Towards Explainable Collaborative Filtering with Taste Clusters Learning

Yuntao Du[1], Jianxun Lian[2], Jing Yao[2], Xiting Wang[2], Mingqi Wu[3], Lu Chen[1], Yunjun Gao[1] and Xing Xie[2]

[1]*College of Computer Science, Zhejiang University, China*

[2] *Microsoft Research Asia, China*

[3]*Microsoft Gaming, USA*

Motivations

◻ Recommender Systems

- ➢ Help billions of users to make decisions related to their personal lives
- ➢ Collaborative filtering (CF) is the main approach for recommendation

Choose a restaurant

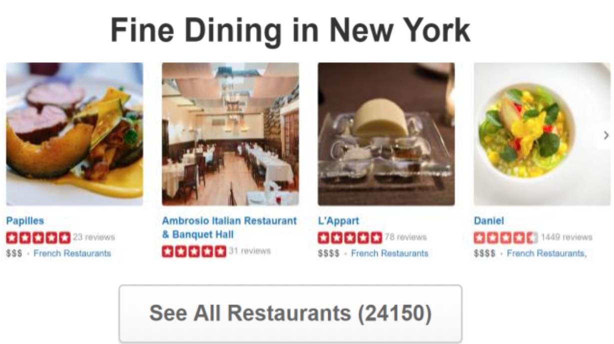

Buy products

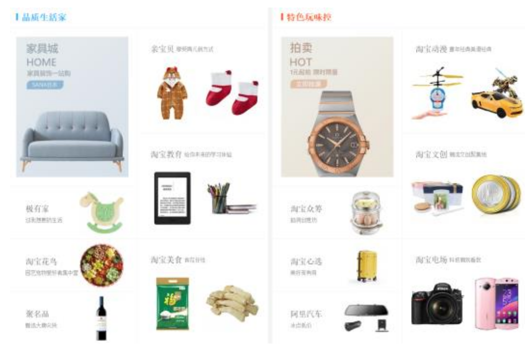

- ➢ A growing need to ensure that the users understand and trust the system

Motivations

Explainable Recommendation

- A growing need to ensure that the users understand and trust the system
- Explanations serve as a **bridge** between recommender systems and users
  - Increase user trust
  - Help users make better decisions (satisfactions)
  - Persuade users to try or buy an item (persuasiveness)
  - Assisting developers in model debugging and abnormal case studies

