中国科学技术大学
University of Science and Technology of China

Cross Pairwise Ranking for Unbiased Item Recommendation

Qi Wan, Xiangnan He, Xiang Wang, Jiancan Wu, Wei Guo, Ruiming Tang
wqq17@mail.ustc.edu.cn

• Popularity Bias in Recommender Systems

- ➢ Observed Interactions affected by previous exposure mechanism, biased towards popular items
- ➢ Recommender trained without consideration of data bias further amplifies the bias

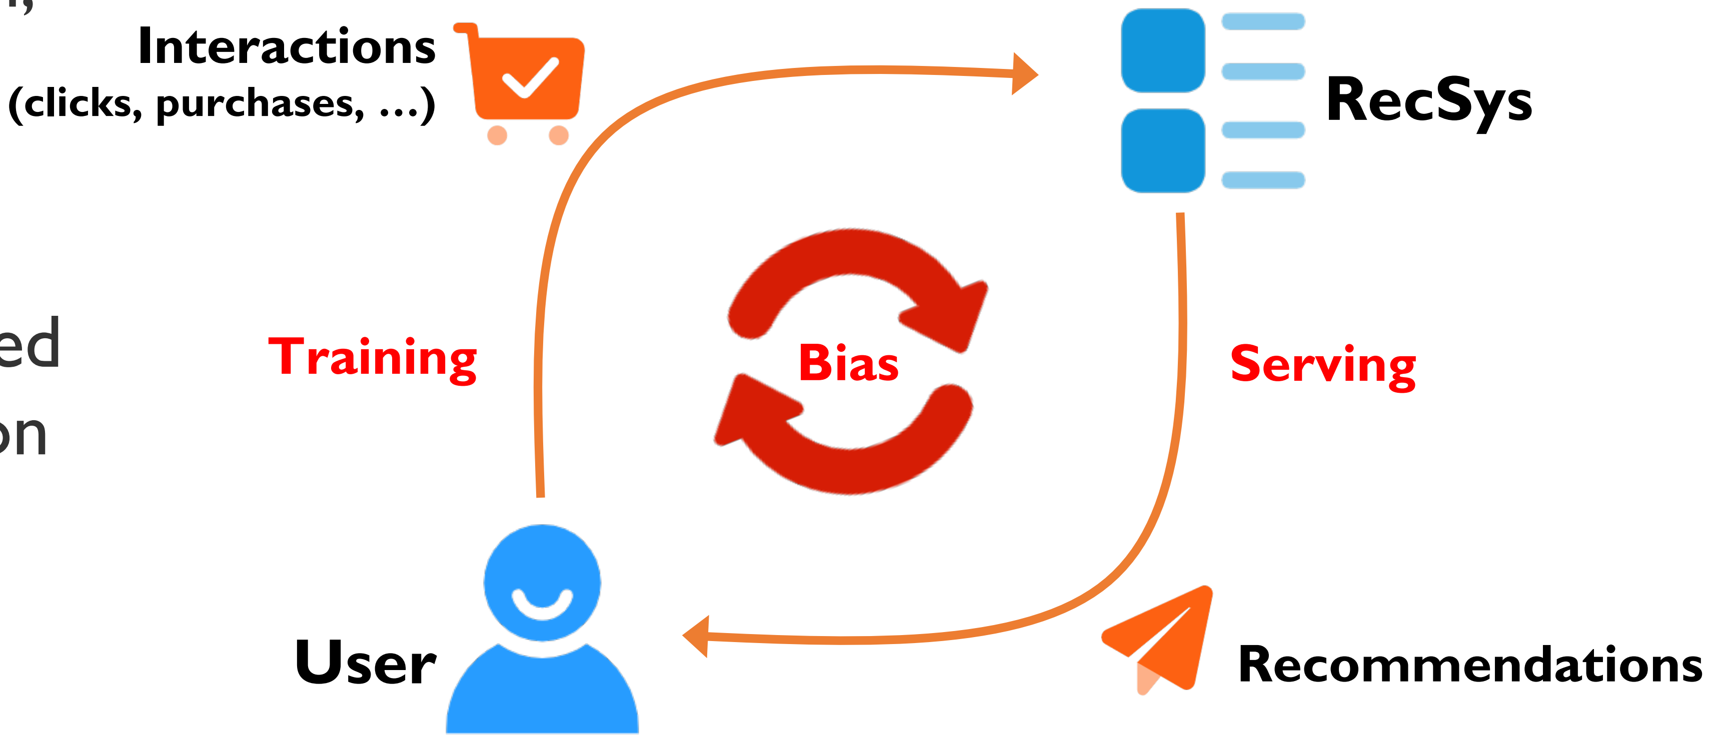

• Non-uniform Exposure Brings the Bias

- Binary variables
  - $Y_{u,i}$: interaction; $R_{u,i}$: relevance; $O_{u,i}$: observation

$$
\begin{aligned}
P(Y_{u,i}=1) &= P(R_{u,i}=1, O_{u,i}=1) \\
&= \underbrace{P(R_{u,i}=1)}_{\text{relevance probability}}\underbrace{P(O_{u,i}=1 \mid R_{u,i}=1)}_{\text{exposure probability}}
\end{aligned}
$$

interaction probability

generally higher for popular items

- Popular items dominate training

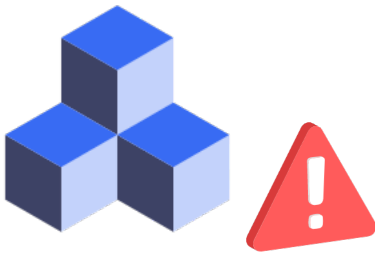

• Mainstream Debiasing Method

Inverse Propensity Scoring (IPS)

- ➢ Reweight samples with estimated propensity scores
  - ➢ Unknown exposure mechanism → Propensities difficult to estimate
- ➢ Aiming to achieve an unbiased expectation of loss
  - ➢ Large variance of loss → expectation-variance trade-off

