

Thinking inside The Box: Learning Hypercube Representations for Group Recommendation

Author

Chen, Tong, Yin, Hongzhi, Long, Jing, Nguyen, Quoc Viet Hung, Wang, Jang, Wang, Meng

Published

2022

Conference Title

SIGIR '22: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval

Version

Accepted Manuscript (AM)

DOI

[10.1145/3477495.3532066](https://doi.org/10.1145/3477495.3532066)

Rights statement

© ACM, 2022. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in SIGIR '22: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 978-1-4503-8732-3, <http://doi.org/10.1145/3477495.3532066>

Downloaded from

<http://hdl.handle.net/10072/417666>

Griffith Research Online

<https://research-repository.griffith.edu.au>

Thinking inside The Box: Learning Hypercube Representations for Group Recommendation

Tong Chen

The University of Queensland
tong.chen@uq.edu.au

Hongzhi Yin*

The University of Queensland
h.yin1@uq.edu.au

Jing Long

The University of Queensland
jing.long@uq.edu.au

Quoc Viet Hung Nguyen

Griffith University
quocviethung.nguyen@griffith.edu.au

Yang Wang

Hefei University of Technology
yangwang@hfut.edu.cn

Meng Wang

Hefei University of Technology
eric.mengwang@gmail.com

ABSTRACT

As a step beyond traditional personalized recommendation, group recommendation is the task of suggesting items that can satisfy a group of users. In group recommendation, the core is to design preference aggregation functions to obtain a quality summary of all group members' preferences. Such user and group preferences are commonly represented as points in the vector space (i.e., embeddings), where multiple user embeddings are compressed into one to facilitate ranking for group-item pairs. However, the resulted group representations, as points, lack adequate flexibility and capacity to account for the multi-faceted user preferences. Also, the point embedding-based preference aggregation is a less faithful reflection of a group's decision-making process, where all users have to agree on a certain value in each embedding dimension instead of a negotiable interval. In this paper, we propose a novel representation of groups via the notion of hypercubes, which are subspaces containing innumerable points in the vector space. Specifically, we design the hypercube recommender (CubeRec) to adaptively learn group hypercubes from user embeddings with minimal information loss during preference aggregation, and to leverage a revamped distance metric to measure the affinity between group hypercubes and item points. Moreover, to counteract the long-standing issue of data sparsity in group recommendation, we make full use of the geometric expressiveness of hypercubes and innovatively incorporate self-supervision by intersecting two groups. Experiments on four real-world datasets have validated the superiority of CubeRec over state-of-the-art baselines.

CCS CONCEPTS

• Information systems → Recommender systems.

*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 ACM 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.
SIGIR '22, July 11–15, 2022, Madrid, Spain.

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-8732-3/22/07...\$15.00
<https://doi.org/10.1145/XXXXXX.XXXXXX>

KEYWORDS

Group Recommendation; Hypercube Representations; Self-supervised Learning

ACM Reference Format:

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. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/XXXXXX.XXXXXX>

1 INTRODUCTION

Group activities are natural habitats for humans, and there is no exception in e-commerce. For example, friends can sign up for the same music festivals on Facebook Events, and online platforms like Meetup allow users to form groups and host activities [42]. As traditional recommender systems only target at suggesting relevant items to individuals, there is an urgent demand in generating recommendations for a group of users, referred to as *group recommendation*. With the increasingly strong synergy between social and commercial functionalities in contemporary e-commerce platforms (e.g., Yelp, Gowalla, Steam, etc.) [40], there has been a recent surge in developing solutions to group recommendation [3, 19, 32].

Similar to conventional recommendation, latent factor models exhibit dominant performance in group recommendation. In a general sense, groups and items are mapped to points in the vector space (i.e., embeddings), then the pairwise group-item affinity can be estimated via intuitive distance/similarity metrics. As such, current group recommenders share a core motivation of learning representative group-level embeddings to capture the joint preferences of users. In rare cases where group-item interactions are dense, group representations can be directly learned via standard collaborative filtering models by substituting users with groups [41]. However, the real group activities are highly sparse [32, 42], restricting the informativeness of the learned group representations and hindering recommendations for new groups. Thus, a common practice is to aggregate a group's representation from its members' user embeddings, which are usually learned from the richer user-item interactions. Such aggregation strategies have evolved from heuristics like the least misery and maximum pleasure [1], to the more recent attentive sum [3] and social influence-based approach [41].

Despite the increasing complexity of these recent preference aggregation approaches, the learned group representations still face a major bottleneck in their expressiveness. This is mainly due

to the amount of information lost when compressing multiple user embeddings into one. In the context of embeddings, a group/user is represented as a point in the vector space, and each entry is the exact coordinate in the corresponding dimension. However, while each individual user’s preference tend to be stable and can be denoted by a certain value in each embedding dimension, the whole group’s overall preference reflects a blend of multiplex personal interests, challenging the capacity of point embeddings to accommodate such diversity. For example, if we associate the embedding dimensions with explicit item features, a group of users may land on different preferred values on the “price” dimension. But, it quickly becomes problematic when using a single point for the entire group due to the lack of considerations for such disparate and diverse user preferences, which commonly exist in web applications [26, 28].

In some cases, though higher dimensionality can be assigned to the group embedding to encode more information and alleviate the problem, this does not resolve the issue that a single embedding is unable to precisely reflect the decision-making process within a user group. Essentially, aggregating user preferences mimics the decision-making process [14] where all group members reach a consensus, reflected by a group-level representation. In reality, the group-level representation should involve synergies of common interests as well as compromises on some personal tastes. For example, for two users that have a large discrepancy on preferred prices, it is more realistic for them to agree upon an approximate price range in between, rather than an exact value. This, unfortunately, cannot be facilitated by existing aggregation schemes where all user preferences are merged into a single point. A point group embedding forces all users to agree upon a fixed decision in each dimension, no matter its value is simply averaged from all users [1] or biased towards more influential group members [3, 41].

To this end, we aim to represent the group-level preferences in a more expressive way, such that the multi-faceted user demands are thoroughly encoded, and the intricate decision-making process among all group members are faithfully preserved. Hence, we advocate addressing the restricted capacity of learned group representations by seeking an alternative to the outgoing point embeddings. Specifically, we introduce the notion of hypercubes as a new way to represent the aggregated group-level preferences in the vector space. In each dimension of the vector space, while a point embedding only has a fixed value, a hypercube covers a range of values that span across a continuous interval [30]. Combining the intervals will shape a “cube”-like subspace, enabling a hypercube to accommodate innumerable points in the vector space [46]. Compared with the traditional group embeddings, hypercubes enable a higher tolerance of different user preferences. Furthermore, as each dimension now encodes a preference interval, this paradigm is a more realistic reflection of how a group of distinct users give concessions to each other and eventually work out a combined criteria for desired items.

However, further challenges are in place when unleashing the full potential of hypercube group representations. Firstly, as traditional aggregation schemes are dedicated to learning groups’ embedding representations, a new paradigm is desired to effectively summarize all group members’ vectorized preferences into a hypercube with minimal information loss in user interests. Secondly, the utilization of group hypercubes voids the applicability of existing point-wise

distance measures like Euclidean distance or cosine similarity for group-item pairs, highlighting the need for a revamped distance metric that can precisely quantify and distinguish the affinity between different hypercubes (i.e., groups) and points (i.e., items). Thirdly, the long-standing problem of severe data sparsity with group-level interactions [32, 42] still persists in our case, which will impede the quality of learned group hypercube representations and amplify the risk of overfitting.

