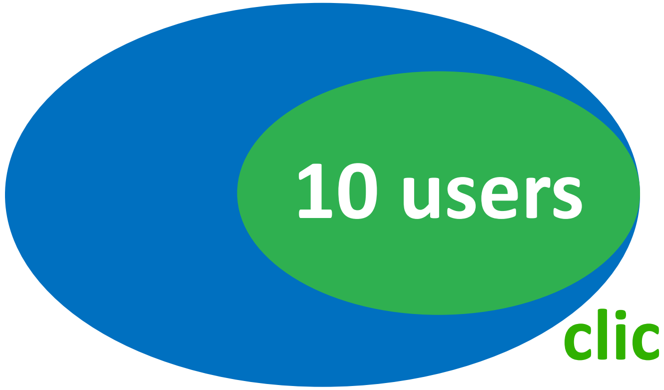

clicked by

$$\boldsymbol{true\ positive\ rate} = \frac{\mathbf{10}}{\mathbf{10}} = \mathbf{100\%}$$

Problem with Conventional Popularity Bias

Conventional concept of popularity bias compares the **recommendation to all users** for items **without considering the ground-truth of user-item matching**.

However, in practice, only the recommendation to matched users can influence the feedback or economic gain items receive.

Solution – Popularity-opportunity Bias

Take user-item matching into consideration, and compare the **probability of being recommended to matched users** for items of different popularity (an equal opportunity based concept).

Example of Popularity-opportunity Bias

This is a real problem, rich-get-richer will happen.

a popular item A

recommended to 100 users

ground truth: 100 matched users

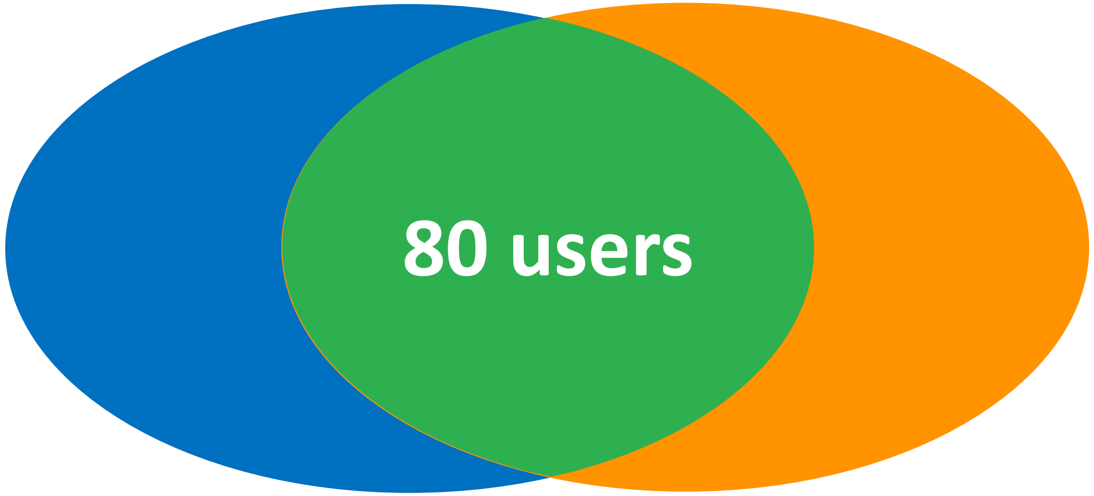

clicked by

$$\boldsymbol{true\ positive\ rate} = \frac{\mathbf{80}}{\mathbf{100}} = \mathbf{80\%}$$

VS

a less popular item B

recommended to 10 users

ground truth: 10 matched users

clicked by

$$\boldsymbol{true\ positive\ rate} = \frac{\mathbf{2}}{\mathbf{10}} = \mathbf{20\%}$$

Popularity-opportunity bias – two views

- User-view popularity-opportunity bias (uPO bias)
- Item-view popularity-opportunity bias (iPO bias)

Popularity-opportunity bias – user-view

- 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*?