New Perspective on Popularity Bias

• Definition of Unbiasedness

- True relevance score $s_{u,i} = \ln P(R_{u,i} = 1)$
- Predicted relevance score $\hat{s}_{u,i}$

Define unbiasedness from ranking perspective:

A loss function L is unbiased if it optimizes the ranking of predicted user-item relevance scores towards that of the true relevance scores:

$$\hat{s}_{u,i} > \hat{s}_{u,j} \Leftrightarrow s_{u,i} > s_{u,j} \text{ when } \mathcal{L} \text{ converges.}$$

• Biasedness of Pointwise and Pairwise Loss

- Pointwise Loss

$$\mathcal{L}_{BCE} = -\sum_{(u,i)\in\mathcal{D}} Y_{u,i}\ln\sigma(\hat{s}_{u,i}) + (1-Y_{u,i})\ln(1-\sigma(\hat{s}_{u,i}))$$

$$\begin{cases} \hat{s}_{u,i} \geq C & \text{if } Y_{u,i} = 1 \\ \hat{s}_{u,i} < C & \text{if } Y_{u,i} = 0 \end{cases}$$

- Pairwise Loss

$$\mathcal{L}_{BPR} = -\sum_{(u,i,j)\in\mathcal{D}_S} \ln\sigma(\hat{s}_{u,i} - \hat{s}_{u,j})$$

$$\hat{s}_{u,i} - \hat{s}_{u,j} > 0,\ \text{if } Y_{u,i} = 1, Y_{u,j} = 0$$

- Real Objective

$$\begin{cases} \ln P(Y_{u,i} = 1) \geq C & \text{if } Y_{u,i} = 1 \\ \ln P(Y_{u,i} = 1) < C & \text{if } Y_{u,i} = 0 \end{cases}$$

$$\hat{s}_{u,i} \Rightarrow s_{u,i} + \ln P(O_{u,i} = 1 \mid R_{u,i} = 1)$$

higher for active users & popular items

The full proof can be found in our paper.

