

Disentangled Modeling of Preferences and Social Influence for Group Recommendation

Guangze Ye^{1, 2, 3}, Wen Wu^{3, 4*}, Guoqing Wang³, Xi Chen⁴, Hong Zheng⁵, Liang He^{1, 2, 3}

¹Lab of Artificial Intelligence for Education, East China Normal University, Shanghai, China

²Shanghai Institute of Artificial Intelligence for Education, East China Normal University, Shanghai, China

³School of Computer Science and Technology, East China Normal University, Shanghai, China

⁴Shanghai Key Laboratory of Mental Health and Psychological Crisis Intervention, School of Psychology and Cognitive Science, East China Normal University, Shanghai, China

⁵Shanghai Changning Mental Health Center, Shanghai, China

{gzye, wgq}@stu.ecnu.edu.cn, ww@cc.ecnu.edu.cn, xchen@psy.ecnu.edu.cn, zhhmm2@163.com, lhe@cs.ecnu.edu.cn

Abstract

The group recommendation (GR) aims to suggest items for a group of users in social networks. Existing work typically considers individual preferences as the sole factor in aggregating group preferences. Actually, social influence is also an important factor in modeling users' contributions to the final group decision. However, existing methods either neglect the social influence of individual members or bundle preferences and social influence together as a unified representation. As a result, these models emphasize the preferences of the majority within the group rather than the actual interaction items, which we refer to as the preference bias issue in GR. Moreover, the self-supervised learning (SSL) strategies they designed to address the issue of group data sparsity fail to account for users' contextual social weights when regulating group representations, leading to suboptimal results. To tackle these issues, we propose a novel model based on **Disentangled Modeling of Preferences and Social Influence for Group Recommendation (DisRec)**. Concretely, we first design a user-level disentangling network to disentangle the preferences and social influence of group members with separate embedding propagation schemes based on (hyper)graph convolution networks. We then introduce a social-based contrastive learning strategy, selectively excluding user nodes based on their social importance to enhance group representations and alleviate the group-level data sparsity issue. The experimental results demonstrate that our model significantly outperforms state-of-the-art methods on two real-world datasets.

Introduction

Humans, as social animals, have a strong desire to participate in group activities. Consequently, a task called GR has emerged (Hu et al. 2014). The GR aims to simulate the group decision-making process and suggest items for a group of users in social networks (Xiao et al. 2017). In existing works (Cao et al. 2018; Sankar et al. 2020; Wang et al. 2020; Zhang et al. 2021; Jia et al. 2021; Wu et al. 2023; Li et al. 2023; Ye et al. 2024), GR generally consists of three steps: (1)

*The corresponding author.

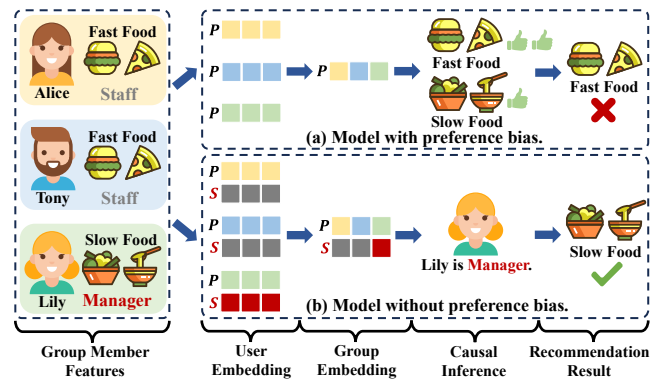


Figure 1: (a) and (b) show the causal inference process of the model with and without preference bias, where P and S are preferences and social embedding, respectively.

Training individual preferences; (2) Aggregating individual preferences to obtain group preferences; and (3) Optimizing group preferences. For example, AGREE (Cao et al. 2018) aggregates member preferences through attention mechanisms, while CubeRec (Chen et al. 2022) utilizes hypercubes to aggregate user representations. Although these methods have achieved satisfactory results, they still suffer from the following two limitations.

On one hand, these methods overly focus on individual preference factors, thereby neglecting other important factors affecting group recommendations, such as social influence, which helps quantify and differentiate the contributions of group members to a group decision (Zheng 2018; Yin et al. 2019). Although some studies (Yin et al. 2019; Guo et al. 2020) consider social information, they fail to disentangle user preferences and social influence but rather bundle them together as a unified representation. Such entanglement leads the model to emphasize the preferences of the majority within the group, which we refer to as the preference bias issue in GR. We provide an example in Figure 1 to further illustrate this bias. Alice and Tony, who are colleagues, prefer fast food, while Lily, their manager, prefers slow food. In (a), the model fails to disentangle preferences




Figure 2: The casual graph of GR.

and social influence, leading to incorrect recommendations of fast food favored by the majority in the group. In (b), the model assigns group members with separate embedding for preferences and social influence, avoiding preference bias issues. Consequently, the model correctly infers that the group’s choice of slow food is due to Lily’s greater influence as a manager. Disentangled user representations thus help the model accurately determine each member’s contribution to group decisions, eliminating preference bias in GR. We illustrate the causal relationships of members’ contributions through Figure 2.

On the other hand, most methods suffer from the group-item interaction data sparsity issue (Sankar et al. 2020). Although a few methods alleviate this issue by designing different SSL objectives (Sankar et al. 2020; Zhang et al. 2021; Li et al. 2023), these SSL strategies fail to account for users’ contextual social weights when regulating group representations, resulting in suboptimal performance. For example, HHGR (Zhang et al. 2021) and SGGCF (Li et al. 2023) achieve data augmentation by randomly removing member nodes from the group and aligning the group representations before and after removal. Since members contribute differently to the group’s final decisions, randomly removing member nodes can alter the decision-making process. In this scenario, such an alignment strategy is harmful.

To address these challenges, we propose a novel model based on **Dis**entangled Modeling of Preferences and Social Influence for Group **Re**commendation (DisRec). Firstly, we design an innovative user-level disentangling network to tackle the preference bias issue in GR. Specifically, we start by devising a disentangled embedding layer to separately encode user preferences and social influence, with two individual graph structures developed to capture their respective features. To comprehensively capture users’ social influence, we innovatively integrate the user-user social graph with the user-item hypergraph (Bai, Zhang, and Torr 2021; Feng et al. 2019) into a unified social hypergraph, employing an attention convolution mechanism for embedding propagation. Next, we concatenate the disentangled embedding to form final representations for users and items, deriving user-level group representations through an attention aggregator.