To address these issues, we present our solution to group recommendation, namely the *hypercube recommender (CubeRec)*. CubeRec inherits the classic point embeddings for individual users/items whose preferences/attributes are deterministic and can be learned from the relatively rich individual-level interactions. Meanwhile, we innovatively represent each group with the more expressive hypercube, which is first composed by the embeddings of its user members, and then optimized through observed group-level interactions. In CubeRec, we propose two alternative strategies to merge individual user embeddings into group hypercubes, specifically a geometric bounding approach that physically encapsulates all group members’ points, and an attentive approach that composes hypercubes by investigating the semantics within user embeddings. To allow for accurate item ranking, we put forward a distance function to measure the distance between each item point and group hypercube. The distance metric accounts for distances from both the inside and outside of a hypercube, thus accommodating different scenarios and maintaining a discriminative characteristic. Furthermore, we make full use of the geometry of group hypercubes and propose a novel self-supervised learning paradigm to tackle the data sparsity in group recommendation. In short, we define an intersection operation between two group representations, and leverage the common users shared between two groups as supervision signals to enrich the information encoded within every hypercube.

In summary, the main contributions of our work are:

- We define a new schema based on hypercubes to represent group-level preferences for group recommendation. Hypercubes bypass the limited expressiveness and flexibility of conventional point embeddings, allowing the group representations to encode multi-faceted user preferences while fully imitating a group’s decision-making process.
- We propose CubeRec, a group recommender that comes with novel designs for learning group-level hypercubes and measuring group-to-item distances. To alleviate data sparsity, CubeRec further incorporates self-supervision by innovatively utilizing the intersection between groups.
- We conduct extensive experiments on four real-world benchmark datasets. Experimental results show that CubeRec yields superior recommendation performance compared with state-of-the-art baselines.

2 PRELIMINARIES

In this section, we mathematically define the *hypercube representations* of groups, which are the key building block of CubeRec.

One can think of a hypercube as extending a rectangle into a higher-dimensional space, where each edge of the hypercube corresponds to a real-valued, closed interval on each dimension [30, 46]. For group \mathcal{G} which is a subset of users, we define its

representation \mathbb{G} via the following center-offset format:

$$\mathbb{G} \equiv \{\mathbf{g} \in \mathbb{R}^d : \mathbf{c} - \mathbf{o} \leq \mathbf{g} \leq \mathbf{c} + \mathbf{o}\}, \quad (1)$$

where $\mathbf{c} \in \mathbb{R}^d$ is the geometric center of the hypercube, and $\mathbf{o} \in \mathbb{R}_{\geq 0}^d$ is the non-negative offset vector. Both the center and offset are d -dimensional vectors learned within the same embedding space, where the center defines the Essentially, in the z -th dimension of \mathbb{R}^d , \mathbb{G} will have its z -th edge spanning across $[c_z - o_z, c_z + o_z]$. For notation convenience, we define the hypercube representation for the n -th group \mathcal{G}_n as a tuple $\mathbb{G}_n = (\mathbf{c}_n, \mathbf{o}_n)$ with its group-specific center \mathbf{c}_n and offset \mathbf{o}_n . We use this tuple as a shorthand for Eq.(1) in the rest parts of our paper. Compared with point embeddings, relaxing the fixed value in each dimension into a flexible range brings an increased tolerance on the variety among group members. It also reflects a more realistic decision process, as the group does not have to pre-approve a specific value per dimension before ranking candidate items.

3 THE HYPERCUBE RECOMMENDER

To provide an overview, CubeRec first learns user and item embeddings based on user-level interactions with items. Then, we propose our solution to building and learning hypercube representations at the group level. To further compensate for the high sparsity of group-item interactions when learning hypercube representations, we introduce a novel learning objective by leveraging self-supervision signals from overlapping groups. In what follows, we unfold the design details of CubeRec.

3.1 Learning User and Item Embeddings

In CubeRec, we employ LightGCN [15] to learn user and item embeddings. Compared with earlier collaborative filtering (CF) methods, adopting LightGCN facilitates the modelling of higher-order collaborative signals through message passing, and offers simplicity via the omission of excessive nonlinear components [38]. We keep this part brief as it is not our core contribution and can actually be replaced by arbitrary latent factor models.

Modelling User-Item Interactions. Treating each user/item as a distinct node, the user-item interactions can be formulated as a bipartite graph, where we let $\mathcal{N}(u_i)$ and $\mathcal{N}(v_j)$ denote the one-hop neighbor sets of user u_i and item v_j , respectively. The embedding of an arbitrary user/item is updated by propagating its neighbors' embeddings into u_i/v_j :

$$\mathbf{u}_i^{(l)} = \sum_{j \in \mathcal{G}(u_i)} \eta_{ij} \mathbf{v}_j^{(l-1)}, \quad \mathbf{v}_j^{(l)} = \sum_{i \in \mathcal{G}(v_j)} \eta_{ij} \mathbf{u}_i^{(l-1)}, \quad (2)$$

where $l \leq L$ denotes the propagation layer, $\eta_{ij} = (\sqrt{|\mathcal{G}(u_i)| \cdot |\mathcal{G}(v_j)|})^{-1}$ is a graph Laplacian normalization term [15, 23], while $\mathbf{u}_i^{(0)}$ and $\mathbf{v}_j^{(0)}$ are randomly initialized and will be trained via back-propagation. Following [15], the final embeddings are obtained via mean pooling across all layers, i.e., $\mathbf{u}_i = \frac{1}{L+1} \sum_{l=0}^L \mathbf{u}_i^{(l)} \in \mathbb{R}^d$ and $\mathbf{v}_j = \frac{1}{L+1} \sum_{l=0}^L \mathbf{v}_j^{(l)} \in \mathbb{R}^d$, which are used in the subsequent recommendation phases.

Handling Explicit Social Ties. Obviously, increasing L can help capture higher-order relationships including implicit social relationships like user-item-user paths, as L essentially controls the number of hops to be reached during the neighborhood aggregation.

As group recommendation services are commonly provided on social platforms, the explicitly available social connections between users (e.g., following relationship on social media) carry strong signals for understanding a user's preferences. However, the plain LightGCN does not account for such explicit social ties. In CubeRec, we use $\mathbf{R} \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{V}|}$ to denote the interactions between users from set \mathcal{U} and items from set \mathcal{V} , and $\mathbf{S} \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{U}|}$ to denote the social relation adjacency matrix of all users. Then, we formulate the adjacency matrix of the user-item graph as the following:

$$\mathbf{A} = \begin{bmatrix} \mathbf{S} & \mathbf{R} \\ \mathbf{R}^T & \mathbf{0} \end{bmatrix}, \quad (3)$$

where $\mathbf{0}$ is a $|\mathcal{V}| \times |\mathcal{V}|$ matrix consisting of 0s, and the resulted adjacency matrix $\mathbf{A} \in \{0, 1\}^{(|\mathcal{U}|+|\mathcal{V}|) \times (|\mathcal{U}|+|\mathcal{V}|)}$ now carries both user-item interactions and user-user social ties. With that, we summarize the socially enhanced LightGCN in the matrix form:

$$\mathbf{E}^{(l)} = (\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}) \mathbf{E}^{(l-1)}, \quad (4)$$

where $\mathbf{D} \in \mathbb{N}^{(|\mathcal{U}|+|\mathcal{V}|) \times (|\mathcal{U}|+|\mathcal{V}|)}$ is a diagonal degree matrix in correspondence to \mathbf{A} , while $\mathbf{E} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{V}|) \times d}$ stacks all user and item embeddings.

User-level Recommendation Loss. We optimize all user and item embeddings with a distance-based objective on user-level interactions, namely the hinge loss [6, 18]:

$$L_{user} = \sum_{(u_i, v_{j^+}, v_{j^-}) \in \mathcal{D}_{user}} \max(0, \lambda + d_{ij^+} - d_{ij^-}), \quad (5)$$

where \mathcal{D}_{user} denotes the training set, λ is the safety margin size to be predefined, and $d_{ij} = \|\mathbf{u}_i - \mathbf{v}_j\|_2^2$ is the squared Euclidean distance between the embeddings of u_i and v_j . The rationale of the hinge loss is that, in the vector space, each user u_i should always be closer to a visited item v_{j^+} than an unvisited one v_{j^-} . As such, a training sample is defined as a triple (u_i, v_{j^+}, v_{j^-}) , where every observed user-item pair (u_i, v_{j^+}) is matched with a fixed amount of uniformly sampled negative items v_{j^-} to construct \mathcal{D}_{user} .

