

Dual Intents Graph Modeling for User-centric Group Discovery

Xixi Wu
Yun Xiong*
Yao Zhang
21210240043@m.fudan.edu.cn
yunx@fudan.edu.cn
yaozhang@fudan.edu.cn
Shanghai Key Laboratory of Data
Science, School of Computer Science,
Fudan University
Shanghai, China

Yizhu Jiao
yizhu2@illinois.edu
University of Illinois at
Urbana-Champaign
IL, USA

Jiawei Zhang
jiawei@ifmlab.org
IFM Lab, Department of Computer
Science, University of California,
Davis
CA, USA

ABSTRACT

Online groups have become increasingly prevalent, providing users with space to share experiences and explore interests. Therefore, user-centric group discovery task, *i.e.*, recommending groups to users can help both users' online experiences and platforms' long-term developments. Existing recommender methods can not deal with this task as modeling user-group participation into a bipartite graph overlooks their item-side interests. Although there exist a few works attempting to address this task, they still fall short in fully preserving the social context and ensuring effective interest representation learning.

In this paper, we focus on exploring the *intents* that motivate users to participate in groups, which can be categorized into different types, like the *social-intent* and the *personal interest-intent*. The former refers to users joining a group affected by their social links, while the latter relates to users joining groups with like-minded people for interest-oriented self-enjoyment. To comprehend different intents, we propose a novel model, **DiRec**, that first models each intent separately and then fuses them together for predictions. Specifically, for social-intent, we introduce the hypergraph structure to model the relationship between groups and members. This allows for more comprehensive group information preservation, leading to a richer understanding of the social context. As for interest-intent, we employ novel structural refinement on the interactive graph to uncover more intricate user behaviors, item characteristics, and group interests, realizing better representation learning of interests. Furthermore, we also observe the intent overlapping in real-world scenarios and devise a novel self-supervised learning loss that encourages such alignment for final recommendations. Extensive experiments on three public datasets show the significant improvement of DiRec over the state-of-the-art methods.

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '23, October 21–25, 2023, Birmingham, United Kingdom
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0124-5/23/10...\$15.00
<https://doi.org/10.1145/3583780.3614855>




Figure 1: Illustration of User-centric Group Discovery (UGD) task on Douban platform. With group participation and item interactions, UGD aims to suggesting groups for users.

CCS CONCEPTS

• Information systems → Recommender systems; • Computing methodologies → Neural networks.

KEYWORDS

Recommender Systems; Graph Neural Networks; Group Recommendation; Data Mining

ACM Reference Format:

Xixi Wu, Yun Xiong, Yao Zhang, Yizhu Jiao, and Jiawei Zhang. 2023. Dual Intents Graph Modeling for User-centric Group Discovery. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3583780.3614855>

1 INTRODUCTION

With the rise of social media, online groups have become increasingly prevalent, providing users with a space to share experiences, explore new interests, and establish social connections [36, 40]. For instance, on Douban platform¹, people who are following the same TV series often form groups to discuss the ongoing plot, fostering a sense of involvement with the show. Similarly, on Steam platform², players also form groups to discuss about the gaming strategies and plan future play sessions. In terms of platforms, such users' attachment to groups can significantly boost the participation and retention rate of users. Therefore, user-centric group discovery (UGD) task, *i.e.*, recommending groups to users, can help both users'

¹<https://www.douban.com/>

²<https://store.steampowered.com/>

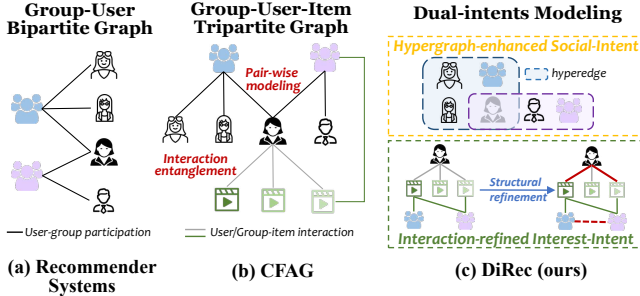


Figure 2: Modeling Method Comparison. (a) Illustration of adopting recommender systems [13, 31, 41] to UGD task. Regarding group as an item and employing recommendation methods to UGD is not ideal as missing interest-side information. (b) CFAG [40] unifies all behavioral data into a tripartite graph and thus has two drawbacks: i) pair-wise modeling of group-member relations leads to incomplete social context preservation; ii) entangled interaction data causes indistinguishable interest learning. (c) Our DiRec proposes a dual intents modeling method including hypergraph-enhanced social-intent and interaction-refined interest-intent.

online experiences and platforms’ long-term developments, which, however, is still under-explored.

Existing recommender systems [13, 14, 22, 31, 34, 41] focus more on discovering relevant items for individual users. Typically, they build user-item bipartite graph based on observed interactions and then employ graph neural networks [12, 17, 28, 37] to propagate collaborative signals for recommendation. However, these approaches may not be appropriate for user-centric group discovery task. Groups contain rich social context and specific interests that cannot be fully captured by modeling interaction histories alone. Additionally, user-group participation can involve various motivations such as social links and personal hobbies, requiring a deeper understanding of users’ *intents* beyond their interactions. Thus, user-centric group discovery task differs from traditional recommendation settings, emphasizing the need for a unique approach that can effectively address these specific requirements.

Although prior work CFAG [40] has explored the integration of user-participation and item-interactions for group discovery, it still has two notable drawbacks: i) **incomplete social context preservation**. As shown in Figure 2, CFAG splits the comprehensive *tuple-wise* group-member relations into multiple *pair-wise* connections between users and groups, leading to group-side information loss and hindering further social signals propagation; ii) **indistinguishable interests representation**. Simply unifying all behavioral data among users, items, groups into a tripartite graph for representation learning may cause different semantics and diverse preference messages to entangle with each other, with propagating deeper layers, ultimately leads to indistinguishable representations [44]. Therefore, a more nuanced approach is needed for the UGD task.

In this paper, we focus on exploring the underlying intents that drive users to participate in groups to improve recommendation. Users’ intents can generally be categorized into *social-intent* and *personal interest-intent*. The former refers to users joining a group due to social impacts from online friends that are connected via

social links. For example, users may join family groups to keep in touch with relatives or connect with friends. Meanwhile, the interest-intent motivates users to join groups with other members who have no social connections but share common interests instead, allowing them to engage and discuss their passions with like-minded individuals. For instance, an ardent Korean-Opera fan may join a Korean-Opera group to discuss the series with other enthusiasts for further enjoyment.

To better understand aforementioned intents, we propose a novel model, **DiRec**, that first models each intent independently and then fuses them together for final predictions. Our approach to intents learning is characterized by two key innovations: i) **Hypergraph-enhanced Social-Intent Modeling**. To fully capture the tuple-wise relationships between groups and corresponding members, we propose a hypergraph structure for modeling by connecting each group node to corresponding members’ nodes within a hyperedge. Based on such expressive group representation, more informative social messages can be generated and then propagated, reinforcing both users’ and groups’ social-side features. In a nutshell, hypergraph modeling enables more comprehensive relationship preservation and more effective representation learning, thereby leading to a richer understanding of social context. ii) **Interaction-refined Interest-Intent Modeling**. As to interest-intent, we mainly mine both users’ and groups’ interests from their interaction histories. However, the raw interaction data can be either noisy [7, 8, 27, 42] or sparse [10], hampering the ability to accurately estimate their preferences. Therefore, we conduct structural refinement on the interactive graph to uncover more intricate user behaviors, item characteristics, and group interests. Through this process, more implicit interactive patterns can be captured, allowing for better interests learning.