• Our Solution: Cross Pairwise Ranking Loss

- Select two positive user-item pairs (u_1, i_1) and (u_2, i_1), such that
- (u_1, i_2) and (u_2, i_1) are negative pairs

$$\mathcal{L}_{CPR} = -\sum_{(u_1,u_2,i_1,i_2)\in\mathcal{D}_2} \ln\sigma\left[\frac{1}{2}\left(\hat{s}_{u_1,i_1} + \hat{s}_{u_2,i_2} - \hat{s}_{u_1,i_2} - \hat{s}_{u_2,i_1}\right)\right]$$

$$\mathcal{D}_2 = \{(u_1, u_2, i_1, i_2) \mid Y_{u_1,i_1} = 1, Y_{u_2,i_2} = 1, Y_{u_1,i_2} = 0, Y_{u_2,i_1} = 0\}$$

$$\hat{s}_{u_1,i_1} + \hat{s}_{u_2,i_2} - \hat{s}_{u_1,i_2} - \hat{s}_{u_2,i_1} > 0$$

$$\text{if } Y_{u_1,i_1} = 1, Y_{u_2,i_2} = 1, Y_{u_1,i_2} = 0, Y_{u_2,i_1} = 0$$

• Unbiasedness of CPR Loss

➢ Assumption: $P(O_{u,i}=1 \mid R_{u,i}=1) = p_u \cdot p_i \cdot P(R_{u,i}=1)^{\alpha}$

User-specific propensity

Item-specific propensity

$\alpha > 0$

Unknown

➢ Unbiased ranking

Pointwise

$$\begin{cases} s_{u,i} \geq \frac{1}{1+\alpha}(C - \ln p_u - \ln p_i) & \text{if } Y_{u,i}=1 \\ s_{u,i} < \frac{1}{1+\alpha}(C - \ln p_u - \ln p_i) & \text{if } Y_{u,i}=0 \end{cases}$$

Pairwise

$$s_{u,i} - s_{u,j} > \frac{1}{1+\alpha}(-\ln p_i + \ln p_j), \text{ if } Y_{u,i}=1, Y_{u,j}=0$$

Cross Pairwise

$$s_{u_1,i_1} + s_{u_2,i_2} - s_{u_1,i_2} - s_{u_2,i_1} > 0$$

$$\text{if } Y_{u_1,i_1}=1, Y_{u_2,i_2}=1, Y_{u_1,i_2}=0, Y_{u_2,i_1}=0$$

$\hat{s}_{u,i} \Rightarrow s_{u,i}$

The full proof can be found in our paper.

• Extending to More Interactions

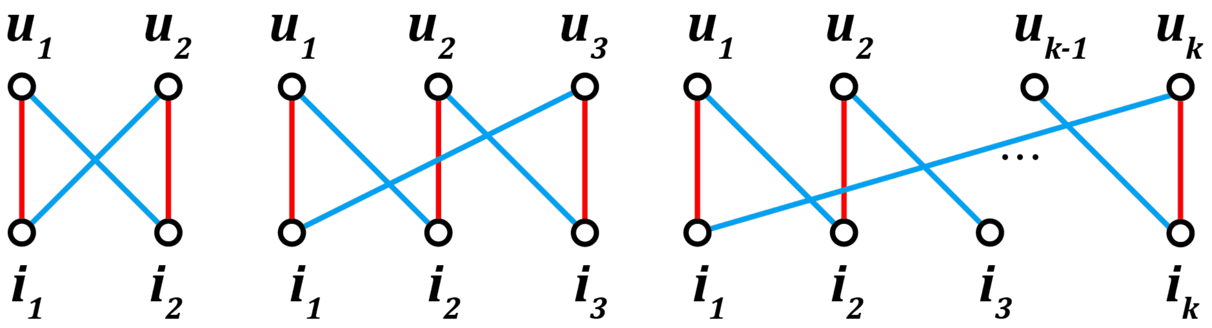

Red lines – positive
Blue lines – negative

$$\mathcal{L}_{CPR} = -\sum_{k} \sum_{(u_1,\ldots,u_k,i_1,\ldots,i_k)\in\mathcal{D}_k} \ln \sigma\left[\frac{1}{k}\cdot\left(\hat{s}_{u_1,i_1}+\cdots+\hat{s}_{u_k,i_k}-\hat{s}_{u_1,i_2}-\cdots-\hat{s}_{u_k,i_1}\right)\right]$$