3.2 Composing Group-level Hypercubes

We perform pretraining with L_{user} to acquire initial user and item embeddings for making recommendations at the group level. As a common practice, a group's representation can be learned by merging all the embeddings of its members via means like attention networks [3, 20] or graph convolution [13, 37], but they inevitably incur limited expressiveness. As discussed earlier, compressing users' multi-faceted preferences into a single point in the vector space leads to significant information loss, and also misaligns with a group's decision-making process in the real world.

In this section, we present our solution to learning expressive and flexible group representations with hypercubes. Specifically, we introduce two different strategies for composing group-level hypercubes using individual users' point embeddings, of which both have their unique properties and are further testified in Section 4.

Geometric Bounding and Projection. Given s user embeddings $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_s \in \mathbb{R}^d$ in group \mathcal{G}_n , the geometric bounding operation is an intuitive way to generate the corresponding group-level hypercube. As Figure 1(a) shows, we first need to find the smallest hypercube covering all s points (i.e., user embeddings) from \mathcal{G}_n .




Figure 1: A schematic view of the key hypercube-based computations in CubeRec. We use \mathbb{R}^2 for demonstration purpose. Corresponding details can be found in Section 3.2 for (a) and (b), Section 3.3 for (c), and Section 3.4 for (d).

Such a hypercube is computed as the following:

$$\tilde{\mathbb{G}}_n = (\tilde{\mathbf{c}}_n, \tilde{\mathbf{o}}_n) = \left(\frac{\mathbf{u}^{max} + \mathbf{u}^{min}}{2}, \frac{|\mathbf{u}^{max} - \mathbf{u}^{min}|}{2} \right), \quad (6)$$

where \mathbf{u}^{max} and \mathbf{u}^{min} are boundary vectors computed via:

$$\begin{aligned} \mathbf{u}^{max} &= \max(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_s), \\ \mathbf{u}^{min} &= \min(\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_s), \end{aligned} \quad (7)$$

where $\max(\cdot)$ and $\min(\cdot)$ both operate element-wise. However, straightforwardly utilizing $\tilde{\mathbb{G}}_n$ as the final group representation is suboptimal due to its geometric property. That is, the more diverse that the users in \mathcal{G}_n are, the larger area that $\tilde{\mathbb{G}}_n$ will cover in the d -dimensional space. As the hypercube shares the same vector space with items, the recommendation will naturally prefer items that falls in the hypercube because it means those items meet all the preference intervals on all latent dimensions. Consequently, a highly diverse group will receive an excessively large hypercube, covering more false positive items that eventually dilute the recommendation accuracy. To alleviate this issue, we learn projection matrices for both the center and offset that help rescale $\tilde{\mathbb{G}}_n$ into the final group hypercube \mathbb{G}_n :

$$\mathbb{G}_n = (\mathbf{c}_n, \mathbf{o}_n) = (\mathbf{W}_c \tilde{\mathbf{c}}_n, \mathbf{W}_o \tilde{\mathbf{o}}_n), \quad (8)$$

where $\mathbf{W}_c \in \mathbb{R}^{d \times d}$ and $\mathbf{W}_o \in \mathbb{R}_{\geq 0}^{d \times d}$ are the corresponding projection weights. Note that \mathbf{W}_o is non-negative such that $\mathbf{o}_n \geq 0$.

Attentive Fusion and Projection. Geometric bounding aims to let a hypercube fairly cover all users' preferences within a group. While this is reasonable in many cases, recent studies also point out that each user in the group may affect the final group decision to a different extent [3, 41] due to varied social influence and/or expertise. To make our learned group representations able to account for every group member's importance discrepancy, we further propose an attentive approach [8, 35] for composing group hypercubes:

$$\tilde{\mathbf{c}}_n = \mathbf{W}_V \mathbf{U}_n \cdot \text{softmax} \left(\frac{(\mathbf{W}_K \mathbf{U}_n)^T \cdot \mathbf{q}}{\sqrt{d}} \right), \quad (9)$$

where $\mathbf{q} \in \mathbb{R}^d$ is the learnable query vector in the above self-attention, $\mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{d \times d}$ are respectively the key and value projection matrices, and $\mathbf{U}_n \in \mathbb{R}^{d \times s}$ stacks all d -dimensional user embeddings in group \mathcal{G}_n . As each $\mathbf{u} \in \mathbf{U}_n$ reflects a user's exact interest, the aggregation in Eq.(9) attentively aggregates all group members' interests, denoted by $\tilde{\mathbf{c}}_n$. Then, we locate the centroid

and offset of the composed hypercube \mathbb{G}_n via:

$$\mathbb{G}_n = (\mathbf{c}_n, \mathbf{o}_n) = (\mathbf{W}_c \tilde{\mathbf{c}}_n, \text{ReLU}(\mathbf{W}_o \tilde{\mathbf{o}}_n + \mathbf{b}_o)), \quad (10)$$

where we slightly abuse notations $\mathbf{W}_c, \mathbf{W}_o \in \mathbb{R}^{d \times d}$ to denote the learnable matrices, and $\mathbf{b}_o \in \mathbb{R}^d$ is the bias vector. Since the attentively aggregated $\tilde{\mathbf{c}}_n$ summarizes the core interest of the group, we directly use its linear projection as the center \mathbf{c}_n . Meanwhile, the process of obtaining $\tilde{\mathbf{c}}_n$ sacrifices users' preference diversity, which should be preserved by the offset. Thus, we employ the more expressive nonlinearity to partially infer and recover such missing information from $\tilde{\mathbf{c}}_n$ and obtain the offset \mathbf{o}_n . The rectified linear unit (ReLU) is used to ensure non-negativity. This process is depicted by Figure 1(b). In short, the attentive approach focuses more on the semantics of user embeddings compared with its geometric counterpart that retains more physical meanings, and is relatively more efficient.

3.3 Learning Hypercubes for Group Recommendation

To optimize the composed hypercube representations, we firstly put forward a distance function that quantifies the affinity between each pair of group and item, and then introduce how we define the loss function at the group level for learning CubeRec.

Computing Item-to-Group Distances. Compared with conventional point group embeddings, each edge of a hypercube defines a range on the corresponding dimension in the latent feature space rather than a single value. Intuitively, for any item point \mathbf{v}_j that falls in the hypercube \mathbb{G}_n , it means each dimension of \mathbf{v}_j satisfies the preference range specified by \mathbb{G}_n , making item v_j a good fit for group \mathcal{G}_n . However, it is not always the case in practice. On one hand, a diverse group tend to have a relatively large coverage with the generated hypercube, containing more irrelevant items. On the other hand, a group with relatively narrow interests can have a rather small hypercube, making it hard to cover any existing item embeddings, especially when d is large. In light of this, we propose an approach that measures the item-to-group distance d_{nj} via:

$$d_{nj} = f_{out}(\mathbb{G}_n, \mathbf{v}_j) + \gamma f_{in}(\mathbb{G}_n, \mathbf{v}_j), \quad (11)$$

where functions $f_{out}(\cdot, \cdot)$ and $f_{in}(\cdot, \cdot)$ measure the outer and inner item-to-group distances, respectively, which are balanced by a fixed coefficient $\gamma \geq 0$. Taking the group hypercube \mathbb{G}_n and item embedding \mathbf{v}_j as inputs, these two functions are formulated as follows

using squared Euclidean distance:

$$\begin{aligned} f_{out}(\mathbb{G}_n, \mathbf{v}_j) &= \|\max(\mathbf{v}_j - \mathbf{g}_n^{\mathbf{p}}, \mathbf{0}) + \max(\mathbf{g}_n^{\mathbf{a}} - \mathbf{v}_j, \mathbf{0})\|_2^2, \\ f_{in}(\mathbb{G}_n, \mathbf{v}_j) &= \|\mathbf{c}_n - \min(\mathbf{g}_n^{\mathbf{p}}, \max(\mathbf{g}_n^{\mathbf{a}}, \mathbf{v}_j))\|_2^2, \end{aligned} \quad (12)$$

where two points $\mathbf{g}_n^{\mathbf{a}}, \mathbf{g}_n^{\mathbf{p}} \in \mathbb{R}^d$ are geometrically the lower-left and upper-right corners of hypercube \mathbb{G}_n , respectively:

$$\mathbf{g}_n^{\mathbf{a}} = \mathbf{c}_n - \mathbf{o}_n, \quad \mathbf{g}_n^{\mathbf{p}} = \mathbf{c}_n + \mathbf{o}_n. \quad (13)$$