Though we separately model users’ social-intent and interest-intent to capture each intent’s distinction, in real-world scenarios, it is common for these intents to overlap. For example, a group of colleagues who work together in the same industry may primarily join a group for social intent, such as staying connected outside of work, while also having an interest in exploring the latest innovations in their field. This overlapping of social and interest intents highlights the need to consider their alignment. Therefore, we propose a novel self-supervised loss [4, 11, 43] that encourages possible intent alignment, allowing dual intents fusion for recommendations.

To summarize, our contributions are as follows:

- In this paper, we focus on the user-centric group discovery task and propose a novel model DiRec. Specifically, we uncover the dual intents behind user participation, *i.e.*, social-intent and interest-intent, for making recommendations.
- For social-intent, we introduce the hypergraph structure for tuple-wise group-member relationship modeling, surpassing the limited expressiveness of existing bipartite graph modeling. As for interest-intent, we employ novel structural refinement on the interactive graph to uncover more intricate user behaviors and group interests, allowing better representation of interests.
- To account for the possible overlaps of dual intents, we propose a novel contrastive learning objective.
- We evaluate DiRec on three public datasets and show significant improvements over existing methods.

2 NOTATION PRELIMINARIES

In this section, we first present the concept of hypergraph and then give a formal definition of user-centric group discovery (UGD) task to facilitate comprehension. Formally, we use bold capital letters (e.g., \mathbf{X}) and bold lowercase letters (e.g., \mathbf{x}) to represent matrices and vectors, respectively. We utilize non-bold letters (e.g., x) to denote scalars, and calligraphic letters (e.g., \mathcal{X}) to denote sets. Notations are summarized in Table 1.

Different from simple graphs, a hypergraph is a more general topological structure where the edge (namely hyperedge in the hypergraph) could connect two or more nodes [45]. Formally, a hypergraph can be represented as $HG = (\mathcal{V}, \mathcal{E}, \mathbf{H})$ where \mathcal{V} is the vertex set, \mathcal{E} is the hyperedge set, and $\mathbf{H} = [h_{ve}] \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{E}|}$ depicts the connectivity of the hypergraph as $h_{ve} = 1$ if the hyperedge e connects the vertex v , otherwise $h_{ve} = 0$. For the preservation of group-user affiliations, hypergraph modeling is superior to previous bipartite graph modeling where *tuple-wise* relationships are split into multiple *pair-wise* connections. By connecting each group and its members within a hyperedge, the comprehensive social context can be preserved, benefiting further representation learning stage.

Specifically, let $\mathcal{U} = \{u_1, u_2, \dots, u_M\}$, $\mathcal{I} = \{i_1, i_2, \dots, i_N\}$, and $\mathcal{G} = \{g_1, g_2, \dots, g_K\}$ be the sets users, items, and groups, respectively, where M , N , and K are the sizes of these three sets. There are three types of observed interactions among \mathcal{U} , \mathcal{I} , and \mathcal{G} , namely, user-group participation, user-item interactions, and group-item interactions. We use $\mathbf{X} = [x_{jt}] \in \mathbb{R}^{M \times K}$ to represent the user-group participation, where $x_{jt} = 1$ if user u_j is a member of group g_t , otherwise $x_{jt} = 0$. Analogously, we define user-item and group-item interactive matrices as $\mathbf{Y} \in \mathbb{R}^{M \times N}$ and $\mathbf{Z} \in \mathbb{R}^{K \times N}$, respectively. The element $y_{jk} = 1$ if user u_j has interacted with item i_k , otherwise $y_{jk} = 0$. For the **user-centric group discovery task** [40] studied in this paper, its target is to conduct group recommendation, which is defined as follows:

Definition 2.1 (User-centric Group Discovery). Given user, item, and group sets \mathcal{U} , \mathcal{I} , and \mathcal{G} , and their interactive matrices \mathbf{X} , \mathbf{Y} , and \mathbf{Z} , the user-centric group discovery (UGD) for a user $u \in \mathcal{U}$ is to predict a ranking list of groups $\{g_l, g_m, \dots, g_n\} \subset \mathcal{G}$ where user u has no existing interactions with such groups yet but will be highly likely to be interested in them in the future.

3 METHODOLOGY

In this section, we present our proposed DiRec for the user-centric group discovery task. The overall framework is shown in Figure 3. We start by introducing the embedding layers that will be trained within this framework in Section 3.1. Next, we detail our approach for modeling and learning dual intents in Sections 3.2 and 3.3, respectively. Specifically, for social-intent, we leverage hypergraph for improved relationship preservation. As for interest-intent, we perform structural refinement to better uncover both users' and groups' interests. Finally, in Sections 3.4 and 3.5, we demonstrate the fusion of dual intents for recommendation.

3.1 Embedding Layer

We maintain three embedding tables \mathbf{U} , \mathbf{I} , and \mathbf{G} to represent the trainable embeddings for users, items, and groups, respectively.

Table 1: Important Notations

Symbol	Description
$\mathcal{U}, \mathcal{I}, \mathcal{G}$	Sets of users, items, and groups
M, N, K	Numbers of users, items, and groups
$\mathbf{X}, \mathbf{Y}, \mathbf{Z}$	User-group, user-item, and group-item interactive matrices
$\mathcal{M}_t = \{u_1^t, \dots, u_{ \mathcal{M}_t }^t\}$	Member set of group g_t
$\mathcal{I}_j = \{i_1^j, \dots, i_{ \mathcal{I}_j }^j\}$	Interaction set of user/group j
$\mathcal{N}_k = \{u_1^k, \dots, u_{ \mathcal{N}_k }^k\}$	Neighboring set of item i_k in $G^{(u)}$
$\mathbf{U} = [\mathbf{U}^{(s)} \parallel \mathbf{U}^{(i)}]$	Embedding table of users
$\mathbf{G} = [\mathbf{G}^{(s)} \parallel \mathbf{G}^{(i)}]$	Embedding table of groups
\mathbf{I}	Embedding table of items
$HG^{(s)} = (\mathcal{V}^{(s)}, \mathcal{E}^{(s)}, \mathbf{H}^{(s)})$	Social hypergraph
$G^{(u)} = (\mathcal{V}^{(u)}, \mathcal{E}^{(u)}, \mathbf{A}^{(u)})$	User interactive graph
$G^{(g)} = (\mathcal{V}^{(g)}, \mathcal{E}^{(g)}, \mathbf{A}^{(g)})$	Group interactive graph
$\bar{\mathbf{U}} = [\bar{\mathbf{U}}^{(s)} \parallel \bar{\mathbf{U}}^{(i)}]$	Refined users' embedding
$\bar{\mathbf{G}} = [\bar{\mathbf{G}}^{(s)} \parallel \bar{\mathbf{G}}^{(i)}]$	Refined groups' embedding

Given that we infer a user's inclination towards groups via dual intents, we consider each user and group containing dual information, namely, social-side and interest-side representation. To extract a user u_j 's dual information from raw embedding $\mathbf{u}_j = \mathbf{U}_{j,:}$, we apply the following procedure:

$$\mathbf{u}_j^{(s)}, \mathbf{u}_j^{(i)} = \text{ExtractLayer}(\mathbf{u}_j),$$

where $\mathbf{u}_j^{(s)}$ and $\mathbf{u}_j^{(i)}$ represent user u_j 's social-side and interest-side representation, respectively. Similarly, we obtain group g_t 's social-side and interest-side embedding as $\mathbf{g}_t^{(s)}, \mathbf{g}_t^{(i)} = \text{ExtractLayer}(\mathbf{g}_t)$, where $\mathbf{g}_t = \mathbf{G}_{t,:}$ corresponds to the raw embedding of group g_t .

