Revealing, Analyzing, and Mitigating the Adverse Impact of Shared Accounts on Personalized Recommendations for Single-User Accounts

The 1st Workshop on Human-Centered Recommender Systems

The ACM Web Conference 2025, April 28 – May 2, Sydney

Xingbo Jing, Tao Lian*, & Li Wang

Taiyuan University of Technology

Motivation

Background

- Most collaborative filtering algorithms assume each account is used by a single user.
- However, there exist many shared accounts in some applications (*e.g.*, IPTV).

Problem

- Shared accounts can breed spurious associations between items that are actually consumed by distinct members, leading to inaccurate recommendations for other single-user accounts.

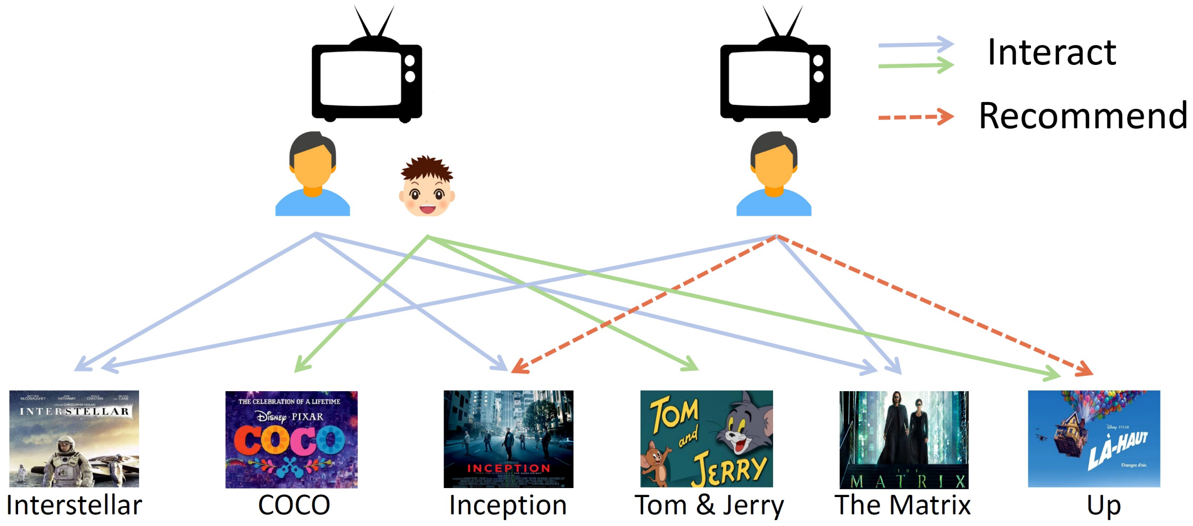

Research Questions

1. How to quantitatively reveal the adverse impact of shared accounts on recommendation results for other single-user accounts?
   - It is hard to know *how a model performs for the same user in two counterfactual settings with and without shared accounts*.
2. What are the underlying causes of such impact?
   - We conjecture that shared accounts hurt item embeddings in model-based CF, yet *how the proximity relations between item embeddings change across settings* remains unexplored.
3. How to mitigate the aforementioned adverse impact?
   - *How to deal with spurious associations* between a pair of items consumed by two members with distinct interests in a shared account?

Revealing the Adverse Impact on Recommendation Performance

- Simulate two settings with the same test set for a fair comparison
  - Controlled Setting：only single-user accounts
  - Experimental Setting：both shared accounts and single-user accounts