As explained in Figure 1(c), assuming \mathbf{v}_j locates outside \mathbb{G}_n , what Eq.(12) does is to: (1) locate the anchor point on \mathbb{G}_n 's surface that is the closest to \mathbf{v}_j ; (2) use $f_{out}(\cdot, \cdot)$ to measure the outer distance from \mathbf{v}_j to this anchor point; and (3) use $f_{in}(\cdot, \cdot)$ to measure the inner distance from the centroid \mathbf{c}_n to this anchor point. The outer and inner distances are then combined into d_{nj} via Eq.(11). When \mathbf{v}_j is inside or on the surface of \mathbb{G}_n , $f_{out}(\mathbb{G}_n, \mathbf{v}_j) = 0$, and $f_{in}(\mathbb{G}_n, \mathbf{v}_j)$ will reduce to the distance between item and centroid because $\min(\mathbf{g}_n^{\mathbf{p}}, \max(\mathbf{g}_n^{\mathbf{a}}, \mathbf{v}_j)) = \mathbf{v}_j$. This is also shown in Figure 1(c).

Group-level Recommendation Loss. Following the definition of item-to-group distance in Section 3.2, the higher the affinity between group \mathbb{G}_n and item v_j , the smaller the distance d_{nj} will be. Analogously to Eq.(5), we adopt the following hinge loss:

$$L_{group} = \sum_{(\mathbb{G}_n, v_{j^+}, v_{j^-}) \in \mathcal{D}_{group}} \max(0, \lambda' + d_{nj^+} - d_{nj^-}), \quad (14)$$

where λ' denotes the safety margin, $(u_i, v_{j^+}, v_{j^-}) \in \mathcal{D}_{group}$ are training instances constructed by negative sampling. Essentially, for cube \mathbb{G}_n , L_{group} aims to pull a positive item's embedding \mathbf{v}_{j^+} into \mathbb{G}_n and closer to its center, while pushing the negative item's embeddings \mathbf{v}_{j^-} away from it.

3.4 Self-supervision via Hypercube Intersections

Unfortunately, as pointed out by prior studies [32, 41], a long-lasting obstacle for learning high-quality group representations is that, the learning of group representations is purely dependant on the interacted items, which are highly sparse at the group level (see statistics in Section 4.1). As such, we propose to take advantage of the rich self-supervision signals from user-group associations for learning representative group hypercubes. Specifically, for group \mathbb{G}_n we firstly draw a different group $\mathbb{G}_{n'}$, where $\mathbb{G}_{n'}$ shares at least one user in common with \mathbb{G}_n (details on how we deal with disjoint groups will follow). Users shared between \mathbb{G}_n and $\mathbb{G}_{n'}$ are termed *relay users*. Then intuitively, the intersection of two group hypercubes denoted by $int(\mathbb{G}_n, \mathbb{G}_{n'})$, which is also a hypercube, is supposed to reflect the properties of all the relay users.

With Figure 1(d), we illustrate our goal of self-supervised learning in CubeRec. In a geometric sense, two hypercubes' intersection is ideally the subspace that hosts the point embeddings of relay users. On this basis, we formulate the self-supervision loss below:

$$L_{self} = \sum_{\forall \mathbb{G}_n} \sum_{\substack{\forall u_{i^+} \in (\mathbb{G}_n \cap \mathbb{G}_{n'}) \\ u_{i^-} \sim p(u_i^- | \mathbb{G}_n, \mathbb{G}_{n'})}} \max(0, \lambda'' + d_{nn'i^+} - d_{nn'i^-}), \quad (15)$$

where $d_{nn'i}$ is a shorthand for the distance between the intersection and user u_i , i.e., $f_{out}(int(\mathbb{G}_n, \mathbb{G}_{n'}), \mathbf{u}_i) + \gamma f_{in}(int(\mathbb{G}_n, \mathbb{G}_{n'}), \mathbf{u}_i)$. For each positive user u_{i^+} that falls in the group intersection $\mathbb{G}_n \cap \mathbb{G}_{n'}$, we sample a negative user u_{i^-} from $\mathcal{U} \setminus (\mathbb{G}_n \cap \mathbb{G}_{n'})$ with uniform

distribution $p(u_{i^-} | \mathbb{G}_n, \mathbb{G}_{n'}) = |\mathcal{U} \setminus (\mathbb{G}_n \cap \mathbb{G}_{n'})|^{-1}$. By doing so, the representations of groups \mathbb{G}_n and $\mathbb{G}_{n'}$ will be regularized by the relay users, so as to capture the common interests between them. Correspondingly, the learned group representations will be substantially more informative than those learned purely with sparse item-level interactions.

Hypercube Intersection Operation. We hereby define how the hypercube intersection $int(\mathbb{G}_n, \mathbb{G}_{n'})$ is calculated. For convenience, we let $int(\mathbb{G}_n, \mathbb{G}_{n'}) = \mathbb{G}_{nn'} = (\mathbf{c}_{nn'}, \mathbf{o}_{nn'})$. Because straightforwardly using the geometric intersection is only applicable when two hypercubes have actual overlaps in each dimension of \mathbb{R}^d , we resort to a relaxed intersection computation via a neural approach. Firstly, center $\mathbf{c}_{nn'}$ is computed by performing element-wise attention over the two group centers:

$$\begin{aligned} \mathbf{c}_{nn'} &= \mathbf{a}_n \odot \mathbf{c}_n + \mathbf{a}_{n'} \odot \mathbf{c}_{n'}, \\ \mathbf{a}_* &= \frac{\exp(\phi(\mathbf{c}_*))}{\exp(\phi(\mathbf{c}_n)) + \exp(\phi(\mathbf{c}_{n'}))} \quad \text{for } * \in \{n, n'\}, \end{aligned} \quad (16)$$

where \odot denotes element-wise multiplication, $\phi(\cdot) : \mathbb{R}^d \mapsto \mathbb{R}^d$ is a multilayer perceptron (MLP) with d neurons throughout all layers. Meanwhile, the offset $\mathbf{o}_{nn'}$ is determined via:

$$\mathbf{o}_{nn'} = \min(\mathbf{o}_n, \mathbf{o}_{n'}) \odot \text{sigmoid}(\psi(\mathbf{o}_n + \mathbf{o}_{n'})), \quad (17)$$

where $\psi(\cdot) : \mathbb{R}^d \mapsto \mathbb{R}^d$ is another MLP with a distinct set of parameters. Essentially, Eq.(16) places the intersection center $\mathbf{c}_{nn'}$ somewhere between \mathbf{c}_n and $\mathbf{c}_{n'}$. Then, with the scaling effect of sigmoid function, Eq.(17) shrinks the offsets of hypercubes \mathbb{G}_n and $\mathbb{G}_{n'}$ to obtain the new offset of the intersection $\mathbf{o}_{nn'}$.

Handling Isolated Groups. Despite the heavily overlapping nature of user groups on social e-commerce platforms, we must take into account the situation where a group \mathbb{G}_n might be disjoint with all other groups in the dataset. On this occasion, we propose the two following strategies to generate a dummy user group $\tilde{\mathbb{G}}_{n'}$ having overlapping users with \mathbb{G}_n :

- I. Proportional Swap (PS): With a specified proportion ρ , we swap $\rho|\mathbb{G}_n|$ users in \mathbb{G}_n with users uniformly sampled from $\mathcal{U} \setminus \mathbb{G}_n$ with replacement. We denote this as $\tilde{\mathbb{G}}_{n'} = \text{PS}(\mathbb{G}_n)$.
- II. Proportional Imputation (PI): With a specified proportion ρ , we inject $\rho|\mathbb{G}_n|$ uniformly sampled users from $\mathcal{U} \setminus \mathbb{G}_n$ (with replacement) into \mathbb{G}_n . We denote this as $\tilde{\mathbb{G}}_{n'} = \text{PI}(\mathbb{G}_n)$.