There are different ways to implement $\text{ExtractLayer}(\cdot)$. For instance, it can be linear transformation, i.e., $\mathbf{u}_j^{(s)} = \mathbf{W}^{(s)} \mathbf{u}_j$, $\mathbf{u}_j^{(i)} = \mathbf{W}^{(i)} \mathbf{u}_j$, where $\mathbf{W}^{(s)}$ and $\mathbf{W}^{(i)}$ are two trainable weight parameter matrices. Moreover, we can also simply define $\text{ExtractLayer}(\cdot)$ as a function that splits a vector into two subvector components, i.e., $[\mathbf{u}_j^{(s)} \parallel \mathbf{u}_j^{(i)}] = \mathbf{u}_j$ where \parallel is the concatenation operation. Both function definition will work for our method, and, by default, we opt for the latter approach due to its simplicity and effectiveness. The shape of item embedding matrix \mathbf{I} is $\mathbb{R}^{N \times d}$ while the shapes of \mathbf{U} and \mathbf{G} are $\mathbb{R}^{M \times 2d}$ and $\mathbb{R}^{K \times 2d}$ (d is the embedding dimension), respectively.

3.2 Social-Intent Learning

In this subsection, we elaborate on social-intent learning process. To be more specific, we introduce a novel hypergraph structure for modeling tuple-wise relationships between groups and corresponding members. Through hypergraph neural networks, we are able to enhance the social-side representations of both users and groups by fully preserving their social context.

3.2.1 Social Hypergraph Construction. We construct a social hypergraph $HG^{(s)} = (\mathcal{V}^{(s)}, \mathcal{E}^{(s)}, \mathbf{H}^{(s)})$ for modeling existing user-group participation. Specifically, the nodes within this hypergraph

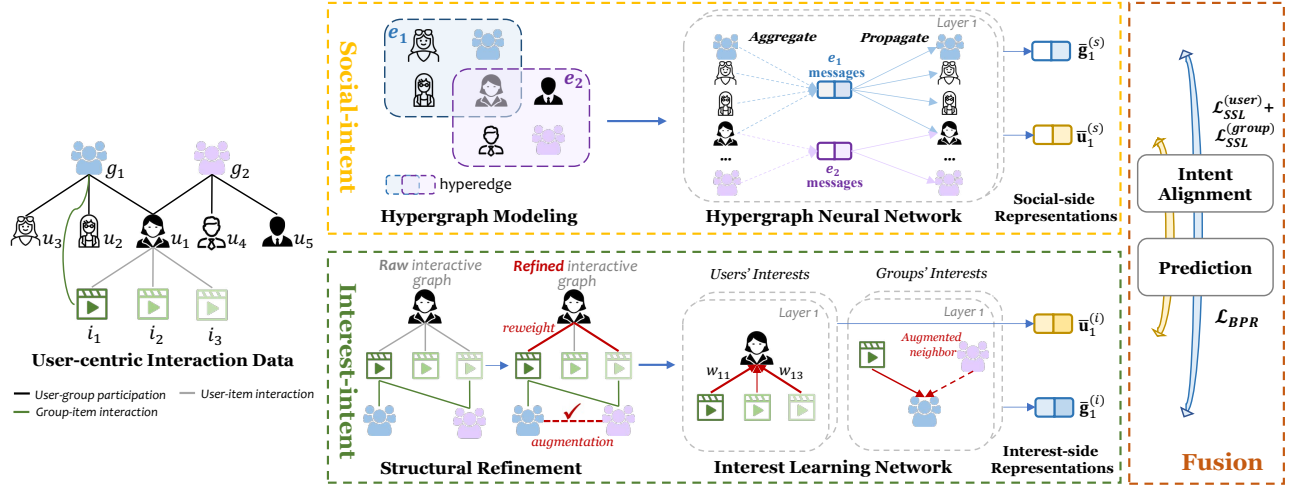


Figure 3: DiRec overview. In the left part, we present the interactive data between users, groups, and items (omit some group-item edges to avoid link interference). We model users' inclination towards groups from dual intents, including social-intent and interest-intent. For social-intent, we leverage hypergraph for relationship preservation and employ hypergraph neural networks for representation learning. As to interest-intent, we conduct structural refinement to uncover intricate user behavior patterns and item characteristics, leading to better interests learning. Finally, dual intents are fused for optimization.

contain all groups and users as $\mathcal{V}^{(s)} = \mathcal{U} \cup \mathcal{G}$, and each group node and their members' nodes are connected via a unique hyperedge as $\mathcal{E}^{(s)} = \mathcal{G}$. And $\mathbf{H}^{(s)} = [\mathbf{h}_{vt}^{(s)}] \in \mathbb{R}^{(M+K) \times K}$ is the social hypergraph adjacency matrix as $h_{vt}^{(s)} = 1$ if node $v \in \mathcal{M}_t \cup \{g_t\}$, where \mathcal{M}_t is the member set of group g_t . For instance, as shown in Figure 3, group g_1 contains three members, therefore, nodes g_1 , u_1 , u_2 , and u_3 are connected within the hyperedge e_1 .

We would like to emphasize that hypergraph modeling is superior to previous **bipartite graph modeling** where tuple-wise relationships are split into multiple pair-wise connections. By connecting each group and its member within a hyperedge, the comprehensive social context of each group can be preserved [18], allowing for group-level social signal propagating, and, finally, benefiting the representation learning. Empirical study also shows the effectiveness of hypergraph modeling than bipartite graph modeling, which will be shown in Section 4.3.2.

3.2.2 Social Context Representation Learning. Based on the constructed social hypergraph, we further conduct representation learning to obtain refined users' and groups' social-side representations.

In hypergraph neural networks [6, 9, 15, 33], hyperedges serve as mediums for processing and transferring information. More specifically, each hyperedge first aggregates the information of its connected nodes to generate messages, which are then used to update nodes' representations [36]. As previously mentioned, we connect each group and its members into a single hyperedge. During the message aggregation stage, the messages carried by each hyperedge are highly expressive, as incorporating both member compositional information and the group's distinct social representation. Therefore, both groups' and users' final social representations can be enhanced from these expressive hyperedge messages.

Formally, we feed the concatenation of users' and groups' social-side embeddings $[\mathbf{U}^{(s)} \parallel \mathbf{G}^{(s)}]$ as well as social hypergraph adjacency matrix $\mathbf{H}^{(s)}$ to hypergraph neural networks [9]. After propagating several layers, we obtain the final users' and groups' social-side representations as follows:

$$\bar{\mathbf{U}}^{(s)}, \bar{\mathbf{G}}^{(s)} = \text{HyperGNN}_{\Theta_1}([\mathbf{U}^{(s)} \parallel \mathbf{G}^{(s)}], \mathbf{H}^{(s)}),$$

where Θ_1 denotes the parameters of hypergraph neural networks. We employ the classical hypergraph neural network [9] due to its simple implementation and generally good performance. Its computation is that $\mathbf{X}^{(l+1)} = \sigma(\mathbf{D}_{\mathcal{V}}^{-\frac{1}{2}} \mathbf{H} \mathbf{D}_{\mathcal{E}}^{-1} \mathbf{H}^T \mathbf{D}_{\mathcal{V}}^{-\frac{1}{2}} \mathbf{X}^{(l)} \Theta^{(l)})$, where $\mathbf{X}^{(l)}$ denotes the features in the l -th layer, $\Theta^{(l)}$ denotes the weight matrix in the l -th layer, and $\mathbf{D}_{\mathcal{V}}$ and $\mathbf{D}_{\mathcal{E}}$ denote the degree matrix of nodes and hyperedges, respectively. Although more sophisticated hypergraph neural networks exist [6, 33], we leave their application as a future work since the focus of our paper is the overall design of dual intents.

3.3 Interest-Intent Learning

Despite social aspect, users also join groups with like-minded people for personal enjoyment. Therefore, we propose the interest intent that involves mining interests from interaction histories. Specifically, we conduct structural refinement to better uncover their interests. However, dynamically learning structures may result in large parameter sizes, slow convergence, and unstable performance [19]. These issues can hinder the scalability and extensibility of our framework. Therefore, we conduct structural refinement as a pre-computation stage, which can be easily implemented and does not impact overall efficiency.