Secondly, we construct a group co-occurrence graph to capture the inherent preferences of the group independently of individual preferences, inspired by previous works (Cao et al. 2018; Jia et al. 2021; Wu et al. 2023). We then use a gate aggregator to merge user-level and group-level representations into the final group representations.

Thirdly, we introduce a novel social-based contrastive learning strategy to enhance group representations and alleviate the group-level data sparsity issue. Unlike previous

methods that randomly remove user nodes (Zhang et al. 2021; Li et al. 2023), we selectively exclude user nodes based on their contextual social importance. Our key technical insight is that only highly influential users significantly affect group decisions. Thus, we remove the most influential user to maximize the divergence between the resultant and original group representations, and the least influential user to minimize this divergence.

To summarize, the contributions are listed as follows:

- We are the first to identify the issue of preference bias in GR and to formulate a causal graph for GR, where the group’s final decision stems from two independent factors, namely user preferences and social influence.
- We propose a novel method, DisRec, which addresses the issue of preference bias in GR by disentangling user representations. Besides, it introduces a social-based SSL strategy that alleviates group-level data sparsity.
- We conduct extensive experiments on two real-world datasets to verify the superiority of DisRec over state-of-the-art baselines. The code for DisRec is available at (<https://github.com/DisRec/DisRec>).

Related Work

Group Recommendation

As mentioned before, existing GR methods first learned member preferences from U-I interaction data, then aggregated these member preferences into group preferences, and finally optimized the group preferences using group-item interaction data. Early methods used predefined aggregation strategies, such as average (Rakesh, Lee, and Reddy 2016), least misery (Baltrunas, Makcinskas, and Ricci 2010), and maximum pleasure (Boratto and Carta 2011). Later works (Cao et al. 2018; Vinh Tran et al. 2019; He, Chow, and Zhang 2020; Yin et al. 2019; Guo et al. 2020; Deng et al. 2021; Zhang, Zheng, and Wang 2022; Ji et al. 2023) adopted attention networks to explicitly model the importance of group members. Recent research employed more complex structures to achieve better user and group representations, such as hypercubes (Chen et al. 2022) and hypergraphs (Guo et al. 2021; Zhang et al. 2021; Jia et al. 2021; Wu et al. 2023). However, these methods either ignored the social influence of individual members or used a unified representation for users, thereby failing to disentangle user preferences and social influence. Furthermore, other methods attempted to design SSL strategies to address the sparsity of group data in GR (Zhang et al. 2021; Sankar et al. 2020; Li et al. 2023; Chen et al. 2022). For example, SGGCF (Li et al. 2023) created different views by randomly dropping out user nodes and edges, which cannot guarantee the quality of generated positive and negative examples.

Disentangled Representation Learning in RS

There were some recent studies (Zheng et al. 2021; Wang et al. 2022; Li et al. 2022; Wu et al. 2022; Chen et al. 2023; Fang et al. 2024) that utilized disentangled representation learning to alleviate bias in recommendation systems. For example, DICE (Zheng et al. 2021) addressed popularity

bias by disentangling user interests and conformity, while DISGCN (Li et al. 2022) focused on separating social homophily and influence in social recommendations. In this work, we disentangle users’ preferences and social influence to eliminate preference bias in GR.

Preliminaries

Let $\mathcal{U} = \{u_1, u_2, \dots, u_M\}$, $\mathcal{I} = \{i_1, i_2, \dots, i_N\}$ and $\mathcal{G} = \{g_1, g_2, \dots, g_K\}$ to represent the sets of M users, N items and K groups, respectively. There are three kinds of observed interaction data among \mathcal{U} , \mathcal{I} and \mathcal{G} : user-item interactions, group-item interactions and user-user interactions. We use $\mathbf{Y}_U \in \mathbb{R}^{M \times N}$ to denote the user-item interactions, where the element $\mathbf{Y}_U(i, j) = 1$ if user u_i has interacted with item i_j otherwise $\mathbf{Y}_U(i, j) = 0$. Likewise, we use $\mathbf{Y}_G \in \mathbb{R}^{K \times N}$ and $\mathbf{R} \in \mathbb{R}^{M \times M}$ to denote the group-item interactions and user-user interactions, respectively. The t -th group $g_t \in \mathcal{G}$ consists of a set of users, and we use $\mathcal{G}_t = \{u_1, u_2, \dots, u_s, \dots, u_{|\mathcal{G}_t|}\}$ to denote the user set of g_t , where $u_s \in \mathcal{U}$, and $|\mathcal{G}_t|$ is the size of group g_t . We denote the set of items which are interacted by g_t as $\mathcal{Y}_t = \{i_1, i_2, \dots, i_j, \dots, i_{|\mathcal{Y}_t|}\}$, where $i_j \in \mathcal{I}$, $|\mathcal{Y}_t|$ is the size of \mathcal{Y}_t . Then, given a target group g_t (or target user u_j), our goal is to generate a ranked list of items that g_t (or target user u_j) may be interested in.

Definition 1. Hypergraph. Let $\mathcal{G}^h = (\mathcal{V}^h, \mathcal{E}^h)$ denote a hypergraph, where \mathcal{V}^h is the vertex set and \mathcal{E}^h is the hyperedge set. Each hyperedge $e \in \mathcal{E}^h$ contains two or more vertices and is assigned a positive weight w_e . $\mathbf{W} \in \mathbb{R}^{|\mathcal{E}^h| \times |\mathcal{E}^h|}$ is the diagonal matrix of the weight of the hyperedge. The hypergraph can be represented by an incidence matrix $\mathbf{H} \in \mathbb{R}^{|\mathcal{V}^h| \times |\mathcal{E}^h|}$, where $h_{i,e} = 1$ if the hyperedge $e \in \mathcal{E}^h$ contains a vertex $v_i \in \mathcal{V}^h$, otherwise 0. For each hypergraph, \mathbf{D} is the diagonal degree matrices of vertex, and \mathbf{B} is the diagonal degree matrices of hyperedge.

Definition 2. Social Hypergraph. On the basis of Definition 1, we further define the social hypergraph. The social hypergraph is a heterogeneous graph that includes two sets of edges: hyperedges and social edges. Let $\mathcal{G}^s = (\mathcal{V}^s, \mathcal{E}^h, \mathcal{E}^s)$ denote a social hypergraph, where $\mathcal{E}^s = \{e_s = \{v_i, v_j\} \mid v_i, v_j \in \mathcal{V}^s, v_i \neq v_j\}$ is the social edge set.