$$\mathcal{D}_k = \{(u_1,\ldots,u_k,i_1,\ldots,i_k) \mid Y_{u_1,i_1}=1, Y_{u_2,i_2}=1,\ldots,\\ Y_{u_k,i_k}=1, Y_{u_1,i_2}=0, Y_{u_2,i_3}=0,\ldots,Y_{u_k,i_1}=0\}, k\geq 2$$

• Dynamic Sampling

- Goal: Speed up the training & Improve the performance.
- Inspired by DNS [], select harder samples with higher probabilities.
- For a sample $S_j = (u_1, \ldots, u_k, i_1, \ldots, i_k)$, measure its difficulty by

$$O_j = \frac{1}{k} \cdot (\hat{\boldsymbol{s}}_{u_1,i_1} + \hat{\boldsymbol{s}}_{u_2,i_2} + \cdots + \hat{\boldsymbol{s}}_{u_k,i_k} - \hat{\boldsymbol{s}}_{u_1,i_2} - \hat{\boldsymbol{s}}_{u_2,i_3} - \cdots - \hat{\boldsymbol{s}}_{u_k,i_1})$$

- Dynamic sampling rate $\beta \geq 1$
- Only use the hardest $\frac{1}{\beta}$ of selected data as training samples

Zhang, Weinan, et al. "Optimizing top-n collaborative filtering via dynamic negative item sampling." *SIGIR*. 2013.

• Experimental Settings

- MovieLens-10M, Netflix Prize and Alibaba iFashion
- Create simulated unbiased data by lowering the sampling rate of popular items → Validation / Testing set
- Evaluation metrics:
  - Overall: Recall@K, NDCG@K
  - Biasedness: ARP@K (Average Recommendation Popularity)
- Baselines
  - Traditional: BPR, Mult-VAE
  - Debiasing: CausE, Rel-MF, UBPR, DICE

• Overall Comparison

Method	Backbone	MovieLens			Netflix			iFashion		
		Recall↑	NDCG↑	ARP↓	Recall	NDCG	ARP	Recall	NDCG	ARP
BPR	MF	0.1579	0.0939	4084	0.1255	0.0952	3341	0.0278	0.0122	436
Mult-VAE	-	0.1676	0.1004	4468	0.1242	0.0921	3554	<u>0.0309</u>	<u>0.0142</u>	415
Debiasing approaches:										
CausE	MF	0.1440	0.0805	3814	0.1080	0.0730	3206	0.0214	0.0094	408
Rel-MF	MF	0.1463	0.0829	4107	0.1138	0.0765	3254	0.0250	0.0113	421
UBPR	MF	0.1682	0.1000	2691	0.1315	0.1007	2289	0.0281	0.0126	404
DICE	MF	<u>0.1835</u>	<u>0.1101</u>	<u>1173</u>	<u>0.1317</u>	<u>0.1007</u>	<u>1244</u>	0.0275	0.0124	<u>**250**</u>
CPR-rand	MF	0.1938	0.1192	**1055**	0.1462	0.1143	1270	0.0307	0.0137	398
CPR	MF	**0.2003**★	**0.1223**★	1138	**0.1511**★	**0.1190**★	**1204**	**0.0332**★	**0.0151**★	359
%Improv.		9.16%	11.08%	-	14.73%	18.17%	-	7.44%	6.34%	-

Without dynamic sampling

➢ CPR-rand beats all the baselines on MovieLens and Netflix, achieves better overall performance on iFashion

➢ CPR further improves performances on all datasets

• Training Efficiency

- ➢ Recall curves
  - ➢ CPR converges to the best performance with the least number of epochs