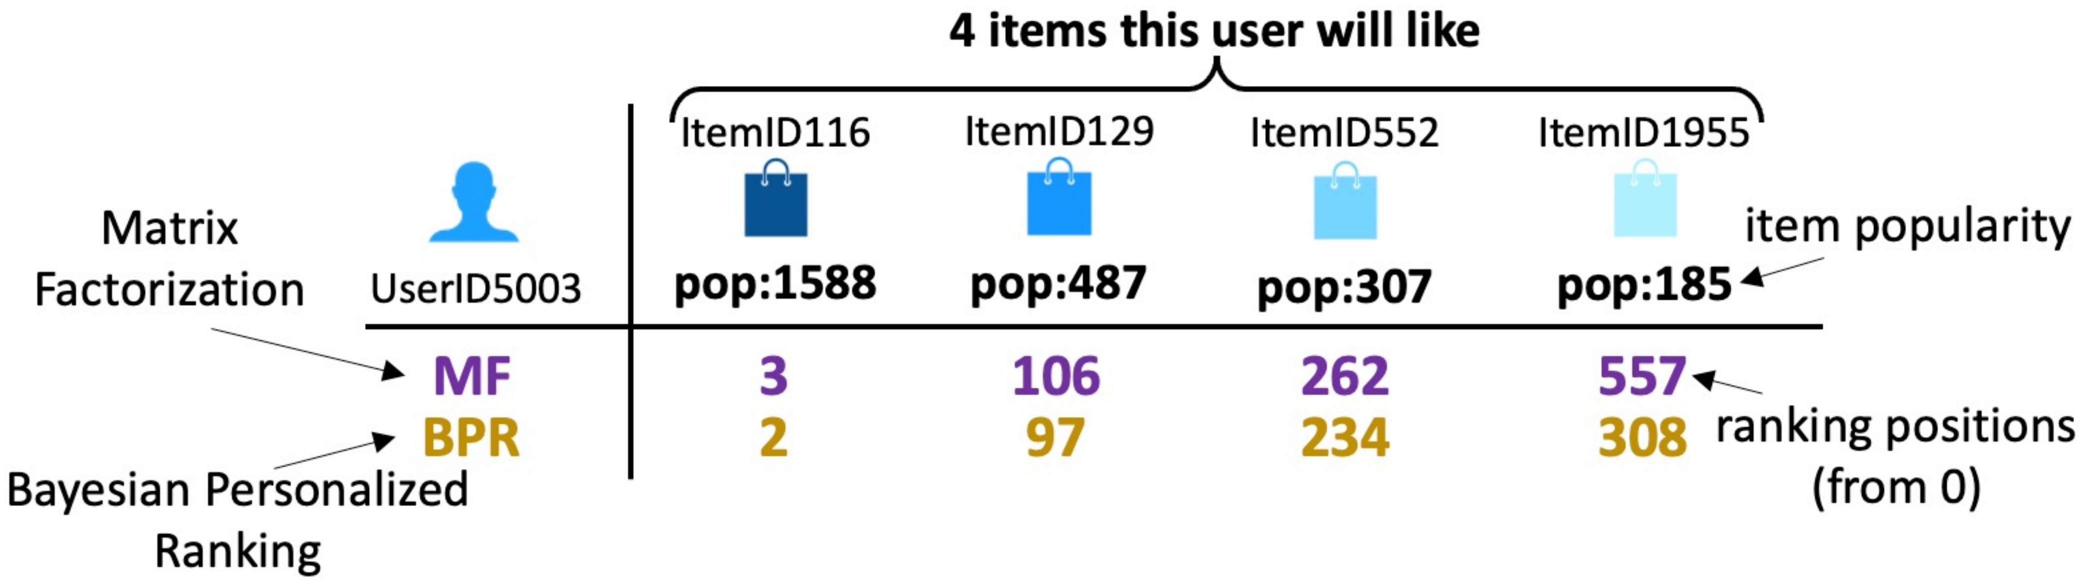

Popularity-opportunity bias – user-view

- User-view popularity-opportunity bias (uPO bias)
  Measure uPO bias by popularity-rank correlation for users (**PRU**).
  From 0 to 1: higher value represents severer bias.