Explanations: why the items are recommended

Existing Methods – Post Hoc

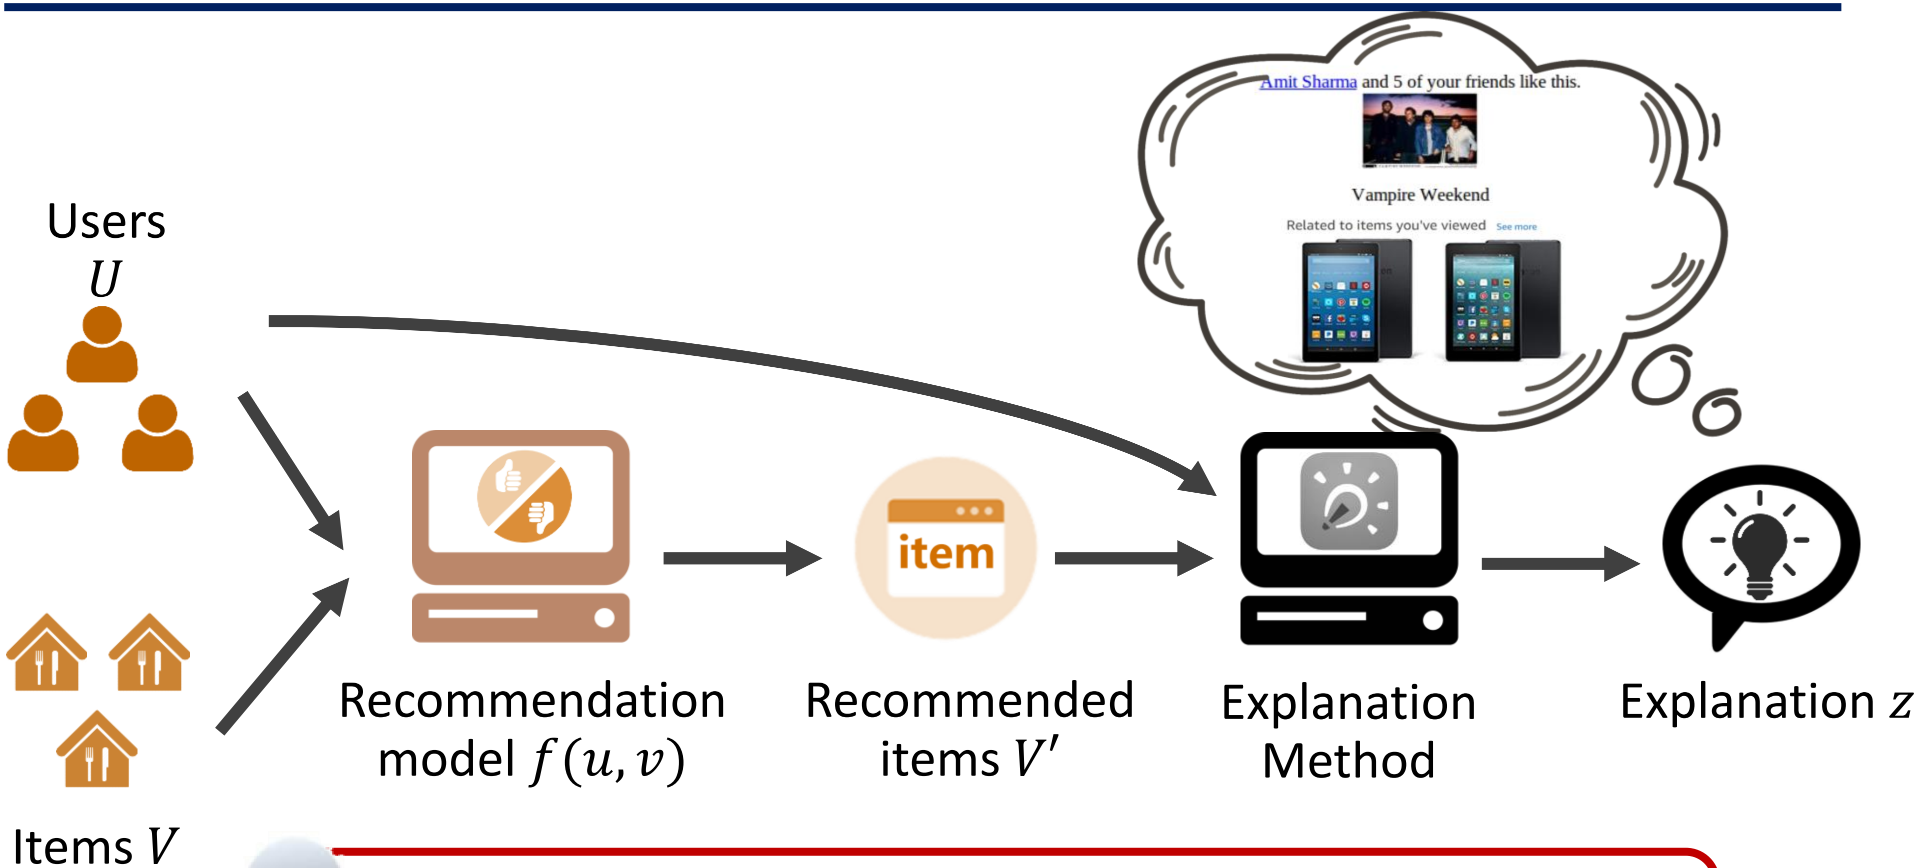

Ignore the working mechanism of the model
Failed to integrated for decision making process of the model