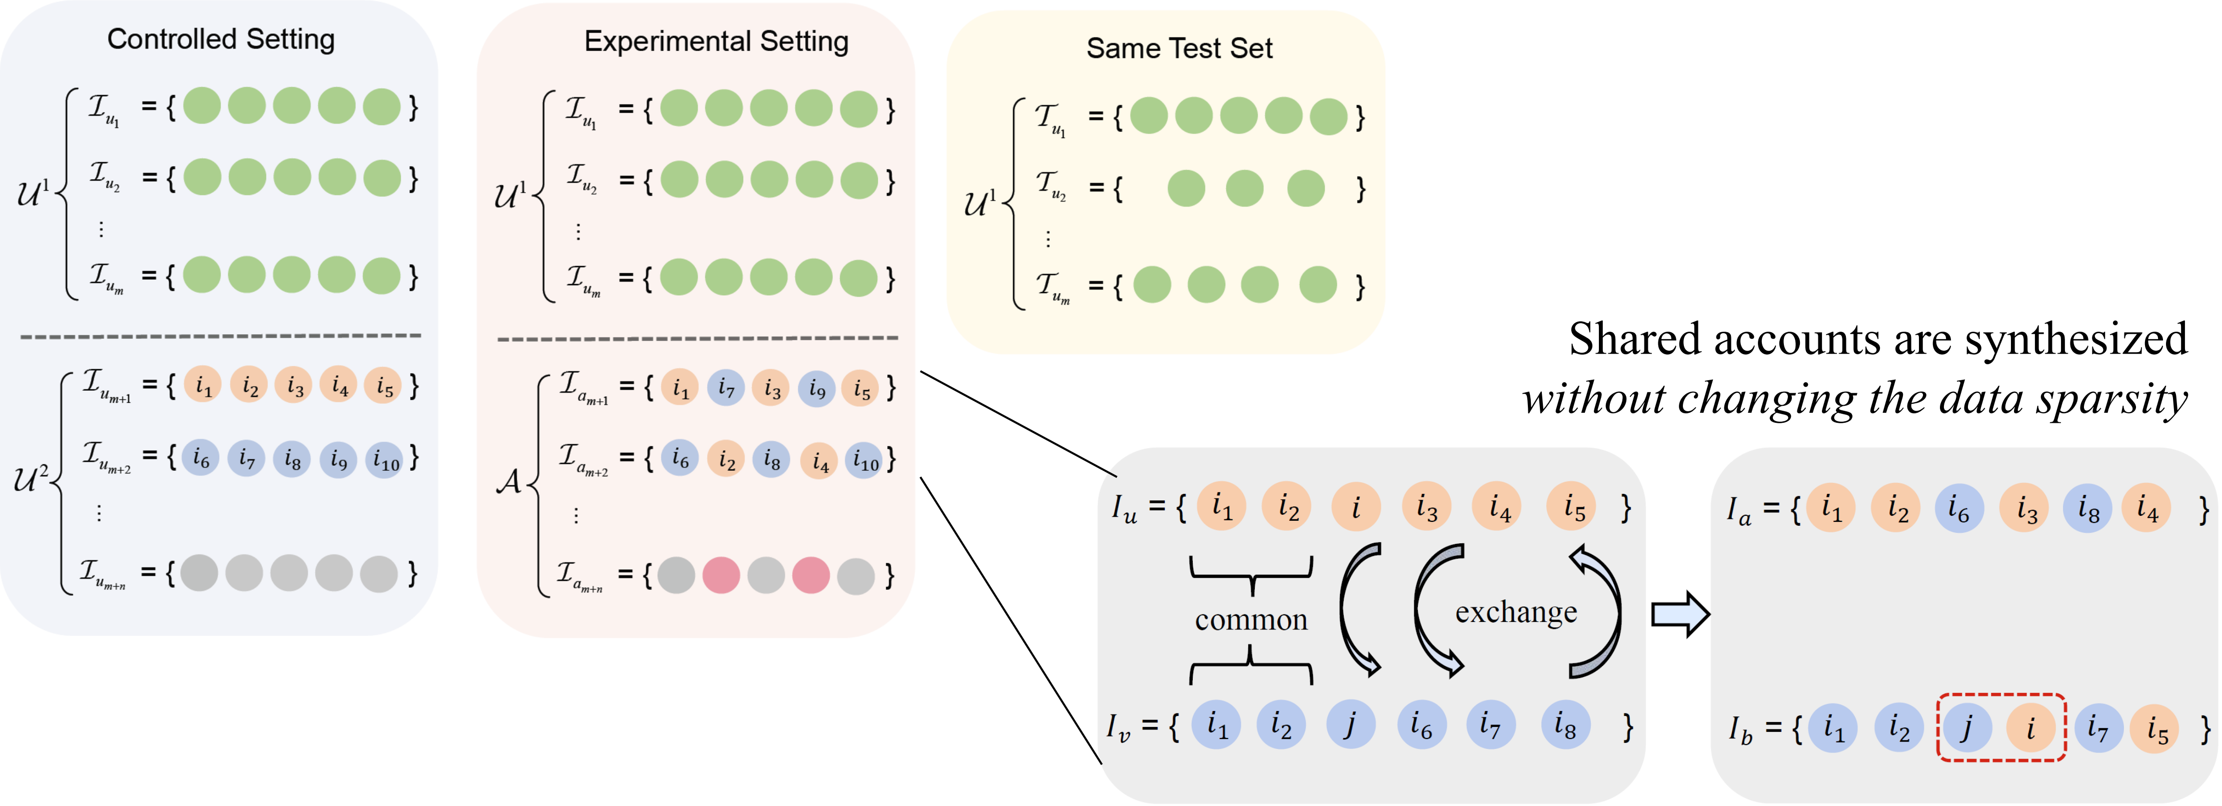

Revealing the Adverse Impact on Recommendation Performance

The presence of shared accounts leads to significant performance degradation for other single-user accounts.

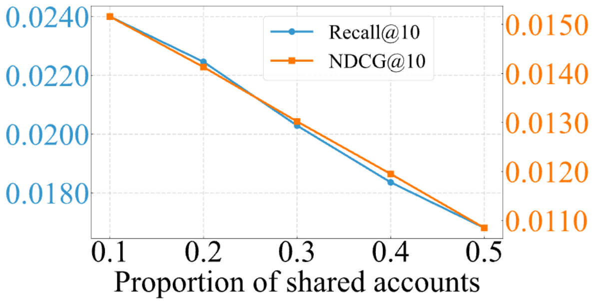

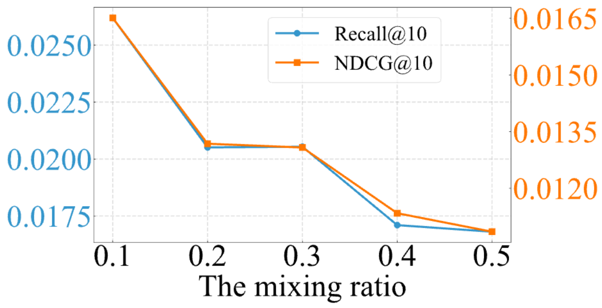

Controlled Setting vs. Experimental Setting

Dataset	Metric	Control	Experiment	Decrease↓
Clothing	Recall@10	0.0269	0.0168	37.55%
	Recall@20	0.0403	0.0240	40.45%
	NDCG@10	0.0168	0.0108	35.71%
	NDCG@20	0.0208	0.0130	37.50%
Beauty	Recall@10	0.0795	0.0640	19.50%
	Recall@20	0.1164	0.0927	20.36%
	NDCG@10	0.0508	0.0416	18.11%
	NDCG@20	0.0623	0.0505	18.94%
MircoLens	Recall@10	0.0371	0.0319	14.02%
	Recall@20	0.0632	0.0560	11.39%
	NDCG@10	0.0340	0.0285	16.18%
	NDCG@20	0.0447	0.0382	14.54%

Impact of the number of item pairs with strong spurious associations

Dataset	p	% shared accounts	Recall@20	NDCG@20
Clothing	0	0	0.0403	0.0208
	50	1.39%	0.0396	0.0205
	100	2.69%	0.0392	0.0203
Beauty	0	0	0.1164	0.0623
	50	3.23%	0.1135	0.0616
	100	6.13%	0.1117	0.0609
MicroLens	0	0	0.0632	0.0447
	50	6.00%	0.0578	0.0411
	100	10.04%	0.0550	0.0382

Analyzing the Adverse Impact on Item Embeddings

- Demoted relevant items: $\mathcal{T}_u^{\downarrow} = \left\{ i \in \mathcal{L}_u^S \mid i \in \mathcal{T}_u \wedge \pi\left(\mathcal{L}_u^S, i\right) < \pi\left(\mathcal{L}_u^A, i\right) \right\}$
- Promoted irrelevant items: $\mathcal{L}_u^{\uparrow} = \left\{ i \in \mathcal{L}_u^A \mid i \notin \mathcal{T}_u \wedge \pi\left(\mathcal{L}_u^A, i\right) < \pi\left(\mathcal{L}_u^S, i\right) \right\}$