The generated $\tilde{\mathbb{G}}_{n'}$ then serves as $\mathbb{G}_{n'}$ to facilitate self-supervision via Eq.(15). We adopt a predefined value $\rho = 0.5$ in our approach, where we use PS and PI alternately during training.

3.5 Optimizing CubeRec

As all components of CubeRec are end-to-end differentiable, we learn the model parameters on multiple objectives with the mini-batch stochastic gradient descent algorithm Adam [22]. As depicted by Algorithm 1, we adopt a two-stage training strategy for CubeRec. To be specific, we first obtain pretrained user and item embeddings by optimizing the user-level loss L_{user} . Then, we fine-tune the pretrained embeddings and learn rest of the parameters in CubeRec by optimizing towards the combined loss $L_{group} + \mu L_{self}$ ($\mu \geq 0$). In both training stages, we update corresponding model parameters in each iteration and repeat the entire training process until the loss converges or is sufficiently small.

Algorithm 1 Training Procedure of CubeRec

```
1: Input:  $\mathcal{U}, \mathcal{G}, \mathcal{V}$ 
2: Output: All model parameters collectively referred to as  $\Theta$ 
3: Randomly initialize  $\Theta$ ;
4: repeat
5:   Draw a mini-batch from  $\mathcal{D}_{user}$ ;
6:   Take a gradient step to update  $\mathbf{u}$  and  $\mathbf{v}$  w.r.t.  $L_{user}$ ;
7: until convergence  $\triangleright$  Pretraining user and item embeddings
8: repeat
9:   Draw a mini-batch from  $\mathcal{D}_{group}$  and compute  $L_{group}$ ;
10:  for each  $\mathcal{G}_n$  in mini-batch do
11:    if  $\exists \mathcal{G}_{n'} \text{ s.t. } \mathcal{G}_n \cap \mathcal{G}_{n'} \neq \emptyset$  then
12:      Compute  $L_{self}$  with a sampled  $\mathcal{G}_{n'}$ ;
13:    else
14:      Generate  $\tilde{\mathcal{G}}_{n'}$  via either PS( $\mathcal{G}_n$ ) or PI( $\mathcal{G}_n$ ) by coin-flipping;
15:      Compute  $L_{self}$  with  $\tilde{\mathcal{G}}_{n'}$ ;
16:  Take a gradient step to update  $\Theta$  w.r.t.  $L_{group} + \mu L_{self}$ ;
17: until convergence
```

We tune the hyperparameters using grid search. Specifically, the latent dimension d is searched in $\{16, 32, 64, 128, 256\}$, and the distance coefficient γ and self-supervision weight μ are searched in $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. For optimizing the three losses L_{user} , L_{group} , L_{self} , we draw 5 negative samples for each positive ground truth cases to construct corresponding training sets. The safety margins on those losses are set to $\lambda = \lambda' = \lambda'' = 0.5$ following common practices for hinge loss [6, 46]. We set the the learning rate to 1×10^{-3} and batch size to 256 according to device capacity. To prevent overfitting, we adopt a dropout [34] ratio of 0.2 on all deep layers of CubeRec during training.

4 EXPERIMENTS

In this section, we conduct experiments to verify the effectiveness of CubeRec in group recommendation. Specifically, we aim to answer the following research questions (RQs):

RQ1: Is CubeRec the new state-of-the-art?

RQ2: Are the major components proposed in CubeRec effective?

RQ3: How does CubeRec perform w.r.t. different group sizes?

RQ4: What is the impact of CubeRec’s key hyperparameters?

4.1 Datasets

We adopt four real-world datasets collected from different event-based social networks, namely **Meetup**, **Yelp**, **Gowalla**, and **Douban**. Both Meetup and Douban contain group events held in different venues respectively in New York and Beijing, where event venues are items to be recommended. Meanwhile, Yelp and Gowalla are typical check-in datasets on different restaurants (i.e., items in our

Table 1: Statistics of experimental datasets.

	Meetup	Yelp	Gowalla	Douban
#users	24,631	34,504	60,805	70,743
#groups	13,552	24,103	78,453	109,538
#items	19,031	22,611	8,984	60,028
#user-item interactions	126,813	482,273	1,625,817	3,422,266
#group-item interactions	19,031	26,883	208,336	164,153
average #items per user	5.15	13.98	26.74	48.38
average #items per group	1.40	1.12	2.66	1.50
average #groups per user	4.83	3.11	3.23	7.53
average group size	8.79	4.45	2.31	4.86

case) based in the US. Since Yelp and Gowalla do not originally contain group information, we follow the widely adopted procedure [32, 41] to construct group interactions by finding overlaps on both check-in times and social relations. That is, we assume if a set of users who are connected on the social network visit the same venue at the same time, then they are regarded as members of a group, and the corresponding activities are group activities [41]. The use of explicit social connections and spatiotemporal tags ensures the quality of discovered user groups in Yelp and Gowalla. We provide the key statistics of the four datasets in Table 1. Each dataset is split with a ratio of 8:1:1 for training, validation, and test, respectively.

4.2 Baselines and Evaluation Protocols

We compare CubeRec¹ with the following state-of-the-art baselines:

- **MF-AVG** [36]: This baseline takes the average of all group-wise user representations learned via user-item matrix factorization (MF) [24] to compose group embeddings.
- **NCF-AVG** [16]: It also represents groups with average pooling on user embeddings, where user embeddings are learned via the neural collaborative filtering instead of MF.
- **AGREE** [3]: This approach utilizes attention networks, thus accounting for different importance of group members when learning group embeddings.
- **SIGR** [41]: It exploits and integrates users’ global and local social influence to improve the group recommendation.
- **CAGR** [42]: This method firstly learns centrality-aware user representations, and then learns a group recommender via two-stage optimization.
- **GroupIM** [32]: This state-of-the-art group recommender adopts mutual information maximization between users and groups to overcome the sparse group-level interactions.

It is worth noting that in CubeRec, we implement both the *geometric* and *attentive* approaches described in Section 3.2 for composing group hypercubes, which are respectively marked as **CubeRec-G** and **CubeRec-A**. We set $L = 3$ for learning user and item embeddings with LightGCN, and implement the two MLPs for hypercube intersection with a 3-layer setting and ReLU activation. Based on the hyperparameter search described in Section 3.5, we fix $d = 64$, $\gamma = 0.3$, and $\mu = 0.7$ across all datasets. The impact of these hyperparameters to the recommendation performance will be further discussed in Section 4.6.

We leverage two metrics, namely recall at rank K (Recall@ K) [7, 10] and normalized discounted cumulative gain at rank K (NDCG@ K) [9, 47] that are widely adopted in recommendation research. We adopt $K = 10, 20$ where all items unvisited by each group are taken as negative samples for evaluation. In short, Recall@ K measures the ratio of the ground truth items that are present on the top- K list, and NDCG@ K evaluates whether the model can rank the ground truth items as highly as possible.

4.3 Recommendation Effectiveness (RQ1)

Table 2 summarizes the performance comparison among all the group recommenders, where each method’s results are averaged over five runs on every dataset. With the recommendation results, we discuss our key findings below.

¹We release our implementation at: <https://github.com/jinglong0407/CubeRec.git>

Table 2: Recommendation results. Numbers in bold face are the best results for corresponding metrics.