Existing Methods - Embedded

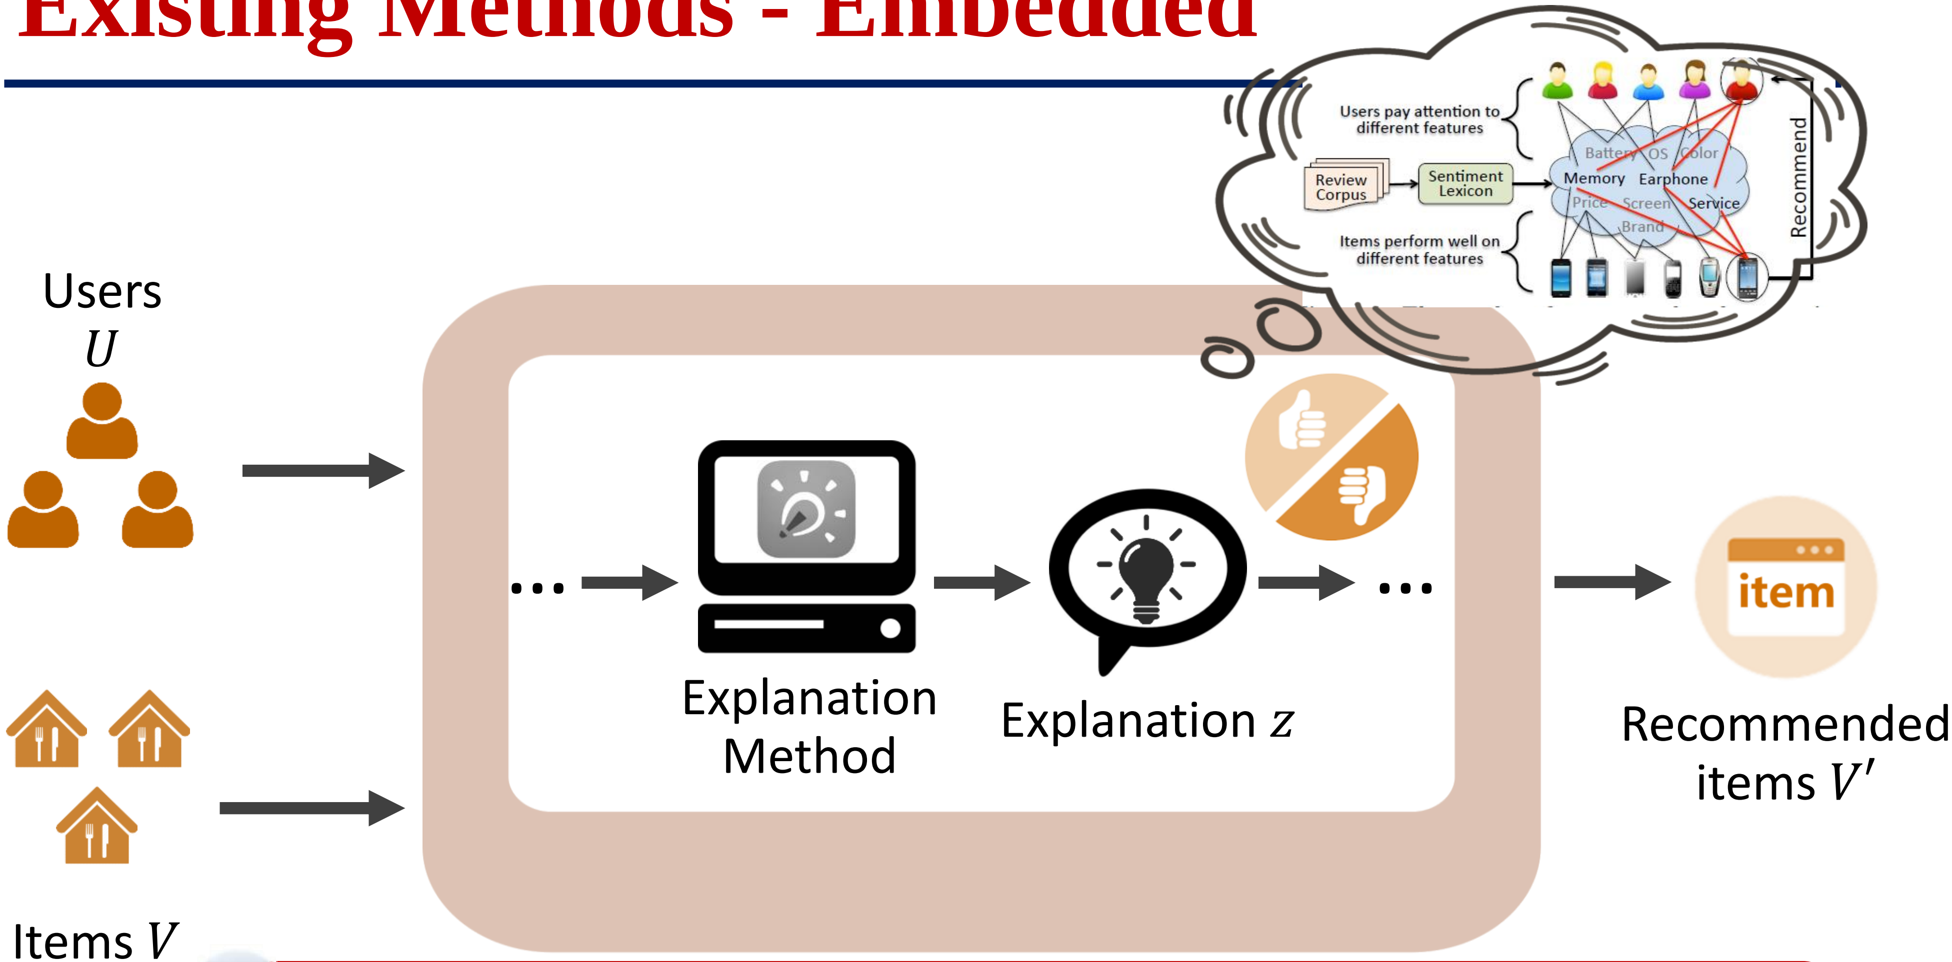

Lack of flexibility
At the cost of recommendation accuracy

Desirable Properties for Explainable CF

- Flexibility
  - The dimension of latent embeddings and the number of interpretable features/topics do not necessarily match (RecSys'13 fails on this)

- Coherence
  - A model's interpretable modules and predictive modules should be aligned during predictive decision making rather than being decoupled as independent modules (SIGIR'15 fails on this)

- Intrinsic explainability