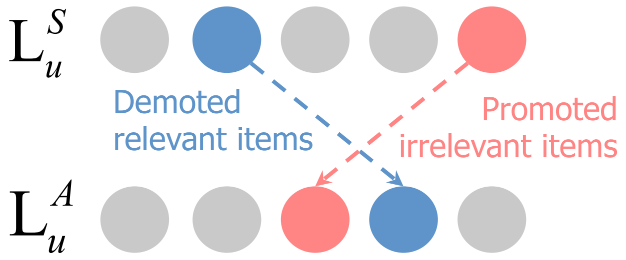

- Two metrics: their proximity to the user's consumed items in the latent space

$$\mathrm{Push}\left(i, \mathcal{I}_u\right) = \left| \left\{ j \in \mathcal{I}_u \mid \pi\left(\mathcal{N}_j^S, i\right) < \pi\left(\mathcal{N}_j^A, i\right) \right\} \right|$$

$$\mathrm{Pull}\left(i, \mathcal{I}_u\right) = \left| \left\{ j \in \mathcal{I}_u \mid \pi\left(\mathcal{N}_j^S, i\right) > \pi\left(\mathcal{N}_j^A, i\right) \right\} \right|$$

$$\mathrm{ratio\_push} = \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{T}_u^{\downarrow}} \mathbb{I}\left(\mathrm{Push}\left(i, \mathcal{I}_u\right) > \mathrm{Pull}\left(i, \mathcal{I}_u\right)\right)}{\sum_{u \in \mathcal{U}} \left|\mathcal{T}_u^{\downarrow}\right|}$$

$$\mathrm{ratio\_pull} = \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{L}_u^{\uparrow}} \mathbb{I}\left(\mathrm{Push}\left(i, \mathcal{I}_u\right) < \mathrm{Pull}\left(i, \mathcal{I}_u\right)\right)}{\sum_{u \in \mathcal{U}} \left|\mathcal{L}_u^{\uparrow}\right|}$$

KNN analysis of item embeddings

Dataset	K	ratio_push	ratio_pull
Clothing	10	0.8665	0.8648
	20	0.8660	0.8578
	50	0.8861	0.8455
Beauty	10	0.8290	0.8375
	20	0.8306	0.8310
	50	0.8357	0.8208
MicroLens	10	0.7891	0.7543
	20	0.8410	0.7676
	50	0.9014	0.7910

The proximity relations between item embeddings are greatly distorted due to the existence of shared accounts.

Mitigating the Adverse Impact by Using Item Multimodal Info.