Definition 3. Co-occurrence Graph. The co-occurrence graph uses hyperedges from a hypergraph as nodes. If two hyperedges share at least one common node, there is an edge between them in the co-occurrence graph. Formally, given a hypergraph $\mathcal{G}^h = (\mathcal{V}^h, \mathcal{E}^h)$, a co-occurrence graph is denoted as $\mathcal{G}^g = (\mathcal{V}^g, \mathcal{E}^g)$ where $\mathcal{V}^g = \{e \mid e \in \mathcal{E}^h\}$ and $\mathcal{E}^g = \{(e_p, e_q) \mid e_p, e_q \in \mathcal{E}^h, |e_p \cap e_q| \geq 1\}$. Each edge is assigned a weight $w_{p,q} = \frac{|e_p \cap e_q|}{|e_p \cup e_q|}$, and the adjacency matrix is defined as $\mathbf{A} \in \mathbb{R}^{M \times M}$ where $a_{p,q} = w_{p,q}$.

Proposed Model

User-Level Disentangling Network

Graph Construction. To disentangle the modeling of group members’ preferences and social influence, we extract two types of graphs from U-I and G-I interactions: the user

preference graph and the user social hypergraph, as shown on the left side of Figure 3. Specifically, the construction details of the two graphs are as follows:

- **User Preference Graph:** To capture each user’s pure preferences, we construct the user preference graph $\mathcal{G}^p = (\mathcal{V}^p, \mathcal{E}^p)$, where $\mathcal{V}^p = \mathcal{U}$ represents the node set of users, and $\mathcal{E}^p = \{(u_t, i_j) \mid u_t \in \mathcal{U}, i_j \in \mathcal{I}, \mathbf{Y}_U(t, j) = 1\}$ denotes the edge set comprising interactions between users and items. The adjacency matrix $\hat{\mathbf{A}} \in \mathbb{R}^{(M+N) \times (M+N)}$ represents the structure of \mathcal{G}^p .
- **User Social Hypergraph:** To disentangle the social influence of members, we construct a social hypergraph $\mathcal{G}^s = (\mathcal{V}^s, \mathcal{E}^h, \mathcal{E}^s)$ where $\mathcal{V}^s = \mathcal{U}$ denotes the node set of users, $\mathcal{E}^h = \mathcal{G}$ denotes the hyperedge set, and $\mathcal{E}^s = \{e_s = \{v_i, v_j\} \mid v_i, v_j \in \mathcal{U}, v_i \neq v_j\}$ denotes the social edge set. For user u_i in group g_t , they are connected to other group members and interacted items via hyperedge $\mathcal{E}_t^h = \mathcal{G}_t \cup \mathcal{Y}_t$, and to their social friends via social edges $\mathcal{E}_{u_i}^s$.

Disentangled Embedding Layer. Unlike existing work, we model two distinct factors using two sets of embeddings for users and items. Specifically, we use two embedding matrices, $\mathbf{U}^P \in \mathbb{R}^{d \times |\mathcal{U}|}$ and $\mathbf{U}^S \in \mathbb{R}^{d \times |\mathcal{U}|}$, to respectively represent user preferences and social influence, where d is the embedding size. Similarly, to ensure consistency in the latent space, we assign embeddings for items that correspond to those of users: $\mathbf{I}^P \in \mathbb{R}^{d \times |\mathcal{I}|}$ and $\mathbf{I}^S \in \mathbb{R}^{d \times |\mathcal{I}|}$.

Preference Graph Propagation. We utilize graph neural networks (He et al. 2020) for embedding propagation to represent the pure preferences of group members for items by capturing the collaborative signals between users and items. Specifically, we feed the concatenation of user preference embeddings $\mathbf{U}^P \in \mathbb{R}^{d \times |\mathcal{U}|}$ and item preference embeddings $\mathbf{I}^P \in \mathbb{R}^{d \times |\mathcal{I}|}$ to the graph convolutional network, denoted as $\mathbf{X}_P^{(0)} = [\mathbf{U}^P; \mathbf{I}^P]$. Referring to the spectral graph convolution (He et al. 2020), we define our graph convolution in the l -th layer as:

$$\mathbf{X}_P^{(l+1)} = \hat{\mathbf{D}}^{-\frac{1}{2}} \hat{\mathbf{A}} \hat{\mathbf{D}}^{-\frac{1}{2}} \mathbf{X}_P^{(l)}, \quad (1)$$

where $\mathbf{X}_P^{(l)}$ denotes the node representations in the l -th layer and $\hat{\mathbf{D}}$ is the diagonal node degree matrix of adjacency matrix $\hat{\mathbf{A}}$. After passing $\mathbf{X}_P^{(0)}$ through L convolutional layers, we average the embeddings obtained at each layer to get the final representation $\hat{\mathbf{X}}_P = \frac{1}{L+1} \sum_{l=0}^L \mathbf{X}_P^{(l)} = [\hat{\mathbf{U}}^P; \hat{\mathbf{I}}^P]$. Here, $\hat{\mathbf{U}}^P$ encodes information about the types of items that users will be interested in, while $\hat{\mathbf{I}}^P$ encodes information about the types of users that will like the items.

Social Hypergraph Propagation. The ability of hyperedges to model group representations is widely considered in existing methods (Zhang et al. 2021; Wu et al. 2023). However, these methods are limited to using hyperedges only to represent group preferences, overlooking the social influence contained within hyperedges. Our method extracts social influence within groups from hyperedges and outside groups from the social graph. As for items, they can also