  - A model can provide interpretable clues that truly reveal the model's running logic, instead of learning a post-hoc model for explanation (IJCAI'18 fails on this)

Our Method

Explainable Collaborative Filtering (ECF)

- The first framework that satisfy all three properties
- Core idea: mining various taste clusters, and map users/items to corresponding clusters
- A group of items which are not only similar in users' latent interest space, but also explicitly share some common tags

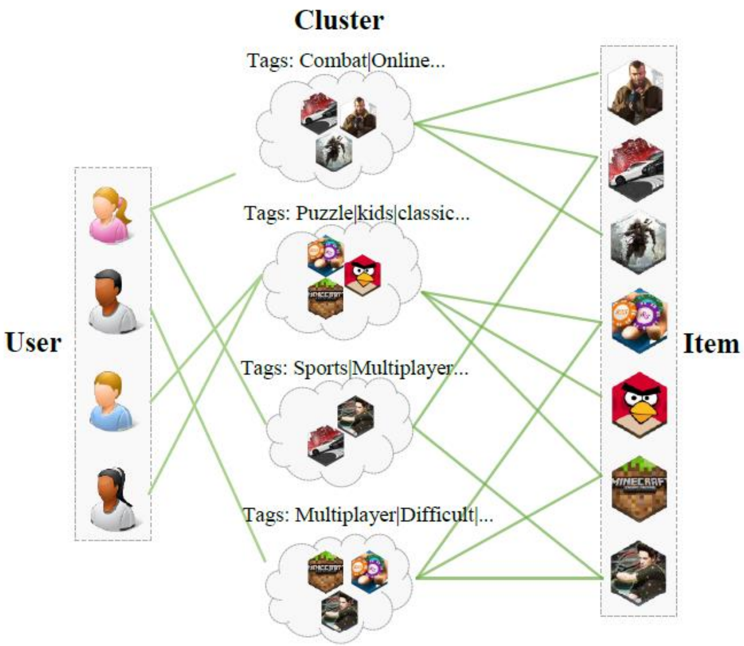

Recommendation process of ECF

Item recommendation

- Prediction score of user u and item i can be calculated by multiplying their affiliations with taste clusters