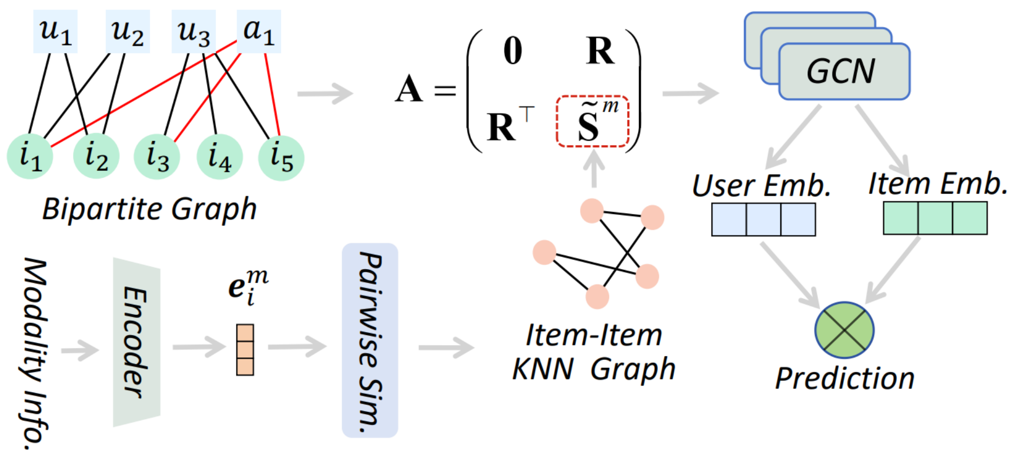

Dataset	Metric	Controlled LightGCN	Experimental LightGCN	VBPR	MMGCN	GRCN	BM3	Ours-V	Ours-T
Clothing	Recall@10	0.0269	0.0168	0.0094	0.0083	0.0248	0.0275	0.0253	**0.0339**
	Recall@20	0.0403	0.0240	0.0142	0.0153	0.0387	0.0408	0.0381	**0.0515**
	NDCG@10	0.0168	0.0108	0.0059	0.0045	0.0148	0.0175	0.0161	**0.0212**
	NDCG@20	0.0208	0.0130	0.0073	0.0065	0.0188	0.0214	0.0198	**0.0264**
Beauty	Recall@10	0.0795	0.0640	0.0389	0.0394	0.0628	0.0699	0.0750	**0.0765**
	Recall@20	0.1164	0.0927	0.0578	0.0632	0.0950	0.1070	0.1084	**0.1108**
	NDCG@10	0.0508	0.0416	0.0245	0.0236	0.0394	0.0419	0.0484	**0.0493**
	NDCG@20	0.0623	0.0505	0.0304	0.0307	0.0491	0.0526	0.0587	**0.0600**
MicroLens	Recall@10	0.0371	0.0319	0.0263	0.0207	0.0330	0.0259	0.0343	**0.0365**
	Recall@20	0.0632	0.0560	0.0462	0.0400	0.0577	0.0473	0.0581	**0.0616**
	NDCG@10	0.0340	0.0285	0.0238	0.0194	0.0288	0.0236	0.0310	**0.0331**
	NDCG@20	0.0447	0.0382	0.0319	0.0273	0.0389	0.0317	0.0407	**0.0431**

The performance decline is effectively alleviated by leveraging the reliable item-item semantic relations from image or text modality to counteract spurious second-order item-item associations in the bipartite behavior graph.

Mitigating the Adverse Impact by Using Item Multimodal Info.

- Fixing rank changes in the recommendation list

$$\text{fix\_demot} = \frac{\sum_{u\in\mathcal{U}}\sum_{i\in\mathcal{T}_u^{\downarrow}} \frac{1}{\log_2(\pi(\mathcal{L}_u^M,i)+1)} - \frac{1}{\log_2(\pi(\mathcal{L}_u^A,i)+1)}}{\sum_{u\in\mathcal{U}}\left|\mathcal{T}_u^{\downarrow}\right|}$$

$$\text{fix\_promo} = \frac{\sum_{u\in\mathcal{U}}\sum_{i\in\mathcal{L}_u^{\uparrow}} \frac{1}{\log_2(\pi(\mathcal{L}_u^M,i)+1)} - \frac{1}{\log_2(\pi(\mathcal{L}_u^A,i)+1)}}{\sum_{u\in\mathcal{U}}\left|\mathcal{L}_u^{\uparrow}\right|}$$

Dataset	Clothing	Beauty	MicroLens
fix_demot	0.0559	0.0601	0.0538
fix_promo	-0.2678	-0.2546	-0.2125

- Demoted relevant items are now moved forward by our mitigation method.
- Promoted irrelevant items are now moved backward by our mitigation method.

Mitigating the Adverse Impact by Using Item Multimodal Info.

□ Ameliorating the item embeddings

$$\mathrm{Push}\left(i, \mathcal{I}_u\right) = \left|\left\{ j \in \mathcal{I}_u \mid \pi\left(\mathcal{N}_j^S, i\right) \prec \pi\left(\mathcal{N}_j^M, i\right)\right\}\right|;$$

$$\mathrm{Pull}\left(i, \mathcal{I}_u\right) = \left|\left\{ j \in \mathcal{I}_u \mid \pi\left(\mathcal{N}_j^S, i\right) \succ \pi\left(\mathcal{N}_j^M, i\right)\right\}\right|.$$

$$\text{ratio\_push} = \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{T}_u^{\downarrow}} \mathbb{I}\left(\mathrm{Push}\left(i, \mathcal{I}_u\right) > \mathrm{Pull}\left(i, \mathcal{I}_u\right)\right)}{\sum_{u \in \mathcal{U}} \left|\mathcal{T}_u^{\downarrow}\right|}$$

$$\text{ratio\_pull} = \frac{\sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{L}_u^{\uparrow}} \mathbb{I}\left(\mathrm{Push}\left(i, \mathcal{I}_u\right) < \mathrm{Pull}\left(i, \mathcal{I}_u\right)\right)}{\sum_{u \in \mathcal{U}} \left|\mathcal{L}_u^{\uparrow}\right|}$$