Method	Yelp				Douban				Gowalla				Meetup			
	Recall@K		NDCG@K		Recall@K		NDCG@K		Recall@K		NDCG@K		Recall@K		NDCG@K	
	K = 10	K = 20	K = 10	K = 20	K = 10	K = 20	K = 10	K = 20	K = 10	K = 20	K = 10	K = 20	K = 10	K = 20	K = 10	K = 20
NCF-AVG [16]	0.0042	0.0090	0.0018	0.0030	0.0023	0.0042	0.0010	0.0015	0.0146	0.0222	0.0088	0.0108	0.0087	0.0092	0.0074	0.0075
FM-AVG [36]	0.0161	0.0249	0.0075	0.0097	0.0032	0.0065	0.0014	0.0023	0.0127	0.0212	0.0061	0.0082	0.0128	0.0232	0.0063	0.0089
AGREE [3]	0.0168	0.0286	0.0073	0.0103	0.0031	0.0047	0.0013	0.0017	0.0125	0.0228	0.0063	0.0088	0.0081	0.0117	0.0033	0.0041
SIGR [41]	0.0194	0.0311	0.0082	0.0114	0.0035	0.0056	0.0016	0.0020	0.0143	0.0266	0.0070	0.0099	0.0090	0.0139	0.0038	0.0049
CAGR [42]	0.0215	0.0345	0.0106	0.0138	0.0026	0.0072	0.0010	0.0021	0.0169	0.0320	0.0090	0.0129	0.0028	0.0029	0.0012	0.0013
GroupIM [32]	0.0424	0.0535	0.0207	0.0235	0.0212	0.0316	0.0098	0.0124	0.0281	0.0384	0.0113	0.0136	0.0407	0.0525	0.0237	0.0266
CubeRec-G	0.0455	0.0592	0.0241	0.0298	0.0243	0.0354	0.0114	0.0167	0.0361	0.0575	0.0171	0.0194	0.0437	0.0616	0.0261	0.0306
CubeRec-A	0.0417	0.0520	0.0206	0.0224	0.0211	0.0309	0.0098	0.0158	0.0343	0.0495	0.0135	0.0164	0.0418	0.0525	0.0202	0.0211




Figure 2: Visualization of CubeRec-G (left) and CubeRec-A (right) for the same user group and recommended positive item on Meetup. Note that only the first two dimensions of \mathbb{R}^d are used to facilitate visualization. The amount of visualized negative items is different as all representations are learned from two independent vector spaces, and embeddings out of the $[-3, 3]$ range are clipped out to ensure clarity.

Apparently, CubeRec-G consistently outperforms all baselines by a significant margin. Compared with GroupIM which is the best baseline, the average improvements on Recall@10 and NDCG@10 brought by CubeRec-G are 14.4% and 23.6%, respectively. Unlike the straightforward mean pooling from user embeddings in NCF-AVG and FM-AVG, the selective preference aggregation schemes in AGREE, SIGR, CAGR and GroupIM generally achieve better performance when learning group-level representations for recommendation. However, these baselines are still subject to limited flexibility and resilience of the point embeddings used for representing groups. In contrast, our hypercube-based group representations can bypass the inherent limitations of point embeddings, thus achieving the best recommendation effectiveness.

Meanwhile, although CubeRec-A is not the best-performing variant of CubeRec, it still yields highly competitive results which are on par or stronger than the best baseline GroupIM. On one hand, its strong performance further demonstrates the benefit of replacing point embeddings with the more expressive hypercube representations. On the other hand, a possible reason that CubeRec-G is more advantageous than CubeRec-A is that, the group hypercubes learned via geometric bounding has a higher tolerance to the multi-faceted user preferences, and are overall more comprehensive representations of group-level interests. To qualitatively compare CubeRec-G and CubeRec-A, in Figure 2, we visualize the learned hypercubes of a randomly picked group from Meetup. For clarity, we only use the first two dimensions of \mathbb{R}^d . The hypercubes and

Table 3: Ablation test with different model architectures (Recall@10 is demonstrated). Numbers in bold face are the best results from each model, and “↓” marks a severe (over 5%) performance drop compared with the best results.

Method	Architecture	Yelp	Douban	Gowalla	Meetup
CubeRec-G	Default	0.0455	0.0243	0.0361	0.0437
	Remove SR	0.0440	0.0222↓	0.0353	0.0414↓
	Point Distance	0.0204↓	0.0098↓	0.0156↓	0.0181↓
	Remove SSL	0.0442	0.0232	0.0342↓	0.0407↓
CubeRec-A	Default	0.0417	0.0211	0.0343	0.0418
	Remove SR	0.0391↓	0.0209	0.0326	0.0371↓
	Point Distance	0.0202↓	0.0099↓	0.0159↓	0.0163↓
	Remove SSL	0.0382↓	0.0195↓	0.0342	0.0399

users in two plots belong to the same group in the dataset, and the same positive item has been successfully recommended in the top-10 list in both evaluation cases. As the visualization suggests, the hypercube learned by CubeRec-G is a tighter and more inclusive representation of all its members’ preferences. Also, CubeRec-G appears to be more capable of pulling the positive item into the group hypercube than CubeRec-A, leading to better recommendation performance.

Another observation is the relatively inferior performance of all methods on Douban compared with other datasets. Despite the dense user-item interactions for learning individual users’ preferences, the recommendation performance is largely impaired by the sparse group-item interactions and the largest item set. It is worth mentioning that methods with self-supervised learning (i.e., GroupIM and CubeRec) are still able to leverage augmented supervision signals to learn quality group representations, providing the best recommendation results even for this challenging dataset.

4.4 Ablation Study (RQ2)

To better understand the performance gain from the core components proposed in CubeRec, we perform ablation analysis on different degraded versions of CubeRec. Table 3 summarizes the recommendation outcomes in terms of Recall@10. In what follows, we introduce all variants and discuss the effectiveness of corresponding model components.

Removing Social Relations (Remove SR). In CubeRec, we infuse social relations when modelling user preferences via Eq.(3). To testify the usefulness of social relations in CubeRec, we remove the social relation matrix \mathbf{S} from the adjacency matrix by replacing




Figure 3: Recommendation performance w.r.t. different group sizes.




Figure 4: Recommendation performance w.r.t. different hyperparameters.

it with a $|\mathcal{U}| \times |\mathcal{U}|$ matrix filled with 0s. The resulted variant models, have experienced noticeable performance drops for both CubeRec-G and CubeRec-A, especially on Meetup dataset with the highest sparsity of user-item interactions. Hence, the inclusion of social relations is important for CubeRec to learn representative user embeddings, which are the building block for group hypercubes.

Using Only The Distance Between Item and Group Center Points (Point Distance). A crucial difference between CubeRec and point embedding-based group recommenders is the use of its hypercubes. Essentially, a group hypercube specifies a subspace in \mathbb{R}^d , where the recommended item points either locate in this subspace or maintain short distances with it. As a core ranking mechanism for group-item pairs in CubeRec, we verify the efficacy of such point-to-hypercube distance metric by replacing Eq.(11) with the conventional point distance. That is, $d_{nj} = \|\mathbf{c}_n - \mathbf{v}_j\|_2^2$ where squared Euclidean distance between the group center and item points are used. As a result, there is a significant decrease in recommendation for CubeRec across all datasets. This reflects that our revamped distance metric in CubeRec is a highly distinctive measurement for the affinity between item and groups.

Removing Self-supervised Learning (Remove SSL). This variant disables the self-supervised loss when optimizing CubeRec, i.e., the model is only optimized towards L_{group} instead of $L_{group} + \mu L_{self}$. As CubeRec can no longer leverage the common users shared between groups as a supervision signal to counteract the group-level data sparsity and regularize the learned group hypercubes, it suffers from inferior recommendation accuracy on all datasets. Specifically, significant performance drops (over 5%) are observed from CubeRec-G on Gowalla and Meetup, and from CubeRec-A on Yelp and Douban. As such, the novel self-supervised learning scheme showcases its strong contribution to the performance gain in CubeRec.

4.5 Effect of Different Group Sizes (RQ3)

In group recommendation, different group sizes will bring varied impact to the recommendation accuracy, and it is crucial for a recommender to keep robust in different circumstances. In this study, we further evaluate the performance of CubeRec on various group sizes. Following [32], we segment all groups into 5 categories, namely groups with 2-3, 3-4, 6-7, 8-9, and 10 or more users. We additionally choose GroupIM as the peer method for comparison because it is the most performant baseline.