$$\hat{y}_{ui} = \text{sparse\_dot}(\mathbf{a}_u, \mathbf{x}_i),$$

Personalized explanation

- For each prediction $\hat{y}_{ui}$, ECF is able to generate explanation by measuring the coherence between users' and items' taste cluster affiliations:

$$C_{ui} = S(\mathbf{a}_u) \cap S(\mathbf{x}_i),$$

- And importance score w_{ui}^c is introduced to quantify the contribution of each taste cluster in C_{ui}:

$$w_{ui}^c = a_{uc} \times x_{ic}.$$

Learning Sparse Affiliation

- Directly learning the affiliation matrix from data is hard
  - Due to its sparsity nature for readability

- Initialize the users/items and taste clusters with embedding

$$\tilde{x}_{ic} = \cos(\mathbf{v}_i, \mathbf{h}_c),$$

$$m_{ic} = \begin{cases} 1 & \text{if } c \in \text{argTopm}(\tilde{\mathbf{x}}_i) \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbf{x}_i = \sigma(\tilde{\mathbf{x}}_i) \odot \mathbf{m}_i,$$

- Learn it with reparameterized trick

$$m_{ic} \approx \tilde{m}_{ic} = \frac{\exp(\cos(\mathbf{v}_i, \mathbf{h}_c)/T)}{\sum_c \exp(\cos(\mathbf{v}_i, \mathbf{h}_c)/T)},$$

$$\hat{m}_{ic} = \tilde{m}_{ic} + \text{detach\_gradient}(m_{ic} - \tilde{m}_{ic}),$$

Optimization of ECF

Reconstruction Loss

- Using user/item-cluster affiliations for prediction:

$$\mathcal{L}_{\text{CS}} = \sum_{(u,i,j)\in O} -\ln \sigma(\hat{y}_{ui} - \hat{y}_{uj}),$$

Tag Similarity Loss

- The items in the same taste clusters should share the similar tags
- Using TF-IDF style to select informative tags for taste clusters:

$$d_{ct} = \tilde{d}_{ct} \times \log(\frac{N}{f_t + \epsilon}), \qquad \beta_{ct} = \frac{\exp(d_{ct}/\tau)}{\sum_{c_j \in \mathcal{T}} \exp(d_{ct}/\tau)},$$

- Maximizing the likelihood of the probabilities of Top-P tags so that the taste clusters can be easily interpreted by those tags:

$$\mathcal{L}_{\text{TS}} = \sum_{c\in C} \sum_{t\in \text{argTopP}(\beta_c)} -\log \beta_{ct},$$

Independence Loss

- Taste clusters should be different to present different user interest space:

$$\mathcal{L}_{\text{IND}} = \sum_{c\in C} -\log \frac{\exp(s(\mathbf{h}_c, \mathbf{h}_c))}{\sum_{c'\in C} \exp(s(\mathbf{h}_c, \mathbf{h}_{c'}))},$$

Metrics for Explainability

In-cluster item coverage

- The proportion of items in the taste cluster that the selected tags can cover

$$\text{Cov.} = \frac{1}{Z} \sum_{c \in C} \sum_{i \in c} \frac{\mathbb{1}(\mathcal{T}_i \cap \mathcal{T}_c)}{|c|},$$

Tag utilization

- how many unique tags are used for interpreting taste clusters

$$\text{Util.} = \frac{1}{|\mathcal{T}|} \bigcup_{c \in C} \mathcal{T}_c,$$

Silhouette

- Similarity difference between intra-cluster items and inter-cluster items

$$\text{Sil.} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \frac{b(i) - a(i)}{\max\{a(i), b(i)\}},$$