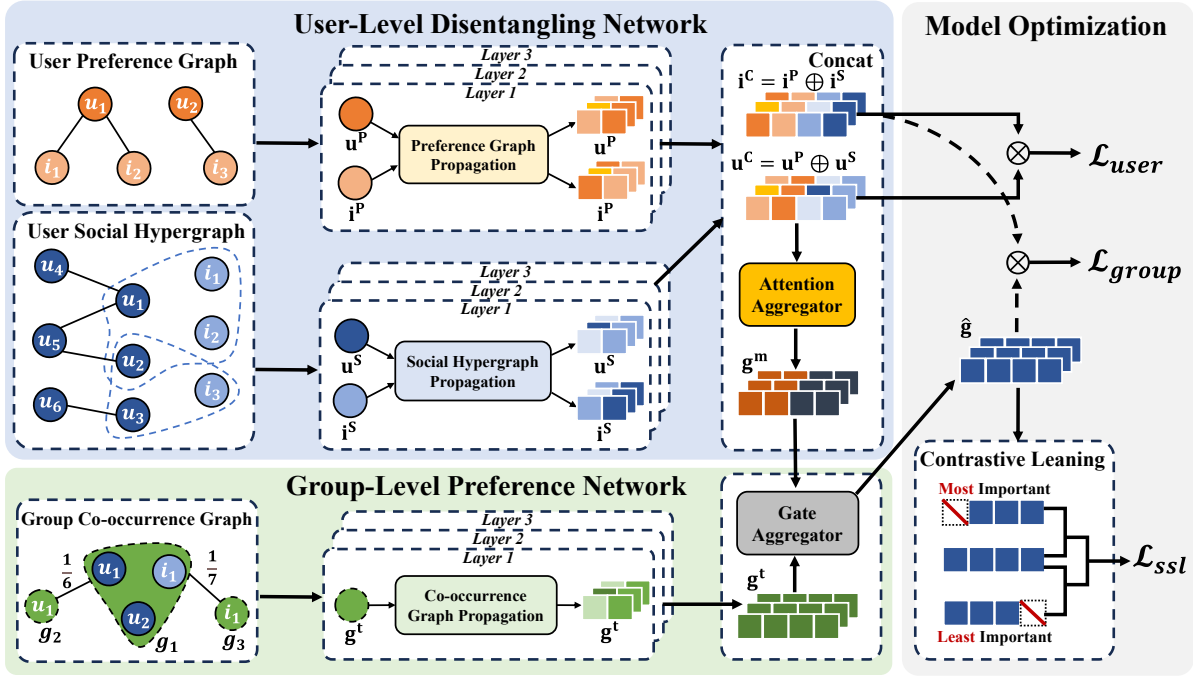


Figure 3: Framework of Disentangled Modeling of Preferences and Social Influence for Group Recommendation (DisRec).

fully leverage hyperedge representations to capture item features.

Therefore, we first apply the hypergraph convolution operation to extract users' social influence within groups and items' social features. Specifically, we feed the concatenation of user preference embeddings $\mathbf{U}^S \in \mathbb{R}^{d \times |\mathcal{U}|}$ and item preference embeddings $\mathbf{I}^S \in \mathbb{R}^{d \times |\mathcal{I}|}$ to the graph convolutional network, denoted as $\mathbf{X}_S^{(0)} = [\mathbf{U}^S; \mathbf{I}^S]$. Referring to the spectral hypergraph convolution (Feng et al. 2019; Wu et al. 2019), we define our hypergraph convolution as:

$$\mathbf{X}_S^{(l+1)} = \mathbf{D}^{-1} \mathbf{H} \mathbf{W} \mathbf{B}^{-1} \mathbf{H}^T \mathbf{X}_S^{(l)} \Psi^{(l)}, \quad (2)$$

where $\mathbf{X}_S^{(l)}$ is the node representations in the l -th layer and $\Psi \in \mathbb{R}^{d \times d}$ is the learnable weight matrix between two convolutional layers. In this process, the multiplication of \mathbf{H}^T and $\mathbf{X}_S^{(l)}$ extracts the social information of each group, while the multiplication with \mathbf{H} propagates this social information to each individual, forming their social influence within groups. We then perform a weighted sum of user friends' social influence within groups to generate users' social influence outside the group:

$$\mathbf{u}_{S,i}^{(l+1)} = \mathbf{u}_{S,i}^{(l)} + \sum_{j \in \mathcal{N}(i)} \alpha_j \mathbf{W}_1 \mathbf{u}_{S,j}^{(l)}, \quad (3)$$

$$\alpha_j = \frac{\exp(\mathbf{h}_1^T \mathbf{W}_1 \mathbf{u}_{S,j}^{(l)})}{\sum_{j' \in \mathcal{N}(i)} \exp(\mathbf{h}_1^T \mathbf{W}_1 \mathbf{u}_{S,j'}^{(l)})}, \quad (4)$$

where $\mathbf{u}_{S,i}^{(l)} \in \mathbf{U}^S$ denotes the node representation of user u_i in the l -th layer, α_j indicates the importance of the social

relation from user u_j to u_i , \mathbf{W}_1 and \mathbf{h}_1^T are learnable parameters, and $\mathcal{N}(i)$ is the social friends of user u_i . Finally, we average the embedding obtained at each layer to generate the final representation $\hat{\mathbf{X}}_S = \frac{1}{L+1} \sum_{l=0}^L \mathbf{X}_S^{(l)} = [\hat{\mathbf{U}}^S; \hat{\mathbf{I}}^S]$. Here, $\hat{\mathbf{U}}^S$ can be considered as a representation of how users will influence others in group decision-making process, while $\hat{\mathbf{I}}^S$ signifies the influence associated with items.

User-Level Representation Aggregation. We concatenate the thoroughly disentangled representations of preferences and social influences to obtain the final representations of users and items: $\hat{\mathbf{U}} = \hat{\mathbf{U}}^P \parallel \hat{\mathbf{U}}^S$, $\hat{\mathbf{I}} = \hat{\mathbf{I}}^P \parallel \hat{\mathbf{I}}^S$, where \parallel means concatenation of two embeddings.

Considering the varying contributions of individual users to the final group decision-making, we employ an attention aggregator to weightedly aggregate members' representations:

$$\mathbf{g}_t^m = \sum_{j \in \mathcal{G}_t} \beta_j \mathbf{W}_2 \hat{\mathbf{u}}_j, \quad (5)$$

$$\beta_j = \frac{\exp(\mathbf{h}_2^T \mathbf{W}_2 \hat{\mathbf{u}}_j)}{\sum_{j' \in \mathcal{G}_t} \exp(\mathbf{h}_2^T \mathbf{W}_2 \hat{\mathbf{u}}_{j'})}, \quad (6)$$

where \mathbf{g}_t^m denotes the user-level representation of group g_t , β_j indicates the weight of user u_j in the group decision, and \mathbf{W}_2 along with \mathbf{h}_2^T are learnable parameters. By stacking all the \mathbf{g}_t^m ($t = 1, \dots, K$), we can get the groups' user-level representations $\hat{\mathbf{G}}^m \in \mathbb{R}^{2d \times K}$.