3.3.1 Structural Refinement on User-Item Interactions. The challenges of distilling users' interest lie in noisy user-item interactions

as they may contain certain noises that do not necessarily indicate his/her real personal preference, *e.g.*, random click or surrogate shopping [8, 20, 21, 27]. Therefore, it is necessary to differentiate between various interactive weights on the user-item interactive graph in order to better identify a user's true intention.

Our intuition is that, co-occurrence between items indicates collaborative signals, making it useful for capturing users' interests [32]. For example, as shown in Figure 3, we find that in user-item interactive graph, items i_1 and i_3 are commonly consumed by some users, reflecting their similarities and user u_1 's interests. Consequently, we would assign more weights to interactions $u_1 - i_1$ and $u_1 - i_3$ to preserve such signal. Formally, we first measure the co-occurrence ratio between two items i_k and i_j via Salton Cosine Similarity (SC) [24] as $\mathbf{SC}(k, j) = \frac{|\mathcal{N}_k \cap \mathcal{N}_j|}{\sqrt{|\mathcal{N}_k \cup \mathcal{N}_j|}}$ where \mathcal{N}_k and \mathcal{N}_j refer to the neighboring set of item i_k and i_j , respectively, in the user-item interaction graph. Then, for re-assigning the interaction weight between user u and item i_j , we would compute the normalized version of above similarities among user u 's all interacted items as $w_{uj} = \frac{1}{|\mathcal{I}_u|} \sum_{i_k \in \mathcal{I}_u} \mathbf{SC}(k, j)$ where \mathcal{I}_u denotes the interaction set of user u . Based on all observed interactions, we could obtain a new weight matrix $\mathbf{W}^{(u)} = [w_{uj}] \in \mathbb{R}^{M \times N}$. Finally, we refine the user-item interactive matrix with this newly computed matrix as $\mathbf{Y}' = \mathbf{Y} + \lambda_1 \mathbf{W}^{(u)}$ where \mathbf{Y}' denotes the refined user-item interactive matrix and λ_1 represents the reweighting coefficient. Using \mathbf{Y}' , we construct refined user-item interactive bipartite graph $G^{(u)} = (\mathcal{V}^{(u)}, \mathcal{E}^{(u)}, \mathbf{A}^{(u)})$ where $\mathcal{V}^{(u)} = \mathcal{U} \cup \mathcal{I}$, $\mathcal{E}^{(u)} = \{(u_j, i_k) | u_j \in \mathcal{U}, i_k \in \mathcal{I}, \text{ and } y_{jk} = 1\}$, and $\mathbf{A}^{(u)} = \begin{bmatrix} \mathbf{0} & \mathbf{Y}' \\ \mathbf{Y}'^T & \mathbf{0} \end{bmatrix}$.

3.3.2 Structural Refinement on Group-Item Interactions. Different from noisy user-item interactions, group-item interactive records are extremely sparse [2, 25], making it hard to estimate groups' interest-side representations with limited neighboring nodes on the interactive graph [10, 40]. To overcome this limitation, we consider performing structural augmentation to enrich each group's neighborhood, thereby improving their interest learning with a larger number of neighbors.

In user recommendation scenarios, if two users have purchased the same item, they are likely to share similar interests [14]. This idea of collaborative filtering can be extend to group-level: if two groups have once interacted with the same item, they may share similar preferences and can reinforce each other's representations toward similar interests. Therefore, on the group-item interactive graph, besides observed group-item links, we connect two groups if they share common interactions. In this way, we obtain a refined group-item graph $G^{(g)} = (\mathcal{V}^{(g)}, \mathcal{E}^{(g)})$ where $\mathcal{V}^{(g)} = \mathcal{G} \cup \mathcal{I}$, $\mathcal{E}^{(g)} = \{(g_t, i_k) | g_t \in \mathcal{G}, i_k \in \mathcal{I}, \text{ and } z_{tk} = 1\} \cup \{(g_p, g_q) | g_p, g_q \in \mathcal{G}, |\mathcal{I}_p \cap \mathcal{I}_q| \geq 1\}$ (\mathcal{I}_p denotes the interaction set of group g_p). Based on augmented edges, we can obtain the adjacency matrix $\mathbf{A}^{(g)}$, which will be further leveraged for interests learning.

3.3.3 Interests Learning. Using both the adjusted user-item and group-item interactive graphs, we perform representation learning to obtain refined users' and groups' interest-side embeddings.

Specifically, for users' interests distilling, recall that we have re-assigned the interaction weight and thus obtain a new user

interactive graph $G^{(u)}$, we would feed the concatenation of users' interest-side embeddings $\mathbf{U}^{(i)}$ and items' embeddings \mathbf{I} , together with the adjacency matrix $\mathbf{A}^{(u)}$ to graph neural networks [13, 31] to refine users' final interests as follows:

$$\bar{\mathbf{U}}^{(i)} = \text{GNN}_{\Theta_2}([\mathbf{U}^{(i)} \parallel \mathbf{I}], \mathbf{A}^{(u)}),$$

where Θ_2 denotes the trainable parameters of graph neural networks and $\bar{\mathbf{U}}^{(i)}$ represents users' final interest-side representations. In this way, users' interests can be effectively captured from most representative interactions, thus pinpointing his/her actual interests. Additionally, we calculate the interest-side representations of groups in a similar way:

$$\bar{\mathbf{G}}^{(i)} = \text{GNN}_{\Theta_3}([\mathbf{G}^{(i)} \parallel \mathbf{I}], \mathbf{A}^{(g)}),$$

where Θ_3 denotes the trainable parameters and $\mathbf{A}^{(g)}$ denotes the adjacent matrix of refined group's interactive graph. Note that instead of unifying users, groups, items into a single graph, we separately deal with users' and groups' interaction data to better understand their respective interests. We employ **LightGCN** [13] for the implementation of graph neural networks due to its simplicity and effectiveness.

3.4 Combination of Dual Intents

Though we separately model social-intent and interest-intent for capturing each intent's distinction, we note that, in real-world scenarios, dual intents may have certain overlaps. To address such overlap, we propose a self-supervised learning objective to guide intent alignment.

More specifically, the crux of our intent alignment approach is to ensure that the user u_j 's social-side representation $\bar{\mathbf{u}}_j^{(s)}$ is more similar to his/her interest-side representation $\bar{\mathbf{u}}_j^{(i)}$ than other users' interest representations, *e.g.*, $\bar{\mathbf{u}}_k^{(i)}$ as $u_k \in \mathcal{U}$. Therefore, we employ the InfoNCE loss [11] for implementation as follows:

$$\mathcal{L}_{SSL}^{(user)_1} = - \sum_{u_j \in \mathcal{U}} \log \frac{\exp(\text{sim}(\bar{\mathbf{u}}_j^{(s)}, \bar{\mathbf{u}}_j^{(i)}))}{\sum_{u_k \in \mathcal{U}} \exp(\text{sim}(\bar{\mathbf{u}}_j^{(s)}, \bar{\mathbf{u}}_k^{(i)}))},$$

where $\text{sim}(\cdot)$ denotes a certain similarity measurement function, *i.e.*, inner product. And this objective holds for interest-side representations as well:

$$\mathcal{L}_{SSL}^{(user)_2} = - \sum_{u_j \in \mathcal{U}} \log \frac{\exp(\text{sim}(\bar{\mathbf{u}}_j^{(i)}, \bar{\mathbf{u}}_j^{(s)}))}{\sum_{u_k \in \mathcal{U}} \exp(\text{sim}(\bar{\mathbf{u}}_j^{(i)}, \bar{\mathbf{u}}_k^{(s)}))}.$$

Finally, we compute user-side loss as $\mathcal{L}_{SSL}^{(user)} = \mathcal{L}_{SSL}^{(user)_1} + \mathcal{L}_{SSL}^{(user)_2}$. Additionally, we infer a group-side loss $\mathcal{L}_{SSL}^{(group)}$ in a similar manner. By imposing these self-supervised learning objectives, we enable the model to better capture overlaps in dual intents and to generate more informative representations for both users and groups.