Informativeness

- Distinctiveness of selected tags to represent the items in the taste cluster

$$\text{Info.} = \frac{1}{|C|} \sum_{c_i \in C} \frac{|R(\mathcal{T}_c) \cap c|}{|c|},$$

Experimental Evaluation

Datasets

- Industrial datasets (Xbox) and real-world datasets (MovieLens and Last-FM)

Dataset	#Users	#Items	#Interactions	#Tags
Xbox	465,258	330	6,240,251	115
MovieLens	6,033	3,378	836,434	18
Last-FM	53,486	2,062	2,228,949	54

Recommendation performance

- Achieve excellent accuracy performance while providing interpretability
- Our method greatly outperforms the baseline in all metrics across all datasets

	Xbox		MovieLens		Last-FM	
	Recall	NDCG	Recall	NDCG	Recall	NDCG
MF	0.5048	0.3268	0.1603	0.2416	0.0658	0.0506
NCF	0.4746	0.2931	0.1606	0.2406	0.0618	0.0401
CDAE	<u>0.5192</u>	<u>0.3286</u>	0.1627	0.2499	0.0589	0.0534
LightGCN	0.4933	0.3261	<u>0.1854</u>	<u>0.2698</u>	<u>0.0788</u>	<u>0.0675</u>
EFM	0.5070	0.3312	0.1702	0.2525	0.0703	0.0549
AMCF	0.5036	0.3217	0.1604	0.2405	0.0675	0.0516
ECF_{single}	0.4231	0.2331	0.1068	0.1501	0.0467	0.0380
ECF	**0.5922**†	**0.3721**†	**0.2124**†	**0.2903**†	**0.0851**†	**0.0773**†

Experimental Evaluation

Explainability

- **K-means**: similarity-oriented method which utilizes item embedding from FM to perform K-means algorithm
- **TagCluster**: tag-oriented method which collects items with the same tags

Method	Cov.	Util.	Sil.	Info.	Overall
Xbox					
ECF	<u>0.8002</u>	**0.7052**	<u>0.2604</u>	**0.3162**	**1.7463**
TagCluster	**0.9950**	0.2878	-0.1788	<u>0.1579</u>	0.9262
K-means	0.5710	<u>0.3739</u>	**0.4286**	0.0185	<u>1.0563</u>
Random	0.5396	0.1450	-0.3614	0.0125	0.0000
MovieLens					
ECF	<u>0.7992</u>	**0.7778**	<u>0.1964</u>	**0.3131**	**1.5651**
TagCluster	**0.991**	<u>0.5259</u>	-0.2573	<u>0.1517</u>	0.8898
K-means	0.6877	0.4478	**0.3265**	0.0168	<u>0.9573</u>
Random	0.5933	0.3672	-0.4452	0.0061	0.0000
Last-FM					
ECF	<u>0.7648</u>	**0.6259**	<u>0.1584</u>	**0.2996**	**1.5352**
TagCluster	**0.9880**	0.3703	-0.2511	<u>0.1206</u>	0.9143
K-means	0.5667	<u>0.4841</u>	**0.3197**	0.0182	<u>1.0752</u>
Random	0.5385	0.2275	-0.4673	0.0148	0.0000

- ECF takes all aspects into consideration so that it can avoid obvious shortcomings on a certain metric

Case Study

Learned Taste Clusters

- Can be used to correct tags

Taste Cluster c_1	Taste Cluster c_2	Taste Cluster c_3	Taste Cluster c_4	Taste Cluster c_5	...
pop|love|chill|soul	rap|hip-pop|pop|dance	soundtrack|ambient|alternative|epic	female_vocalists|soul|pop|indie	folk|singer_songwriter|acoustic|pop	...
• Kiss Me • How Do I Live • I Will Be • It Will Rain • He Won't Go ...	• Juice Box • Lose Yourself • Stronger • Men in Black • The Way I Are ...	• Panoramic • Hand Covers Bruise • Penetration • You Never Can Tell • Only Hope ...	• Bubbly • Hometown Glory • Daydreamer • Broken-Hearted Girl • Don't Be A Stranger ...	• Forever Young • Sensing Owls • Bubbly • The Kid • Honeymoon Child ...	...