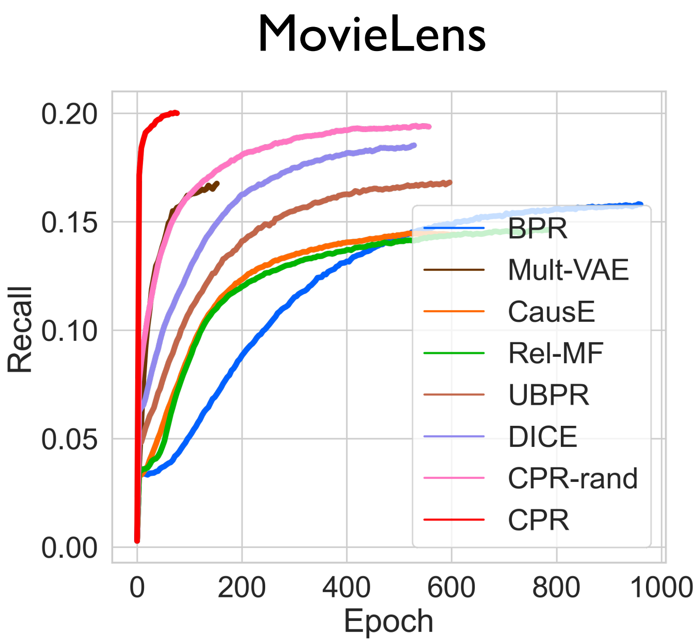

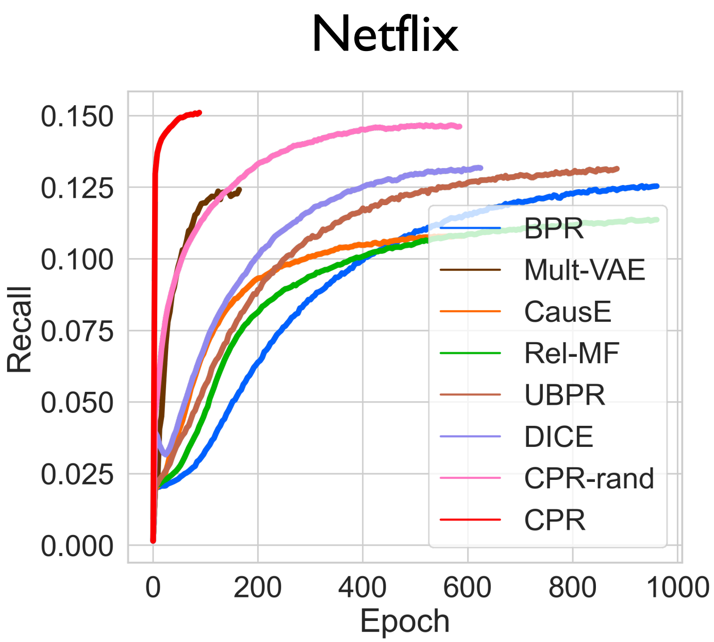

- ➢ Total training time
  - ➢ CPR takes much less time compared with the strongest baseline DICE