3.5 Prediction and Optimization

In this subsection, we show the detailed prediction and optimization process. After obtaining social-side and interest-side representations, we concatenate them to obtain final users' and groups' representations as follows:

$$\bar{\mathbf{U}} = [\bar{\mathbf{U}}^{(s)} \parallel \bar{\mathbf{U}}^{(i)}], \bar{\mathbf{G}} = [\bar{\mathbf{G}}^{(s)} \parallel \bar{\mathbf{G}}^{(i)}],$$

where \parallel denotes concatenation operation. Then for a user u_j and group g_t , the ranking score of this user-group pair is calculated by the inner product:

$$\hat{s}_{jt} = \bar{\mathbf{u}}_j \cdot \bar{\mathbf{g}}_t,$$

where $\bar{\mathbf{u}}_j = \bar{\mathbf{U}}_j$; and $\bar{\mathbf{g}}_t = \bar{\mathbf{G}}_t$, refer to u_j and g_t final representations, respectively. Then, we employ the Bayesian Personalized Ranking (BPR) loss [23] for optimization as follows:

$$\mathcal{L}_{BPR} = \sum_{(u_j, g_t, g_{t'}) \in \mathcal{D}} -\log \sigma(\hat{s}_{jt} - \hat{s}_{jt'}) + \lambda_2 \|\Theta\|_2^2,$$

where $\mathcal{D} = \{(u_j, g_t, g_{t'}) | x_{jt}=1 \text{ and } x_{jt'} = 0\}$ is the training data with observed interactions and random sampled negative pairs. Θ denotes all trainable parameters within this framework, which is regularized by λ_2 . Finally, we unify the ranking loss as well as self-supervised loss into a joint training framework as $\mathcal{L} = \mathcal{L}_{BPR} + \lambda_3 \mathcal{L}_{SSL}^{(user)} + \lambda_4 \mathcal{L}_{SSL}^{(group)}$ where λ_3 and λ_4 are two hyper-parameters.

4 EXPERIMENTS

In this section, we present our experimental setup and empirical results. Our experiments are designed to answer the following research questions (RQs):

- **RQ1 (Overall Performance):** How does DiRec perform compared with both representative recommendation methods and the start-of-the-art UGD method?
- **RQ2 (Ablation Study):** How do the key components of DiRec affect the model performance, *i.e.*, social-intent, interest-intent, hypergraph modeling, structural refinement, and SSL loss?
- **RQ3 (Parameters Study):** How do different settings of hyper-parameters affect the model performance?

4.1 Experimental Settings

4.1.1 Datasets. Following prior work CFAG [40], we conduct experiments on three real-world datasets: Mafengwo, Weeplaces, and Steam. The statistics of these datasets are listed in Table 2. The Mafengwo and Weeplaces datasets comprise the user's travel history with a location-based social network, whereas Steam is a newly-released dataset, collected from the Steam online platform by CFAG. We found CFAG did not split the validation set for model selection, therefore, we **re-split** the dataset into train set, validation set, and test set with the ratios of 70%, 10%, and 20%, respectively.

4.1.2 Baselines. To show the effectiveness of our proposed DiRec, we compare it with the following representative baselines:

Group Recommendation methods: Group Recommendation (GR) task is defined as suggesting suitable items for groups [36]. Therefore, GR methods focus on modeling groups' preferences,

Table 2: Statistics of datasets

Dataset	Mafengwo	Weeplaces	Steam
# Users	1,269	1,501	11,099
# Groups	972	4,651	1,085
# Items	999	6,406	2,351
# User-Group Participation	5,574	12,258	57,654
# User-Item Interactions	8,676	43,942	444,776
# Group-Item Interactions	2,540	6,033	22,318
Avg. # Groups/user	4.39	8.17	5.19
Avg. # Items/user	6.84	29.28	40.07
Avg. # Items/group	2.61	1.29	20.57

overlooking the inclination between users and groups. To empirically show the inferior performance of applying GR on UGD task, we consider the following GR methods. Note that to apply these methods to UGD task, we adapt them by replacing their initial group-item BPR loss with user-group BPR loss:

- **AGREE [1]** This is a classical group recommendation method that leverages attention mechanism to estimate group representations by adaptively aggregating members' representations.
- **ConsRec [36]** This is the state-of-the-art group recommendation method that models group's preferences from member-level, item-level, and group-level views.

Recommender methods: We also consider following item recommendation models to illustrate their infeasibility on the UGD task. As CFAG [40] does, we apply them on UGD task by treating each group as an item and thus only utilizing user-group interactions:

- **MF-BPR [23]** This is a classical recommendation model that employs matrix factorization technique and is optimized by the pair-wise BPR loss.
- **NGCF [31]** This method is a variant of standard GCN [17] that leverages high-order connectivity in a user-item bipartite graph.
- **LightGCN [13]** This is a method based on NGCF with optimization in training efficiency and generation ability by removing feature transformation and non-linear activation.
- **SGL [34]** This work performs contrastive learning on LightGCN model to augment node representations.
- **SimGCL [41]** This is a strong GNN-based recommender model that incorporates an augmentation-free contrastive loss.
- **DCCF [22]** This is the state-of-the-art GNN-based recommender model with disentangled representations.

User-centric Group Discovery (UGD) methods:

- **CFAG [40]** This is the only existing UGD method. It constructs a group-user-item tripartite graph and designs a specific tripartite graph convolution mechanism for representation learning.

4.1.3 Evaluation Metrics. We evaluate UGD task by ranking the test groups with all non-interacted groups of users. We adopt Recall@{5, 10, 20} and NDCG@{5, 10, 20} as evaluation metrics. Higher Recall and NDCG indicate better performance. We evaluate the performance of all methods over the same metrics and test data.

4.1.4 Implementation Details. We implement DiRec in PyTorch. For all baselines, we utilize their official codes released in Github. We conduct all the experiments on GPU machines of Nvidia Tesla V100

Table 3: UGD Performance comparison on three datasets with Recall (R) reported.

Method Type	Dataset Metric	Mafengwo			Weeplaces			Steam		
		R@5	R@10	R@20	R@5	R@10	R@20	R@5	R@10	R@20
Group Recommendation	AGREE _(SIGIR'18)	0.1216	0.1716	0.2368	0.1680	0.2270	0.2911	0.1168	0.1768	0.2577
	ConsRec _(WWW'23)	0.2204	0.3141	0.4063	0.2514	0.3472	0.4391	<u>0.1690</u>	0.2440	0.3330
Recommender Systems	MF-BPR	0.1832	0.2407	0.2973	0.1734	0.2278	0.2822	0.1207	0.1848	0.2671
	NGCF _(SIGIR'19)	0.1999	0.2650	0.3284	0.1787	0.2392	0.2961	0.1330	0.2010	0.2903
	LightGCN _(SIGIR'20)	0.2293	0.2930	0.3596	0.1791	0.2431	0.3069	0.1543	0.2343	0.3273
	SGL _(SIGIR'21)	0.2259	0.2956	0.3555	0.1810	0.2443	0.3046	0.1535	0.2336	0.3276
	SimGCL _(SIGIR'22)	<u>0.2309</u>	0.2925	0.3585	0.1808	0.2477	0.3079	0.1544	0.2332	0.3271
	DCCF _(SIGIR'23)	0.2049	0.2649	0.3216	0.1858	0.2507	0.3171	0.1614	0.2417	0.3423
User-centric Group Discovery	CFAG _(WSDM'23)	0.2274	<u>0.3242</u>	<u>0.4194</u>	<u>0.2855</u>	<u>0.3824</u>	<u>0.4893</u>	0.1597	<u>0.2485</u>	<u>0.3502</u>
	DiRec _{ours}	0.2653	0.3585	0.4549	0.3177	0.4119	0.4987	0.1702	0.2685	0.3708

Table 4: UGD Performance comparison on three datasets with NDCG (N) reported.