Explanations of the recommendation

Listing History of u_{71}

- Penetration
- Hand Covers Bruise
- In My Place
- Homwtown Glory
- He Won't Go
- My Hear Will Go On
- Laura ...

Item: i_{77414} ("Bubbly")

Tags: female_vocalists|pop|folk...

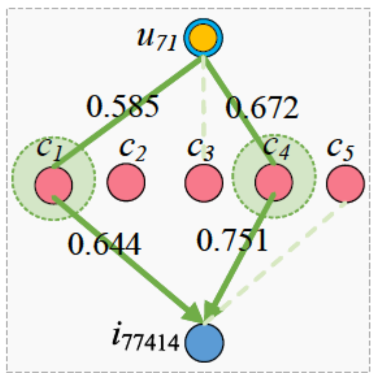

More Potentials of ECF

Taste Cluster Recommendation

- A new but ubiquitous recommendation task like playlist recommendation in Spotify or gamelist recommendation in Xbox

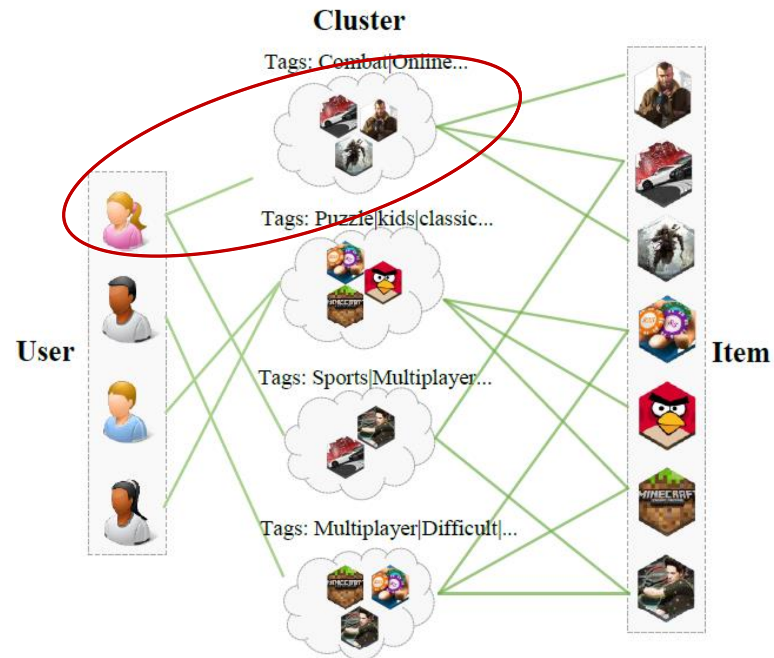

User Profiling

- User-cluster affiliations discovered by ECF can also be used as user profiles directly
- Can be used for user-level predictive tasks, ad audience targeting and look-alike audience extension, etc.

Flexibility

- Applied with other popular embedding-based methods like LightGCN

	Performance		Explainability				
	R@20	N@20	Cov.	Util.	Sil.	Info.	Overall
ECF	0.0851	0.0773	0.7648	0.6259	0.1584	0.2996	1.5352
ECF_{LGN}	**0.0876**	**0.0792**	**0.7831**	**0.6430**	**0.1590**	**0.3042**	**1.5758**

Conclusions

- A neat yet effective explainable framework. ECF leverages interpretable taste clusters and sparse user- and item-cluster affiliations for recommendation in a flexible, coherent, and Intrinsic explainability way

- Optimization for ECF. We present a method ECF to learn high quality taste clusters with informative tags and sparse affiliations simultaneously in an end-to-end manner.

- Quantitative metrics. Comprehensive analysis on the explainability quality of taste clusters.

- Extensive experiments. Considerable experimental results demonstrate the superiority of ECF, and it has been deployed into real recommendation scenarios in Xbox[1].

[1] https://www.xbox.com/en-US/xbox-game-pass

Thank you !

Questions?

ytdu@zju.edu.cn

Towards Explainable Collaborative Filtering with Taste Clusters Learning

https://github.com/zealscott/ECF

WWW 2023