We summarize the recommendation results in Figure 3, where Recall@10 is used for benchmarking. The first observation is that, CubeRec-G outperforms CubeRec-A and GroupIM in almost all scenarios. This also supports the superior performance of CubeRec-G given its capability to deal with different group sizes. Secondly, CubeRec-G yields more performance gain compared with GroupIM when the group sizes are relatively small (i.e., between 2 and 5). This is especially important for group recommendation, as the group sizes follow the typical long-tail distribution in all our datasets, considering it is more natural and feasible for users to form smaller groups in real life. Thirdly, we find that in general, group sizes are positively associated with the recommendation accuracy. Though smaller groups may have weaker discrepancies in user preferences during decision-making, the learning of group-level representations are also restricted by the limited user-item interactions. When the group size grows, the composed group representations tend to encode richer predictive signals for making accurate recommendations. Meanwhile, in some cases like Yelp and Gowalla, when the group size exceeds 9, the recommendation performance starts decreasing. This is largely associated with the excessive noise introduced by the diverse group members.

4.6 Hyperparameter Sensitivity (RQ4)

We answer RQ4 by investigating the performance fluctuations of CubeRec with varying hyperparameters, in particular the latent dimension d , trade-off γ between inner and outer distances in Eq.(11), and the self-supervision strength μ . Based on the standard hyperparameter setup in Section 4.2, we tune the value of one hyperparameter while keeping the others unchanged, and record the new results achieved in Figure 4. Note that we only showcase the hyperparameter sensitivity for CubeRec-G since it is the best-performing of the two, and CubeRec-A exhibits a highly similar trend as CubeRec-G.

Impact of d . The value of d is examined in $\{16, 32, 64, 128, 256\}$. In general, CubeRec benefits from a relatively larger d on all four datasets. But noticeably, the performance improvement stops when d reaches a certain size (64 and 128 in our case) due to overfitting.

Impact of γ . We also study the impact of $\gamma \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ that weighs the contribution of inner distance in our revamped group-to-item distance metric. As γ increases from 0.1 to 0.3, there is a generally upward trend in CubeRec’s performance. However, when γ exceeds 0.5, the performance starts to decrease. Hence, in the revamped distance metric, accounting for slightly more outer distance offers higher resolution when distinguishing a group’s preferences towards different items.

Impact of μ . The self-supervised learning loss is multiplied by μ to control its regularization effect. When the value of μ varies in $\{0.1, 0.3, 0.5, 0.7, 0.9\}$, the best performance of CubeRec is observed when $\mu = 0.7$ in most scenarios, suggesting the necessity of utilizing our proposed self-supervision to counteract the sparsity of group-item interactions. However, CubeRec achieves better performance when $\mu = 0.3$ on Yelp, which possibly attributes to the smallest average group size in this dataset (3.11 users per group), hindering the construction of self-supervision signals via relay users.

5 RELATED WORK

In general, group recommendation methods are designed in different ways to suit two types of audiences [29, 32, 41]: persistent groups and occasional groups. While persistent group recommendation assumes dense interaction records at the group level, our work addresses occasional group recommendation, which is more practical and widely studied as real user groups tend to be formed ad-hoc with limited group-item interactions [41]. Compared with persistent group recommendation where traditional collaborative filtering can be directly applied on group-item interactions [19, 33], a common practice in occasional group recommendation is to first learn individuals’ preferences from user-level interactions, then perform preference aggregation to infer the overall interest of a group. Besides, a widely acknowledged challenge [32, 42] in occasional group recommendation is the highly sparse group-level interactions. Hence, we review two lines of research that contributes to group recommendation, namely methods for preference aggregation and counteracting data sparsity.

5.1 Preference Aggregation

In group recommendation, all group members’ preferences can be aggregated via either *late aggregation* or *early aggregation*. The aim of late aggregation [1] is to generate item recommendations or predicted item-wise affinity scores for each group member, which are combined at the output stage via predefined strategies to produce recommendation results for the group. In this category, the most representative approaches are average satisfaction, least misery and maximum pleasure [1, 2, 29]. Unfortunately, these predefined aggregation strategies heavily rely on heuristics, thus lacking the desired expressiveness and flexibility for optimal performance [14, 21].

As such, recent group recommenders mostly rely on early aggregation methods [3, 27, 36, 42, 44] for learning group-level preference representations. In contrast to late aggregation, such methods first aggregate the preferences of all group members into a group-level representation, and then make group recommendations accordingly. Some early aggregation methods are built upon probabilistic models [12, 27, 44] by considering both a user’s own preferences and her/his impact to the group, but the same type of probability

distribution is assumed across users, which is infeasible in real-world cases. To address this problem, latent factor models [20, 31] are adopted to map groups to embeddings in the vector space. For example, attentive neural networks are proposed in [3, 14] to selectively aggregate user representations within a group, and [36] further captures the fine-grained interactions between group members via a sub-attention network. More recently, there are also studies to incorporate additional information like social connections [4] and knowledge graphs [11] into the learning process of group representations. However, as discussed in Section 1, the point embeddings used in those methods sacrifice the diversity of personal preferences. Besides, compared with alternative solutions like capsule networks [5, 25] that rank items by mapping the learned multi-interest representations to a single embedding, CubeRec recommends items directly with group hypercubes, thus retaining loyalty to the decision-making process in real groups.

5.2 Alleviating Data Sparsity for Groups

When recommending items to occasional groups, the other crucial goal is to mitigate the scarce interactions between groups and items which can heavily hinder the quality of learned group representations. Some research account for the social relationships among users to provide hints on group-level preferences [4, 14, 41], such as [41] that considers the impact of users’ social influences on a group’s decision. Afterwards, [14] proposes an intra-group voting mechanism based on self-attention [35] to simultaneously model the social influence and the pairwise user interactions, and [17] further learns multi-view embeddings from the interactions among groups, users and items with a heterogeneous graph. As self-supervised learning (SSL) has shown its effectiveness in general recommendation tasks [39, 43, 48], attempts are also made to design SSL objectives to counteract the data sparsity in group recommendation. For example, [45] leverages common users among groups and designs a double-scale contrastive learning to enhance the learned user representations. Meanwhile, [32] proposes a user-group mutual information maximization scheme to jointly learn informative user and group representations. Though SSL is proven to benefit group recommendation, existing methods are only compatible to point user/group embeddings in the vector space. In contrast, our design of SSL in CubeRec takes advantage of the geometric properties of hypercubes, allowing for more expressive group representations to be learned despite data sparsity.

6 CONCLUSION

To learn high-quality group representations that have the capacity to encode the multi-faceted user preferences and the flexibility to fully resemble a group’s decision-making process, we move beyond the traditional point embeddings, and propose to learn hypercube group representations in this paper. The proposed solution CubeRec takes advantage of its more expressive hypercubes in the vector space, and comes with a new distance metric and an intersection-based self-supervision paradigm to respectively facilitate group-item pairwise ranking and mitigation of data sparsity. Experimental results have rigorously verified the efficacy of CubeRec, shedding light on its applications in a wider range of recommendation tasks.

ACKNOWLEDGEMENT

This work is supported by UQ New Staff Research Start-up Grant (NS-2103). This work is partly sponsored by Australian Research Council under the streams of Future Fellowship (FT210100624), Discovery Project (DP190101985) and Discover Early Career Researcher Award (DE200101465), and is also supported by National Natural Science Foundation of China (U21A20470).