Group-Level Preference Network

Co-occurrence Graph Propagation. As mentioned before, merely aggregating user-level representations is insufficient to adequately capture the inherent preferences of

the group. Since we only learn the group’s inherent preferences independently, we still apply graph neural networks to encode high-order relationships between groups based on group similarity. Specifically, we feed the group embeddings $\mathbf{G} \in \mathbb{R}^{2d \times K}$ to the graph convolutional network, denoted as $\mathbf{G}^{(0)} = \mathbf{G}$. The propagation mechanism at each layer is similar to Equation 1. After passing $\mathbf{G}^{(0)}$ through L convolutional layers, we average the embeddings obtained at each layer to get the final representation $\hat{\mathbf{G}}^t = \frac{1}{L+1} \sum_{l=0}^L \mathbf{G}^{(l)}$.

Group Gate Aggregation. We have obtained the user-level group representation $\hat{\mathbf{G}}^m$ and the group-level representation $\hat{\mathbf{G}}^t$. To adaptively control their combination, we propose a gate aggregator:

$$\hat{\mathbf{G}} = \gamma \hat{\mathbf{G}}^m + (1 - \gamma) \hat{\mathbf{G}}^t, \quad (7)$$

$$\gamma = \sigma(\mathbf{W}^m \mathbf{G}^m + \mathbf{W}^t \mathbf{G}^t + \mathbf{b}^g), \quad (8)$$

where \mathbf{W}^m , \mathbf{W}^t and \mathbf{b}^g are learnable parameters and σ is the activation function. The γ denote the learned weight that balanced the contributions of $\hat{\mathbf{G}}^m$ and $\hat{\mathbf{G}}^t$. By using the gate aggregator, we can obtain the final group representations $\hat{\mathbf{G}}$.

Model Optimization

To optimize the DisRec parameters, we unify the loss functions of user, group, and SSL. Firstly, we adopt the pairwise learning task loss function (Wang et al. 2017) to optimize user and item representations, which is designed as follows:

$$\mathcal{L}_{user} = - \sum_{(j,h,h') \in \mathcal{O}} (\hat{r}_{jh} - \hat{r}_{jh'} - 1)^2, \hat{r}_{jh} = \mathbf{u}_j \mathbf{i}_h^T, \quad (9)$$

where \mathbf{u}_j represents the embedding of user u_j , \mathbf{i}_h represents the embedding of item i_h . \mathcal{O} represents the training set, in which each one includes the user u_j , interacted item i_h , and unobserved items i'_h . Similarly, the group loss function can be achieved as follows:

$$\mathcal{L}_{group} = - \sum_{(t,h,h') \in \mathcal{O}'} (\hat{y}_{th} - \hat{y}_{th'} - 1)^2, \hat{y}_{th} = \mathbf{g}_t \mathbf{i}_h^T, \quad (10)$$

where \mathbf{g}_t represents the embedding of group g_t , \mathbf{i}_h represents the embedding of item i_h . \mathcal{O}' represents the training set, in which each one includes the group g_t , interacted item i_h , and unobserved items i'_h .

Contrastive learning. To overcome the long-standing issue of data sparsity in GR, we intend to leverage the rich self-supervised signals from social information to enhance group representations. Specifically, we identify the most important person u' and the least important person u'' in group g_t based on Equation 6. By removing u' and u'' respectively, we obtain new group representations g' and g'' . We treat (g_t, g') as a positive example pair and (g_t, g'') as a negative example pair. Our goal is to accurately reflect how changes in group membership impact group decisions. Following S²-HHGR (Zhang et al. 2021), we adopt crossentropy loss to construct the contrastive loss:

$$\mathcal{L}_{ssl} = - \sum_{t \in \mathcal{G}} [\log \sigma(f(\hat{\mathbf{g}}_t, \hat{\mathbf{g}}'_t)) + [\log \sigma(1 - f(\hat{\mathbf{g}}_t, \hat{\mathbf{g}}''_t))]], \quad (11)$$

Dataset	#Users	#Groups	#Items	#U-I interactions	#G-I interactions
Mafengwo	5,275	995	1,513	39,761	3,595
Yelp	34,504	24,103	22,611	482,273	26,883

Table 1: Statistics of datasets.

where $f(\hat{\mathbf{g}}_t, \hat{\mathbf{g}}'_t) = \sigma(\hat{\mathbf{g}}_t \mathbf{W}_c \hat{\mathbf{g}}'^T)$ is a learnable distance function, and $\hat{\mathbf{g}}_t$ is the representation of group g_t . We jointly train \mathcal{L}_{user} , \mathcal{L}_{group} and \mathcal{L}_{ssl} using the Adam algorithm. The overall objective is defined as follows:

$$\mathcal{L} = \mathcal{L}_{user} + \mathcal{L}_{group} + \delta \mathcal{L}_{ssl} \quad (12)$$

where δ is the hyper-parameter to balance the task of self-supervised learning and supervised learning task for GR.

Experiment

Experimental Settings

Datasets. We conduct experiments on two real-world public datasets: Mafengwo, published by (Cao et al. 2018), and Yelp, published by (Yin et al. 2019). Mafengwo is a tourism website where users create or join group travel activities. Yelp is a typical check-in dataset for various restaurants in the US. Table 1 reports the statistics of these datasets.

Evaluation Metrics. To evaluate the performance of the Top-K recommendation for all methods, we employ the widely adopted metrics NDCG@K and HR@K with $K=\{5, 10\}$, as in (Wu et al. 2023). All items are considered as candidates for fairness. The permutation test (Nichols and Holmes 2002) is adopted for significance tests.

Baselines. To evaluate the effectiveness of DisRec, we compared it with the following baselines: **NCF** (He et al. 2017) treats a group as a virtual user and ignores the member information of the group. **SIGR** (Yin et al. 2019) integrates users’ global and local social information to improve GR. **GroupIM** (Sankar et al. 2020) maximizes mutual information between users and groups to overcome the group-level data sparsity. **HCR** (Jia et al. 2021) proposes a dual channel hypergraph convolutional network to capture member-level and group-level preferences. **S²-HHGR** (Zhang et al. 2021) integrates hierarchical hypergraph learning and self-supervised learning for improved group representation. **CubeRec** (Chen et al. 2022) utilizes a hypercube vector space to replace point embeddings for representing group preferences. **SGGCF** (Li et al. 2023) explores two kinds of contrastive learning module to capture the implicit relations between groups and items. **ConsRec** (Wu et al. 2023) designs three views to capture the consensus information within groups.