Method Type	Dataset Metric	Mafengwo			Weeplaces			Steam		
		N@5	N@10	N@20	N@5	N@10	N@20	N@5	N@10	N@20
Group Recommendation	AGREE _(SIGIR'18)	0.0887	0.1051	0.1218	0.1132	0.1330	0.1501	0.0742	0.0937	0.1143
	ConsRec _(WWW'23)	0.1519	0.1832	0.2069	0.1714	0.2037	0.2283	0.1034	<u>0.1335</u>	<u>0.1604</u>
Recommender Systems	MF-BPR	0.1293	0.1484	0.1631	0.1170	0.1358	0.1507	0.0771	0.0979	0.1190
	NGCF _(SIGIR'19)	0.1440	0.1655	0.1820	0.1213	0.1417	0.1572	0.0862	0.1082	0.1311
	LightGCN _(SIGIR'20)	0.1713	0.1922	0.2093	0.1225	0.1443	0.1619	0.1013	0.1274	0.1513
	SGL _(SIGIR'21)	0.1718	<u>0.1945</u>	0.2099	0.1235	0.1448	0.1614	0.1018	0.1278	0.1520
	SimGCL _(SIGIR'22)	<u>0.1725</u>	0.1926	0.2095	0.1233	0.1460	0.1626	0.1015	0.1273	0.1514
	DCCF _(SIGIR'23)	0.1493	0.1691	0.1837	0.1256	0.1476	0.1659	<u>0.1053</u>	0.1314	0.1573
User-centric Group Discovery	CFAG _(WSDM'23)	0.1552	0.1867	<u>0.2111</u>	<u>0.1938</u>	<u>0.2264</u>	<u>0.2551</u>	0.1035	0.1324	0.1584
	DiRec _{ours}	0.1908	0.2208	0.2455	0.2246	0.2565	0.2797	0.1086	0.1405	0.1669

with 32GB memory. We empirically set the weight of regularization loss λ_2 as 0.0001. To obtain the values of λ_3 and λ_4 , we conduct the grid search over $\{0, 0.0001, 0.001, 0.01, 0.1, 1\}$. All reported results are the average of **5 runs**. Our codes, data, and baseline implementations are available at <https://github.com/WxxShirley/CIKM2023DiRec>.

4.2 Overall Performance (RQ1)

We present the overall performance comparison in Tables 3 and 4. From these results, we note the following key observations:

- Our proposed DiRec achieves the best performance on all three datasets over all evaluation metrics, demonstrating its superiority. Compared with the start-of-the-art UGD model CFAG [40], our model still realizes a significant improvement. Unlike CFAG, which unifies all types of nodes and interactive edges into a single graph, we preserve the distinction of dual intents (*i.e.*, social- and interest- intent), resulting in better performance.
- Group Recommendation methods are not suitable for UGD task as both the classical method AGREE [1] and the state-of-the-art method ConsRec [36] perform poorly on all datasets. The reason for this is that GR aims to estimate group interests accurately, failing to fully capture the inclination between users and groups.

- Existing recommender models also fail to address UGD tasks. Modeling observed user-group participation as a bipartite graph and conducting representation learning as NGCF [31], LightGCN [13], SGL [34], SimGCL [41], and DCCF [22] only captures social-intent while interest-intent is overlooked. Their inferior performance highlights the necessity of dual intents modeling.

4.3 Ablation Study (RQ2)

In this subsection, we explore the effectiveness of key components within DiRec to demonstrate the rationale behind our proposal. Specifically, we evaluate the following components: dual intents modeling, the hypergraph modeling, the structural refinement, and the self-supervised learning loss. Due to space limitations, we only present the evaluation metrics of NDCG@{10, 20} in Table 5, while the Recall shows the same trend.

4.3.1 Effectiveness of Dual Intents. To showcase the effectiveness of dual intents modeling, we conduct an ablation study where we predict the ranking score **with only one intent**, *i.e.*, either social-intent or interest-intent. The corresponding results are denoted as “Social-” and “Interest-” in Table 5, respectively. It is obvious that utilizing only one intent leads to a significant degradation in

Table 5: Ablation Study on three datasets with NDCG (N) reported. “Social-” and “Interest-” refer to the variant that only utilizes social-intent and interest-intent, respectively. “w/o. HG” denotes the variant that replaces hypergraph modeling with bipartite graph modeling. And “w/o. UI Re.” and “w/o. GI Aug.” refer to variants that eliminate reweighting and augmentation, respectively.

Dataset Metric	Mafengwo		Weeplaces		Steam	
	N@10	N@20	N@10	N@20	N@10	N@20
Full	0.2208	0.2455	0.2565	0.2797	0.1405	0.1669
Social-Interest-	0.1954	0.2157	0.1593	0.1784	0.1373	0.1628
	0.1561	0.1900	0.2426	0.2671	0.1097	0.1391
w/o. HG	0.2110	0.2325	0.2490	0.2754	0.1253	0.1546
w/o. UI Re.	0.2208	0.2453	0.2558	0.2788	0.1398	0.1664
w/o. GI Aug.	0.2117	0.2393	0.2200	0.2417	0.1371	0.1639
w/o. SSL	0.2194	0.2430	0.2525	0.2752	0.1395	0.1651

performance, highlighting the necessity of dual intents modeling. Moreover, the importance of each intent varies across different datasets. For example, on Weeplaces, predicting only from the interest intent leads to better results, while on Steam and Mafengwo, social-intent plays a more crucial role.

4.3.2 Effectiveness of Hypergraph Modeling. Recall that we construct a hypergraph to represent the relations between groups and user members, to evaluate its effectiveness, we replace the hypergraph with a user-group bipartite graph and then adopt LightGCN [13] for representation learning. The results are denoted as “w/o. HG” in Table 5. By comparing the results with the full DiRec model, we can see that hypergraph modeling leads to a significant performance improvement, while the bipartite graph approach is less effective in preserving the user-group participation.

4.3.3 Effectiveness of Structural Refinement. To handle the distinct properties of user-item and group-item interactions, we perform structural refinement on the raw interactive graph. Specifically, we adopt a re-weighting technique on noisy user-item interactions to extract the most representative interacted items, while using an augmentation method on sparse group-item interactions to learn better group characteristics. To conduct the ablation study, we exclude user-item re-weighting and group-item augmentation techniques individually, and the results are presented as “w/o. UI Re.” and “w/o. GI Aug.” in Table 5, respectively.

UI Reweight By comparing the results of “w/o. UI Re.” and “Full” in Table 5, we can observe that the introduction of UI reweighting leads to performance improvement. However, the improvement is relatively marginal for Mafengwo because the user-item interactions are already sparse and can reflect the true user preferences. On the other hand, for Weeplaces and Steam, where user-level interactions are abundant and may contain noise, performing reweighting results in a more significant improvement.

To further investigate the effectiveness of introducing reweighting on user-item interactions, we perform structural perturbation on raw user-item interactions and then compare the performance of DiRec with its variant that removes the reweighting step. The




Figure 4: Performance comparison between DiRec and the variant that eliminates the reweighting part under varying levels of user-item interaction perturbations.