Dataset	Clothing		Beauty		MicroLens	
	Ori.	Our	Ori.	Our	Ori.	Our
ratio_push	0.8660	0.7579	0.8306	0.7426	0.8410	0.8175
ratio_pull	0.8578	0.2604	0.8310	0.3960	0.7676	0.3529

Our mitigation method reduces ratio_push and ratio_pull significantly compared to the original model.

The item-item semantic graph can partly remedy the detrimental effect on proximity relations between item embeddings caused by shared accounts.

Mitigating the Adverse Impact by Using Item Multimodal Info.

Counteracting strong spurious associations

- For each item pair with strong spurious association
  - s_1: cosine similarity between their embeddings in the controlled setting
  - s_2: cosine similarity between their embeddings in the experimental setting
  - s_3: cosine similarity in the experimental setting after applying our mitigation method

- Positive values of $S_2 - S_1$: Item pairs with strong spurious associations have higher cosine similarities in the experimental setting than in the controlled setting.
- Negative values of $S_3 - S_2$: The reliable item-item semantic graph derived from the multimodal information can effectively counteract the adverse impact of spurious second-order item-item associations caused by shared accounts during graph convolution.

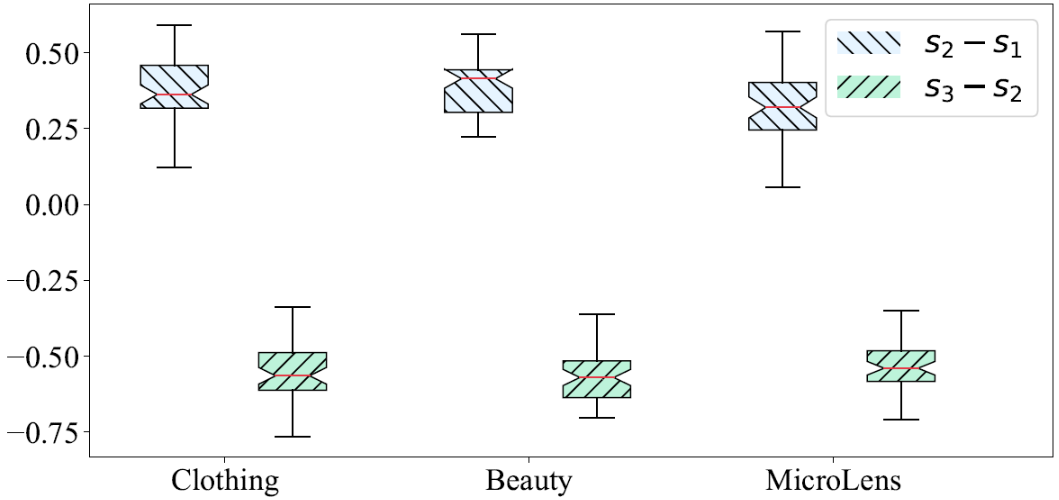

Conclusion & Future Work

Main Findings

- We reveal an unexplored problem that shared accounts can cause spurious associations between items, and adversely affect the recommendation performance of CF algorithms for other single-user accounts.
- A simple yet effective mitigation strategy is proposed to counteract spurious second-order item-item associations by leveraging the semantic similarity embodied in the multimodal information of items.

Future Work

- Explore more advanced multimodal recommendation methods that can leverage the valuable information from multiple modalities, and meanwhile avoid the potential negative effect of fake co-occurrences in the collaborative signal and noise in the multimodal features.
- Extend our methodology for analyzing the distortion of proximity relations between item embeddings to other research topics in robust RS (*e.g.*, attacking, defensing, denoising).

Thank you for you attention!

Q&A