Implementation Details. For the general settings, the embedding size is set to 64, the batch size is 512, and the number of negative samples is 10. For the baseline models, we refer to their best parameter setups reported in the original papers. For our model, we set the convolutional layer L to 3, the SSL weight δ to 0.5 and tune the learning rate in $[1e - 4, 1e - 3]$, and the dropout rate in $[0.1, 0.5]$.

Dataset	Metric	NCF	SIGR	HCR	GroupIM	S ² -HHGR	SGGCF	CubeRec	ConsRec	DisRec	Improv.
Mafengwo	HR@5	0.3681	0.4092	0.4824	0.5297	0.5726	0.6241	0.6824	<u>0.6925</u>	0.7266*	4.92%
	HR@10	0.4365	0.4583	0.5176	0.6040	0.5899	0.7005	0.7699	<u>0.7829</u>	0.8101*	3.47%
	NDCG@5	0.2798	0.3077	0.4557	0.4291	0.5405	0.5069	0.5572	<u>0.5675</u>	0.6045*	6.52%
	NDCG@10	0.3171	0.3347	0.4668	0.4534	0.5473	0.5320	0.5862	<u>0.5969</u>	0.6312*	5.75%
Yelp	HR@5	0.2622	0.3276	0.5086	0.5875	0.5379	0.5428	0.5553	<u>0.5904</u>	0.6157*	4.29%
	HR@10	0.3165	0.4274	0.5638	<u>0.6687</u>	0.5965	0.6074	0.6173	0.6606	0.7061*	5.59%
	NDCG@5	0.2890	0.2407	0.4065	0.4992	0.4115	0.4687	0.4602	<u>0.5051</u>	0.5441*	7.72%
	NDCG@10	0.3198	0.2729	0.4241	0.5056	0.4239	0.4867	0.4831	<u>0.5154</u>	0.5646*	9.55%

Table 2: Performance comparison of all methods on **group recommendation** task. The column of ‘‘Improv.’’ indicates the improvement of the best result of DisRec (boldface) compared with the best baseline (underscore). * indicates the statistical significance for p -value < 0.05 compared with the best baseline.

Dataset	Metric	NCF	SIGR	HCR	GroupIM	S ² -HHGR	SGGCF	CubeRec	ConsRec	DisRec	Improv.
Mafengwo	HR@5	0.4759	0.4623	<u>0.5804</u>	0.0969	0.4734	0.4513	0.1750	0.5736	0.6608*	13.85%
	HR@10	0.5603	0.5569	0.6186	0.1323	0.5580	0.5286	0.3002	0.6346	0.7020*	10.62%
	NDCG@5	0.4447	0.4273	<u>0.5456</u>	0.0686	0.3834	0.3042	0.1045	0.5436	0.5795*	6.21%
	NDCG@10	0.4702	0.4506	0.5775	0.0801	0.4110	0.3370	0.1442	<u>0.5829</u>	0.6115*	4.91%
Yelp	HR@5	0.4029	0.4153	0.4774	0.0998	0.4465	0.4310	0.1611	<u>0.4903</u>	0.5668*	15.60%
	HR@10	0.4991	0.5286	<u>0.5499</u>	0.1560	0.5168	0.5138	0.1957	0.5319	0.6245*	13.57%
	NDCG@5	0.3981	0.3042	0.4375	0.0652	0.3841	0.3253	0.1259	<u>0.4640</u>	0.5028*	8.36%
	NDCG@10	0.4516	0.3370	0.4870	0.0865	0.4772	0.3426	0.1371	<u>0.5083</u>	0.5468*	7.57%

Table 3: Performance comparison of all methods on **user recommendation** task.

Overall Performance

The experimental performance on GR task and user recommendation (UR) task are shown in Table 2 and Table 3, respectively. We observed that in both GR and UR tasks, our proposed DisRec consistently outperforms all baselines on two datasets in terms of HR@K and NDCG@K metrics.

For the GR task, our model gains the greatest advantage over single-strategy methods that use only hypergraphs or SSL (HCR, GroupIM), with average increases of 30.4% in H@10 and 29.8% in N@10 across two datasets. Compared to mixed-strategy methods that integrate both (hyper) graph and SSL (S²-HHGR, SGGCF), our improvements are slightly smaller but still significant, with average increases of 21.9% in H@10 and 20.8% in N@10 across two datasets. This is because DisRec not only designs superior social hypergraphs and more effective social-based SSL, but also overcomes the preference bias issue of previous models by disentangling user preferences and social influence. Additionally, compared to models that also use social information (SIGR and CubeRec), our model achieves improvements of at least 12.3% in N@10. This improvement may be attributed to our social hypergraph’s ability to better capture users’ social influence both within and outside their groups.

For the UR task, our model also demonstrates notable advantages. For instance, compared to the strongest baseline model, ConsRec, our model achieve an average improvement of 12.10% in H@10 and 6.24% in N@10 across two datasets. We attribute this to disentangled user representations, which help mitigate preference bias in GR as well as enhance the performance of the UR task.

Dataset	Metric	w/o. S	w/o. P	w/o. SSL	DisRec
Mafengwo	HR@10	0.7698	0.7888	<u>0.8050</u>	0.8101*
	NDCG@10	0.5916	0.6125	<u>0.6214</u>	0.6312*
Yelp	HR@10	0.6625	0.6869	<u>0.6942</u>	0.7061*
	NDCG@10	0.5267	0.5423	<u>0.5514</u>	0.5646*

Table 4: The comparison between DisRec and its variants.

Ablation Study

We adopt two sets of embeddings to disentangle user preferences and social influence and employ a social-based SSL strategy to enhance group representations. To verify the effectiveness of each main components of DisRec, we conduct the ablation study to analyze their contribution. Due to space limitations, Table 4 only shows the results on H@10 and N@10 for two datasets, where ‘‘w/o. S’’, ‘‘w/o. P’’, and ‘‘w/o. SSL’’ represent the variants of removing social influence embeddings, preferences embeddings, and the SSL strategy, respectively. From Table 4, we observe that removing any part results in a performance degradation across all datasets. For example, on the Yelp dataset, ‘‘w/o. SSL’’ shows a 2.3% decrease in NDCG@10 compared to DisRec. Furthermore, we observe that the variant ‘‘w/o. SSL’’ performs the best among all variants, indicating that disentangled user embeddings have a greater impact on the model performance. However, the variant ‘‘w/o. S’’ outperforms the variant ‘‘w/o. P’’, suggesting that solely modeling user preferences is insufficient and further emphasizing the importance of disentangled user contributions in GR.