Spearman's rank correlation

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

popularity of items user *u* likes in testing set

popularity of items user *u* likes in testing set

Popularity-opportunity bias – item-view

- Item-view popularity-opportunity bias (iPO bias)
  Whether popular items have higher expected ranking to matched users than less popular items?

5 items with different popularity

		ItemID213	ItemID632	ItemID578	ItemID1219	ItemID3001
		pop:1220	**pop:351**	**pop:178**	**pop:95**	**pop:18**
MF	**avg_rank**	**31**	**233**	**468**	**673**	**1915**
BPR	**avg_rank**	**49**	**242**	**565**	**616**	**1467**

average rankings over matched users

Popularity-opportunity bias – item-view

- Item-view popularity-opportunity bias (iPO bias)
  Measure iPO bias by popularity-rank correlation for items (**PRI**).
  From 0 to 1: higher value represents severer bias.

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

popularity of all items

average rankings of items for matched users

Empirically show the prevalence of popularity-opportunity bias

	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

High bias measured for both user and item views for both models and four datasets

Debiasing approach – Popularity Compensation (PC)

Promote less popular items by adding compensation to predicted scores.

$$\mathbf{C}_{u,i} = \frac{1}{pop(i)} \cdot (\beta \cdot \widehat{\mathbf{R}}_{u,i})$$

Calculate compensation based on popularity

$$\widehat{\mathbf{R}}^{*}_{u,i} = \widehat{\mathbf{R}}_{u,i} + \alpha \cdot \mathbf{C}_{u,i} \cdot \frac{\|\mathbf{C}_{u}\|}{\|\mathbf{R}_{u}\|}$$

Add the compensation to predicted score

Debiasing approach – Popularity Compensation (PC)

Promote less popular items by adding compensation to predicted scores.

$$\mathbf{C}_{u,i} = \boxed{\frac{1}{pop(i)}} \cdot (\beta \cdot \boxed{\widehat{\mathbf{R}}_{u,i}})$$

Calculate compensation based on popularity

$$\widehat{\mathbf{R}}^{*}_{u,i} = \widehat{\mathbf{R}}_{u,i} + \alpha \cdot \mathbf{C}_{u,i} \cdot \frac{\|\mathbf{C}_{u}\|}{\|\mathbf{R}_{u}\|}$$

Add the compensation to predicted score

Debiasing approach – Popularity Compensation (PC)

Promote less popular items by adding compensation to predicted scores.

$$\mathbf{C}_{u,i} = \frac{1}{pop(i)} \cdot (\beta \cdot \widehat{\mathbf{R}}_{u,i})$$

Calculate compensation based on popularity

$$\boxed{\widehat{\mathbf{R}}^*_{u,i} = \widehat{\mathbf{R}}_{u,i} + \alpha \cdot \mathbf{C}_{u,i} \cdot \frac{\|\mathbf{C}_u\|}{\|\mathbf{R}_u\|}}$$

Add the compensation to predicted score

Experiment result

		NDCG@k		PRU	PRI
		@20	@50		
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

Methods to reduce the conventional popularity bias

Proposed popularity compensation method

Experiment result

		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

Proposed PC method reduces the popularity-opportunity bias to similar degree as conventional popularity debiasing methods

Experiment result

		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

Proposed PC method preserves utility better than conventional popularity debiasing methods

Conclusions

- Propose the study of **popularity-opportunity bias**;
- Empirically show the vulnerability of two matrix factorization models to the bias by a **data-driven study** on four datasets;
- **Theoretically** show how these two models inherently produce the popularity-opportunity 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 to reduce the popularity-opportunity bias and preserve recommendation utility compared with conventional popularity debiasing methods.

Thank You!

Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang, and James Caverlee
Texas A&M University

Special thank to SIGIR Travel Grants program!