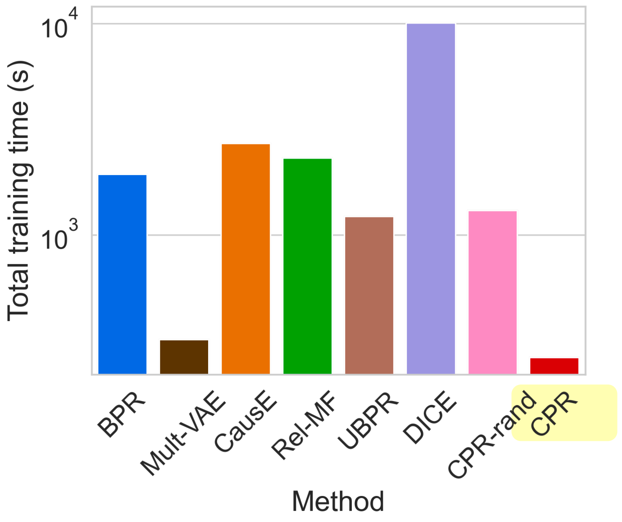

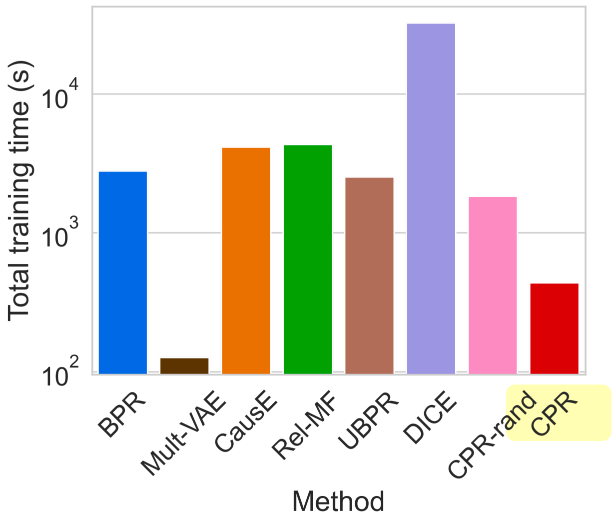

• Debiasing Ability

- ➢ Four groups of item –
  - ➢ sort items by their degrees in the ascending order, and group them such that the sum degrees in each group are approximately equal
  - ➢ Lines: the percentage of degrees contained in each group
  - ➢ Columns: the percentage of the recommended items of each group