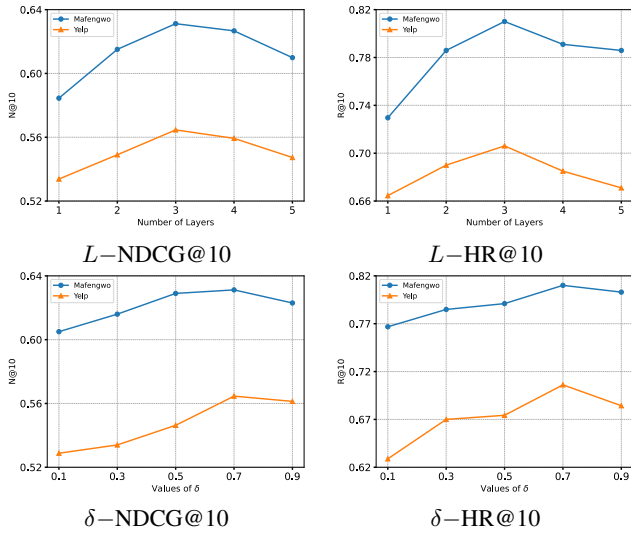


Figure 4: The influence of different model hyperparameters.

Hyperparameter Analysis


To investigate the impact of important hyperparameters in DisRec, we evaluate the number of graph convolution layers L and the balance coefficient of SSL task δ on two datasets. We adjust each hyperparameter while others remain unchanged, and the results are reported in Figure 4.

Impact of L . The performance of graph convolutional networks is influenced by the number of network layers L . We tune L in $\{1, 2, 3, 4, 5\}$. In general, DisRec benefits a larger L across all datasets, achieving the best performance with $L = 3$. Performance declines with excessive layers (e.g., $L = 5$) due to over-smoothing, while too few layers fail to capture complex higher-order relationships.


Impact of δ . The coefficient δ in Eq. (12) controls the SSL regularization effect. To analyze the impact of δ on our model, the coefficient δ is set among $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. DisRec’s performance increases as δ rises, until δ reaches 0.7, after which it decreases with further increases in δ . The worst performance occurs when $\delta = 0.1$, suggesting that utilizing our proposed social-based SSL strategy is necessary to address the issue of group data sparsity. However, excessively high δ values cause the model to overly focus on high-weight individuals, resulting in suboptimal outcomes.

Case Study

To further investigate the effectiveness of DisRec in addressing preference bias in GR, we conduct a case study using a synthetic dataset inspired by (Wang et al. 2021). Specifically, we generate two fake items for each positive group-item pair in the test set of the Mafengwo dataset, with these fake items interacted with by at least 2/3 of the group members. Models with preference bias often favor items aligning with the majority preferences, even if these are not the group’s final choices. We define $Rank_{gap} = Rank_{fake} - Rank_{true}$ to measure the model’s ability to overcome preference bias, where $Rank_{fake}$ and $Rank_{true}$ represent the ranks of the paired fake and real items, respectively. A model



(a) Distribution of $Rank_{gap}$ values.



(b) A visualized example from group g_{127} .

Figure 5: Results of DisRec evaluation on synthetic data.

that effectively mitigates preference bias should have a high $Rank_{gap}$. Figure 5(a) presents the $Rank_{gap}$ values for DisRec and the strongest baseline model, ConsRec, across 719 test samples. It is evident from the figure that DisRec has a lower $Count$ value near a $Rank_{gap}$ of 0 compared to ConsRec. The maximum $Rank$ value for ConsRec peaks around 400, while DisRec’s $Rank$ values are more evenly distributed between 400 and 1200. This is attributed to DisRec’s successful disentangling of user preferences and social influence causes. Additionally, Figure 5(b) visualizes a randomly selected test case $g_{127} - i_{232}$. Here, group g_{127} interacts with item i_{232} , while items i_{477} and i_{552} , which are popular among most members, are set as fake items. The right table shows that DisRec correctly assigns a higher rank to item i_{232} , whereas ConsRec is misled by the fake items, demonstrating DisRec’s effectiveness in mitigating preference bias.

Conclusion

In this paper, we address the issue of preference bias in previous GR models and propose an innovative model DisRec, which solves this problem by disentangling users’ preferences and social influence. Additionally, DisRec adopts a social-based SSL strategy to alleviate the issue of group data sparsity by generating higher quality positive and negative samples to regularize group representations. Through extensive experiments on two widely used datasets, we demonstrate the superiority of our model in both GR and UR tasks. Furthermore, a case study focused on the issue of preference bias confirms DisRec’s ability to distinguish between real and fake items.

Acknowledgments

This work is funded by STI 2030-Major Projects 2021ZD0200500, National Natural Science Foundation of China (under project No. 62377013), and the Fundamental Research Funds for the Central Universities. It is also supported by Natural Science Foundation of Shanghai, China (under project No. 22ZR1419000), the Research Project of Changning District Science and Technology Committee (under project No. CNKW2022Y37), and the Medical Master's and Doctoral Innovation Talent Base Project of Changning District (under project No. RCJD2022S07).