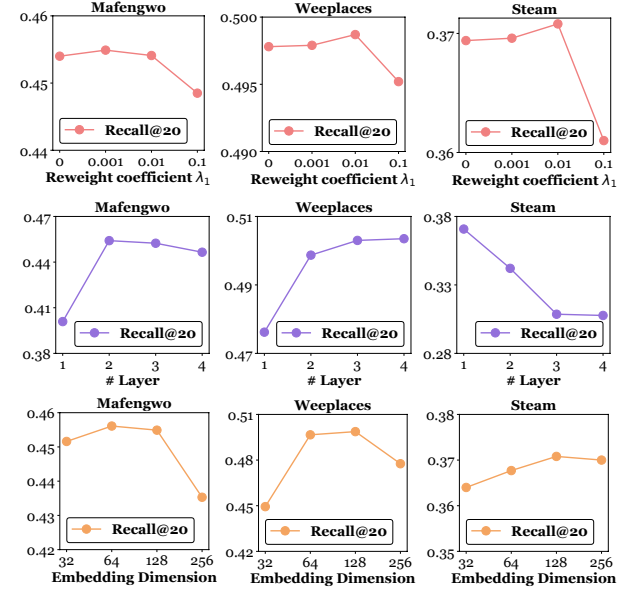


Figure 5: Parameters Study

results, as shown in Figure 4, reveal that when introducing higher levels of noise, the improvement brought by reweighting becomes more evident, demonstrating its effectiveness.

GI Augmentation By comparing the results of “w/o. GI Aug.” and “Full” in Table 5, we can observe that performing augmentation on group-item interactive graph significantly enhance the overall performance. Notably, the improvement is most significant on Weeplaces, with an increase of 16.6% on NDCG@10. This is because on Weeplaces, the group-item interactions are extremely sparse, as the average number of items that a group has interacted with is only 1.29. With so few neighbors, it is challenging to learn effective groups’ interest-side representations. However, with the augmentation method, additional neighbors can help improve the target group’s representations, leading to better performance.

4.3.4 Effectiveness of Self-supervised Learning. We also propose a self-supervised learning component that guides intent alignment for better representation. To evaluate its effectiveness, we exclude this loss and only optimize with BPR loss. This variant is denoted as “w/o. SSL”. As shown in Table 5, removing the self-supervised loss leads to a degradation in performance, indicating the importance of the self-supervised learning design.

4.4 Parameters Study (RQ3)

In this subsection, we investigate the importance of key parameters within DiRec framework, *i.e.*, the coefficient of user-item reweighting matrix λ_1 , number of (hyper)graph convolutional layers L , and embedding dimension d . The results are shown in Figure 5.

4.4.1 The Coefficient of User-item Reweighting Matrix λ_1 . We reassign the weights of user-item interaction matrix as $\mathbf{Y}' = \mathbf{Y} + \lambda_1 \mathbf{W}^{(u)}$ where λ_1 is the coefficient that controls the importance of newly computed interaction weights. To show its influence, we conduct a parameter study that tunes λ_1 across $\{0, 0.001, 0.01, 0.1\}$ and then compare performance difference. Figure 5 illustrates that with suitable λ_1 values, all three datasets experience performance improvement while overly large values of λ_1 may ruin the original interactions and lead to performance degradation. Therefore, for Mafengwo, we set λ_1 to 0.001, and 0.01 for Weeplaces and Steam.

4.4.2 Number of Convolutional Layers L . We leverage both hypergraph neural networks [9] and graph neural networks [13] for representation learning. Therefore, we perform the parameter study to show the influence of different numbers of convolutional layers. As Figure 5 indicates, we choose $L = 2$ for Mafengwo and Weeplaces, while $L = 1$ for Steam.

4.4.3 Embedding Dimension d . We also study the influence of embedding dimensions within DiRec framework. From Figure 5, we observe that the best performance is almost always achieved when the dimension is set to 128 across all datasets. Therefore, we choose 128 as the default setting.

5 RELATED WORKS

5.1 UGD Related Works

The task of user-centric group discovery was first proposed in CFAG [40]. As a pioneering work, CFAG unifies the diverse interactions between users, groups, and items into a tripartite graph and devises a tripartite graph convolutional network for representation learning and prediction. Although CFAG achieves ideal performance, we emphasize that its pair-wise modeling of group-member relations leads to incomplete social context preservation, and entangled interaction data causes indistinguishable interest learning.

Although CFAG is a prominent work in the field of UGD, no previous work has specifically targeted this task. However, there exist some closely related tasks such as (1) community detection, (2) group recommendation, and (3) recommender systems. Hence, we provide a brief introduction to these tasks and highlight the differences between them and UGD task.

5.1.1 Community Detection. Community detection task is defined as partitioning graph nodes into multiple groups, where internal nodes are more similar or more closely-related than the external [16, 26, 35, 38, 39]. Therefore, with observed social links between users, community detection aims to identify new user groups. On the contrary, UGD task differs in the following ways: i) does not necessitate the user-user relations as input, ii) all group information is already known, and iii) the objective is to predict users' ranking score towards a specific group.

5.1.2 Group Recommendation. Group recommendation task aims to suggesting suitable items for groups to boost group-level activities [1, 25, 36]. Existing group recommendation methods focus on aggregating diverse members' interests to estimate the group-level overall interest. Therefore, they mainly model both users and groups from interest-side, leaving social relations totally overlooked.

Besides recommending items to groups, the term *group recommendation* in literature also refers to recommending groups to their potential members. Traditional methods usually apply various algorithms to recover user-group membership matrices, using available side information such as semantic information from group descriptions in [3], visual information from photos in [29], and user behaviors in different time periods [30]. However, these methods require side information, which can lead to degraded performance when recommendation based solely on interaction information [40]. Thus, these methods are not suitable for UGD settings.

5.1.3 Recommender Systems. General recommendation models aim to recommend items for users based on observed user-item interactions. With graph neural networks (GNN)'s powerful capability of learning informative representations in graph data, GNN has been widely leveraged for recommendation scenarios [13, 31, 34, 41]. Most previous works on Graph-based RS only focus on user-item bipartite graph [13]. It is difficult to directly apply these works to UGD tasks with three relations to be managed: user-item, group-item and user-group relations.

5.2 Intent-aware Recommendation

In item recommendation models, the term "intent" typically refers to user's interests either for item contents or behavior purposes [5]. To study such latent and implicit intentions behind users' behaviors, recent methods model intent as a set of fixed number of learnable variables and implement it in different ways. For instance, some methods derive intent via performing clustering on users' representations [5] while others allow intent representations to be entirely learnable [22]. However, in our UGD task, we define "intent" as a user's inclination toward a group, and we differentiate intent into two aspects: social-intent and interest-intent.

6 CONCLUSION

In this paper, we study the user-centric group discovery task and propose a novel model, DiRec. Specifically, we reveal the dual intents behind users to groups, *i.e.*, social-intent and interest-intent. For social-intent, we employ hypergraph for relation preservation. To capture both users' and groups' item-level interests, we perform structural refinement on raw interactive graph. In future, we plan to study more fine-grained intents within social- and interest-intent to better understand users' behaviors. Besides, we plan to explore more complex intent alignment phenomenon and devise a more efficient self-supervised learning loss.

ACKNOWLEDGEMENTS

This work is partially supported by the National Natural Science Foundation of China Projects No. U1936213, No.62206059, China Postdoctoral Science Foundation 2022M710747, and NSF through grants IIS-1763365 and IIS-2106972.