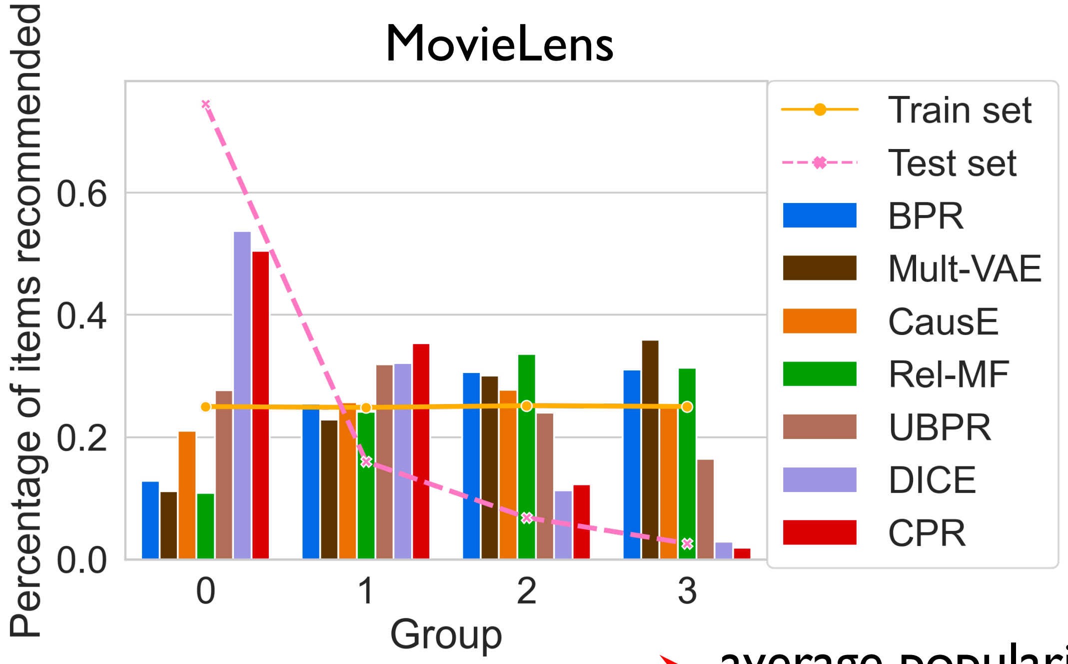

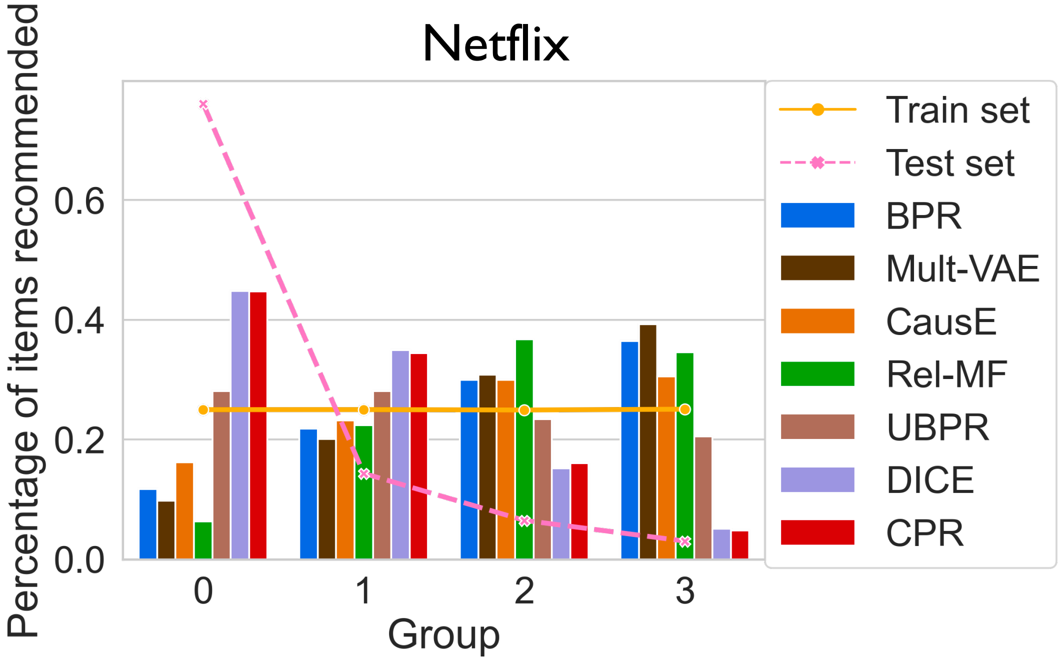

average popularity ↑

• Generalization Ability

Performances *w.r.t.* different backbones

Method	Backbone	MovieLens		Netflix	
		Recall	NDCG	Recall	NDCG
BPR	LightGCN	0.1661	0.0995	0.1310	0.0985
DICE	LightGCN	0.1864	0.1121	0.1355	0.1043
CPR	LightGCN	**0.1963**	**0.1210**	**0.1502**	**0.1180**
BPR	NeuMF	0.1519	0.0890	0.1254	0.0932
DICE	NeuMF	0.1568	0.0914	0.1126	0.0841
CPR	NeuMF	**0.1895**	**0.1155**	**0.1434**	**0.1109**
BPR	NGCF	0.1615	0.0951	0.1199	0.0899
DICE	NGCF	0.1774	0.1071	0.1312	0.0999
CPR	NGCF	**0.1952**	**0.1182**	**0.1498**	**0.1173**

Performances *w.r.t.* different degrees of bias

	θ	Method	MovieLens		Netflix	
			Recall	NDCG	Recall	NDCG
bias amplified	0.5	BPR	0.1217	0.0679	0.0939	0.0654
	0.5	DICE	0.1461	0.0839	0.1048	0.0741
	0.5	CPR	**0.1584**	**0.0927**	**0.1189**	**0.0858**
bias reduced	-0.5	BPR	0.1413	0.0809	0.1086	0.0777
	-0.5	DICE	0.1597	0.0956	0.1133	0.0858
	-0.5	CPR	**0.1737**	**0.1040**	**0.1274**	**0.0965**

➢ CPR consistently performs better than baselines

• Conclusion & Future Work

- Analyze the biasedness of mainstream loss functions from a new perspective, showing that they optimize recommendations towards a biased ranking of user preference
- Propose a new unbiased method CPR

- Future work
  - a more general assumption of exposure probability
  - extend CPR to other scenarios, like alleviating the group fairness bias

Thank you!

Codes and Datasets available at https://github.com/Qcactus/CPR