References

- Bai, S.; Zhang, F.; and Torr, P. H. 2021. Hypergraph convolution and hypergraph attention. *Pattern Recognition*, 110: 107637.
- Baltrunas, L.; Makičinskas, T.; and Ricci, F. 2010. Group recommendations with rank aggregation and collaborative filtering. In *Proceedings of the ACM Conference on Recommender Systems (Recsys)*, 119–126.
- Boratto, L.; and Carta, S. 2011. State-of-the-art in group recommendation and new approaches for automatic identification of groups. *Information Retrieval and Mining in Distributed Environments*, 1–20.
- Cao, D.; He, X.; Miao, L.; An, Y.; Yang, C.; and Hong, R. 2018. Attentive group recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 645–654.
- Chen, J.; Dong, H.; Wang, X.; Feng, F.; Wang, M.; and He, X. 2023. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, 41(3): 1–39.
- Chen, T.; Yin, H.; Long, J.; Nguyen, Q. V. H.; Wang, Y.; and Wang, M. 2022. Thinking inside The Box: Learning Hypercube Representations for Group Recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1664–1673.
- Deng, Z.; Li, C.; Liu, S.; Ali, W.; and Shao, J. 2021. Knowledge-aware group representation learning for group recommendation. In *International Conference on Data Engineering (ICDE)*, 1571–1582. IEEE.
- Fang, J.; Zhang, G.; Cui, Q.; Tang, C.; Gu, L.; Li, L.; Gu, J.; and Zhou, J. 2024. Backdoor Adjustment via Group Adaptation for Debaised Coupon Recommendations. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, 11944–11952.
- Feng, Y.; You, H.; Zhang, Z.; Ji, R.; and Gao, Y. 2019. Hypergraph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, 3558–3565.
- Guo, L.; Yin, H.; Chen, T.; Zhang, X.; and Zheng, K. 2021. Hierarchical hyperedge embedding-based representation learning for group recommendation. *ACM Transactions on Information Systems (TOIS)*, 40(1): 1–27.
- Guo, L.; Yin, H.; Wang, Q.; Cui, B.; Huang, Z.; and Cui, L. 2020. Group recommendation with latent voting mechanism. In *International Conference on Data Engineering (ICDE)*, 121–132.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 639–648.
- He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *Proceedings of the ACM Web Conference (WWW)*, 173–182.
- He, Z.; Chow, C.-Y.; and Zhang, J.-D. 2020. GAME: Learning graphical and attentive multi-view embeddings for occasional group recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 649–658.
- Hu, L.; Cao, J.; Xu, G.; Cao, L.; Gu, Z.; and Cao, W. 2014. Deep modeling of group preferences for group-based recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, volume 28.
- Ji, J.; Pei, J.; Lin, S.; Zhou, T.; He, H.; Jia, J.; and Hu, N. 2023. Multi-granularity attention model for group recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, 3973–3977.
- Jia, R.; Zhou, X.; Dong, L.; and Pan, S. 2021. Hypergraph convolutional network for group recommendation. In *International Conference on Data Mining (ICDM)*, 260–269. IEEE.
- Li, K.; Wang, C.-D.; Lai, J.-H.; and Yuan, H. 2023. Self-Supervised Group Graph Collaborative Filtering for Group Recommendation. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM)*, 69–77.
- Li, N.; Gao, C.; Jin, D.; and Liao, Q. 2022. Disentangled modeling of social homophily and influence for social recommendation. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 35(6): 5738–5751.
- Nichols, T. E.; and Holmes, A. P. 2002. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human Brain Mapping*, 15(1): 1–25.
- Rakesh, V.; Lee, W.-C.; and Reddy, C. K. 2016. Probabilistic group recommendation model for crowdfunding domains. In *Proceedings of the ACM International Conference on Web Search and Data Mining (WSDM)*, 257–266.
- Sankar, A.; Wu, Y.; Wu, Y.; Zhang, W.; Yang, H.; and Sundaram, H. 2020. Groupim: A mutual information maximization framework for neural group recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1279–1288.
- Vinh Tran, L.; Nguyen Pham, T.-A.; Tay, Y.; Liu, Y.; Cong, G.; and Li, X. 2019. Interact and decide: Medley of sub-attention networks for effective group recommendation. In *Proceedings of the International ACM SIGIR Conference*

- on *Research and Development in Information Retrieval (SIGIR)*, 255–264.
- Wang, W.; Feng, F.; He, X.; Zhang, H.; and Chua, T.-S. 2021. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1288–1297.
- Wang, W.; Zhang, W.; Rao, J.; Qiu, Z.; Zhang, B.; Lin, L.; and Zha, H. 2020. Group-aware long-and short-term graph representation learning for sequential group recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 1449–1458.
- Wang, X.; Chen, H.; Zhou, Y.; Ma, J.; and Zhu, W. 2022. Disentangled representation learning for recommendation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1): 408–424.
- Wang, X.; He, X.; Nie, L.; and Chua, T.-S. 2017. Item silk road: Recommending items from information domains to social users. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 185–194.
- Wu, F.; Souza, A.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. 2019. Simplifying graph convolutional networks. In *International Conference on Machine Learning (ICML)*, 6861–6871. PMLR.
- Wu, J.; Fan, W.; Chen, J.; Liu, S.; Li, Q.; and Tang, K. 2022. Disentangled contrastive learning for social recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 4570–4574.
- Wu, X.; Xiong, Y.; Zhang, Y.; Jiao, Y.; Zhang, J.; Zhu, Y.; and Yu, P. S. 2023. ConsRec: Learning Consensus Behind Interactions for Group Recommendation. In *Proceedings of the ACM Web Conference (WWW)*, 240–250.
- Xiao, L.; Min, Z.; Yongfeng, Z.; Zhaoquan, G.; Yiqun, L.; and Shaoping, M. 2017. Fairness-aware group recommendation with pareto-efficiency. In *Proceedings of the ACM Conference on Recommender Systems (Recsys)*, 107–115.
- Ye, G.; Wu, W.; Shi, L.; Hu, W.; Chen, X.; and He, L. 2024. A personality-guided preference aggregator for ephemeral group recommendation. *Applied Soft Computing*, 167: 112274.
- Yin, H.; Wang, Q.; Zheng, K.; Li, Z.; Yang, J.; and Zhou, X. 2019. Social influence-based group representation learning for group recommendation. In *International Conference on Data Engineering (ICDE)*, 566–577.
- Zhang, J.; Gao, M.; Yu, J.; Guo, L.; Li, J.; and Yin, H. 2021. Double-scale self-supervised hypergraph learning for group recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, 2557–2567.
- Zhang, S.; Zheng, N.; and Wang, D. 2022. GBERT: Pre-training User representations for Ephemeral Group Recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM)*, 2631–2639.
- Zheng, Y. 2018. Identifying Dominators and Followers in Group Decision Making based on The Personality Traits. In *IUI Workshops*.
- Zheng, Y.; Gao, C.; Li, X.; He, X.; Li, Y.; and Jin, D. 2021. Disentangling user interest and conformity for recommendation with causal embedding. In *Proceedings of the ACM Web Conference (WWW)*, 2980–2991.