REFERENCES

- [1] Da Cao, Xiangnan He, Lianhai Miao, Yahui An, Chao Yang, and Richang Hong. 2018. Attentive Group Recommendation. *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval* (2018).
- [2] Tong Chen, Hongzhi Yin, Jing Long, Quoc Viet Hung Nguyen, Yang Wang, and Meng Wang. 2022. Thinking inside The Box: Learning Hypercube Representations for Group Recommendation. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2022).
- [3] Wen-Yen Chen, Dong Zhang, and Edward Y. Chang. 2008. Combinational collaborative filtering for personalized community recommendation. In *Knowledge Discovery and Data Mining*.
- [4] Xinlei Chen and Kaiming He. 2020. Exploring Simple Siamese Representation Learning. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 15745–15753.
- [5] Yong-Guang Chen, Zhiwei Liu, Jia Li, Julian McAuley, and Caiming Xiong. 2022. Intent Contrastive Learning for Sequential Recommendation. *Proceedings of the ACM Web Conference 2022* (2022).
- [6] Eli Chien, Chao Pan, Jianhao Peng, and Olgica Milenkovic. 2021. You are AllSet: A Multiset Function Framework for Hypergraph Neural Networks. *ArXiv abs/2106.13264* (2021).
- [7] Wenqi Fan, Xiaorui Liu, Wei Jin, Xiangyu Zhao, Jiliang Tang, and Qing Li. 2022. Graph Trend Filtering Networks for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [8] Ziwei Fan, Ke Xu, Zhang Dong, Hao Peng, Jiawei Zhang, and Philip S. Yu. 2023. Graph Collaborative Signals Denoising and Augmentation for Recommendation. *ArXiv abs/2304.03344* (2023).
- [9] Yifan Feng, Haoxuan You, Zizhao Zhang, R. Ji, and Yue Gao. 2018. Hypergraph Neural Networks. In *AAAI*.
- [10] Wei Guo, Rong Su, Renhao Tan, Huifeng Guo, Yingxue Zhang, Zhirong Liu, Ruiming Tang, and Xiuqiang He. 2021. Dual Graph enhanced Embedding Neural Network for CTR Prediction. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*.
- [11] Michael U Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *International Conference on Artificial Intelligence and Statistics*.
- [12] William L. Hamilton, Zitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *NIPS*.
- [13] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [14] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web*.
- [15] Jing Huang and Jie Yang. 2021. UniGNN: a Unified Framework for Graph and Hypergraph Neural Networks. In *IJCAI*.
- [16] Yuting Jia, Qinqin Zhang, Weinan Zhang, and Xinbing Wang. 2019. Community-GAN: Community Detection with Generative Adversarial Nets. *The World Wide Web Conference* (2019).
- [17] Thomas Kipf and Max Welling. 2016. Semi-Supervised Classification with Graph Convolutional Networks. *ArXiv abs/1609.02907* (2016).
- [18] Geon Lee, Minyoung Choe, and Kijung Shin. 2021. How Do Hyperedges Overlap in Real-World Hypergraphs? - Patterns, Measures, and Generators. *Proceedings of the Web Conference 2021* (2021).
- [19] Yixin Liu, Yu Zheng, Daokun Zhang, Hongxu Chen, Hao Peng, and Shirui Pan. 2022. Towards Unsupervised Deep Graph Structure Learning. *Proceedings of the ACM Web Conference 2022* (2022).
- [20] Bhaskar Mehta, Thomas Hofmann, and Wolfgang Nejdl. 2007. Robust collaborative filtering. In *ACM Conference on Recommender Systems*.
- [21] Michael P. O'Mahony, Neil J. Hurley, and Guénolé C. M. Silvestre. 2006. Detecting noise in recommender system databases. In *International Conference on Intelligent User Interfaces*.
- [22] Xubin Ren, Lianghao Xia, Jiashu Zhao, Dawei Yin, and Chao Huang. 2023. Disentangled Contrastive Collaborative Filtering. *ArXiv abs/2305.02759* (2023).
- [23] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. *ArXiv abs/1205.2618* (2009).
- [24] Gerard Salton. 1989. Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer.
- [25] Aravind Sankar, Yanhong Wu, Yuhang Wu, Wei Zhang, Hao Yang, and H. Sundaram. 2020. GroupIM: A Mutual Information Maximization Framework for Neural Group Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [26] Fan-Yun Sun, Meng Qu, Jordan Hoffmann, Chin-Wei Huang, and Jian Tang. 2019. vGraph: A Generative Model for Joint Community Detection and Node Representation Learning. *ArXiv abs/1906.07159* (2019).
- [27] Changxin Tian, Yuexiang Xie, Yaliang Li, Nan Yang, and Wayne Xin Zhao. 2022. Learning to Denoise Unreliable Interactions for Graph Collaborative Filtering. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [28] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph Attention Networks. *ArXiv abs/1710.10903* (2017).
- [29] Jingdong Wang, Zhe Zhao, Jiazhen Zhou, Hao Wang, Bin Cui, and Guo-Jun Qi. 2012. Recommending Flickr groups with social topic model. *Information Retrieval* 15 (2012), 278–295.
- [30] X. Wang, Roger Donaldson, Christopher Nell, Peter Gorniak, Martin Ester, and Jiajun Bu. 2016. Recommending Groups to Users Using User-Group Engagement and Time-Dependent Matrix Factorization. In *AAAI Conference on Artificial Intelligence*.
- [31] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [32] Yu Wang, Yuying Zhao, Yi Zhang, and Tyler Derr. 2022. Collaboration-Aware Graph Convolutional Network for Recommender Systems. *Proceedings of the ACM Web Conference 2023* (2022).
- [33] Tianxin Wei, Yuning You, Tianlong Chen, Yang Shen, Jingrui He, and Zhangyang Wang. 2022. Augmentations in Hypergraph Contrastive Learning: Fabricated and Generative. *Advances in neural information processing systems* 35 (2022), 1909–1922.
- [34] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2020. Self-supervised Graph Learning for Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [35] Xixi Wu, Yun Xiong, Yao Zhang, Yizhu Jiao, Caihua Shan, Yiheng Sun, Yangyong Zhu, and Philip S. Yu. 2022. CLARE: A Semi-supervised Community Detection Algorithm. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. ACM.
- [36] Xixi Wu, Yun Xiong, Yao Zhang, Yizhu Jiao, Jiawei Zhang, Yangyong Zhu, and Philip S. Yu. 2023. ConsRec: Learning Consensus Behind Interactions for Group Recommendation. In *Proceedings of the ACM Web Conference 2023*.
- [37] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How Powerful are Graph Neural Networks? *ArXiv abs/1810.00826* (2018).
- [38] Jaewon Yang and Jure Leskovec. 2013. Overlapping community detection at scale: a nonnegative matrix factorization approach. *Proceedings of the sixth ACM international conference on Web search and data mining* (2013).
- [39] Jaewon Yang, Julian McAuley, and Jure Leskovec. 2013. Community Detection in Networks with Node Attributes. *2013 IEEE 13th International Conference on Data Mining* (2013), 1151–1156.
- [40] Mingdai Yang, Zhiwei Liu, Liangwei Yang, Xiao lin Liu, Chen Wang, Hao Peng, and Philip S. Yu. 2022. Ranking-based Group Identification via Factorized Attention on Social Tripartite Graph. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*.
- [41] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Li zhen Cui, and Quoc Viet Hung Nguyen. 2021. Are Graph Augmentations Necessary?: Simple Graph Contrastive Learning for Recommendation. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2021).
- [42] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J. Kim. 2019. Graph Transformer Networks. In *Neural Information Processing Systems*.
- [43] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. 2021. Barlow Twins: Self-Supervised Learning via Redundancy Reduction. In *International Conference on Machine Learning*.
- [44] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and N. Chawla. 2019. Heterogeneous Graph Neural Network. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (2019).
- [45] Dengyong Zhou, Jiayuan Huang, and Bernhard Scholkopf. 2006. Learning with Hypergraphs: Clustering, Classification, and Embedding. In *NIPS*.