REFERENCES

- [1] Sihem Amer-Yahia, Senjuti Roy, Ashish Chawlat, Gautam Das, and Cong Yu. 2009. Group recommendation: Semantics and efficiency. *PVLDB* 2, 1 (2009), 754–765.
- [2] Linas Baltrunas, Tadas Makcinskis, and Francesco Ricci. 2010. Group recommendations with rank aggregation and collaborative filtering. In *RecSys*. 119–126.
- [3] Da Cao, Xiangnan He, Lianhai Miao, Yahui An, Chao Yang, and Richang Hong. 2018. Attentive group recommendation. In *SIGIR*. 645–654.
- [4] Da Cao, Xiangnan He, Lianhai Miao, Guangyi Xiao, Hao Chen, and Jiao Xu. 2021. Social-Enhanced Attentive Group Recommendation. *TKDE* 33, 03 (2021), 1195–1209.
- [5] Yukuo Cen, Jianwei Zhang, Xu Zou, Chang Zhou, Hongxia Yang, and Jie Tang. 2020. Controllable multi-interest framework for recommendation. In *SIGKDD*. 2942–2951.
- [6] Hongxu Chen, Hongzhi Yin, Weiqing Wang, Hao Wang, Quoc Viet Hung Nguyen, and Xue Li. 2018. PME: Projected metric embedding on heterogeneous networks for link prediction. In *SIGKDD*.
- [7] Tong Chen, Hongzhi Yin, Hongxu Chen, Rui Yan, Quoc Viet Hung Nguyen, and Xue Li. 2019. AIR: Attentional intention-aware recommender systems. In *ICDE*. 304–315.
- [8] Tong Chen, Hongzhi Yin, Quoc Viet Hung Nguyen, Wen-Chih Peng, Xue Li, and Xiaofang Zhou. 2020. Sequence-Aware Factorization Machines for Temporal Predictive Analytics. *ICDE* (2020).
- [9] Tong Chen, Hongzhi Yin, Guanhua Ye, Zi Huang, Yang Wang, and Meng Wang. 2020. Try this instead: Personalized and interpretable substitute recommendation. In *SIGIR*. 891–900.
- [10] Tong Chen, Hongzhi Yin, Yujia Zheng, Zi Huang, Yang Wang, and Meng Wang. 2021. Learning elastic embeddings for customizing on-device recommenders. In *SIGKDD*. 138–147.
- [11] Zhiyi Deng, Changyu Li, Shujin Liu, Waqar Ali, and Jie Shao. 2021. Knowledge-Aware Group Representation Learning for Group Recommendation. In *ICDE*. 1571–1582.
- [12] Jagadeesh Gorla, Neal Lathia, Stephen Robertson, and Jun Wang. 2013. Probabilistic group recommendation via information matching. In *WWW*. 495–504.
- [13] Lei Guo, Hongzhi Yin, Tong Chen, Xiangliang Zhang, and Kai Zheng. 2021. Hierarchical Hyperedge Embedding-based Representation Learning for Group Recommendation. *TOIS* (2021).
- [14] Lei Guo, Hongzhi Yin, Qinyong Wang, Bin Cui, Zi Huang, and Lizhen Cui. 2020. Group recommendation with latent voting mechanism. In *ICDE*. 121–132.
- [15] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. *SIGIR* (2020).
- [16] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. 173–182.
- [17] Zhixiang He, Chi-Yin Chow, and Jia-Dong Zhang. 2020. GAME: Learning Graphical and Attentive Multi-view Embeddings for Occasional Group Recommendation. In *SIGIR*. 649–658.
- [18] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. 2017. Collaborative metric learning. In *WWW*. 193–201.
- [19] Liang Hu, Jian Cao, Guandong Xu, Longbing Cao, Zhiping Gu, and Wei Cao. 2014. Deep modeling of group preferences for group-based recommendation. In *AAAI*.
- [20] Zhenhua Huang, Xin Xu, Honghao Zhu, and MengChu Zhou. 2020. An efficient group recommendation model with multiattention-based neural networks. *TNNLS* 31, 11 (2020), 4461–4474.
- [21] Renqi Jia, Xiaofei Zhou, Linhua Dong, and Shirui Pan. 2021. Hypergraph Convolutional Network for Group Recommendation. In *ICDM*. 260–269.
- [22] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2015).
- [23] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. *ICLR* (2017).
- [24] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* (2009).
- [25] Chao Li et al. 2019. Multi-interest network with dynamic routing for recommendation at Tmall. In *CIKM*. 2615–2623.
- [26] Yunchuan Li, Yan Zhao, and Kai Zheng. 2021. Preference-aware Group Task Assignment in Spatial Crowdsourcing: A Mutual Information-based Approach. In *ICDM*. 350–359.
- [27] Xingjie Liu, Yuan Tian, Mao Ye, and Wang-Chien Lee. 2012. Exploring personal impact for group recommendation. In *CIKM*. 674–683.
- [28] Quoc Viet Hung Nguyen, Chi Thang Duong, Thanh Tam Nguyen, Matthias Weidlich, Karl Aberer, Hongzhi Yin, and Xiaofang Zhou. 2017. Argument discovery via crowdsourcing. *VLDBJ* 26, 4 (2017), 511–535.
- [29] Elisa Quintarelli, Emanuele Rabosio, and Letizia Tanca. 2016. Recommending new items to ephemeral groups using contextual user influence. In *RecSys*. 285–292.
- [30] Hongyu Ren, Weihua Hu, and Jure Leskovec. 2020. Query2box: Reasoning over knowledge graphs in vector space using box embeddings. *ICLR* (2020).
- [31] Sarina Sajjadi Ghaemmaghami and Amiral Salehi-Abari. 2021. DeepGroup: Group Recommendation with Implicit Feedback. In *CIKM*. 3408–3412.
- [32] Aravind Sankar, Yanhong Wu, Yuhang Wu, Wei Zhang, Hao Yang, and Hari Sundaram. 2020. Groupim: A mutual information maximization framework for neural group recommendation. In *SIGIR*. 1279–1288.
- [33] Shunichi Seko, Takashi Yagi, Manabu Motegi, and Shinyo Muto. 2011. Group recommendation using feature space representing behavioral tendency and power balance among members. In *RecSys*. 101–108.
- [34] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *JMLR* (2014), 1929–1958.
- [35] Ashish Vaswani et al. 2017. Attention is all you need. In *NIPS*. 5998–6008.
- [36] Lucas Vinh Tran, Tuan-Anh Nguyen Pham, Yi Tay, Yiding Liu, Gao Cong, and Xiaoli Li. 2019. Interact and decide: Medley of sub-attention networks for effective group recommendation. In *SIGIR*. 255–264.
- [37] Wen Wang, Wei Zhang, Jun Rao, Zhijie Qiu, Bo Zhang, Leyu Lin, and Hongyuan Zha. 2020. Group-Aware Long-and Short-Term Graph Representation Learning for Sequential Group Recommendation. In *SIGIR*. 1449–1458.
- [38] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR*. 165–174.
- [39] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised graph learning for recommendation. In *SIGIR*. 726–735.
- [40] Hongzhi Yin and Bin Cui. 2016. *Spatio-temporal recommendation in social media*. Springer.
- [41] Hongzhi Yin, Qinyong Wang, Kai Zheng, Zhixu Li, Jiali Yang, and Xiaofang Zhou. 2019. Social influence-based group representation learning for group recommendation. In *ICDE*. 566–577.
- [42] Hongzhi Yin, Qinyong Wang, Kai Zheng, Zhixu Li, and Xiaofang Zhou. 2020. Overcoming data sparsity in group recommendation. *TKDE* (2020).
- [43] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Jundong Li, and Zi Huang. 2022. Self-Supervised Learning for Recommender Systems: A Survey. *arXiv preprint arXiv:2203.15876* (2022).
- [44] Quan Yuan, Gao Cong, and Chin-Yew Lin. 2014. COM: a generative model for group recommendation. In *SIGKDD*. 163–172.
- [45] Junwei Zhang, Min Gao, Junliang Yu, Lei Guo, Jundong Li, and Hongzhi Yin. 2021. Double-Scale Self-Supervised Hypergraph Learning for Group Recommendation. In *CIKM*. 2557–2567.
- [46] Shuai Zhang, Huoyu Liu, Aston Zhang, Yue Hu, Ce Zhang, Yumeng Li, Tanhao Zhu, Shaojian He, and Wenwu Ou. 2021. Learning User Representations with Hypercuboids for Recommender Systems. In *WSDM*. 716–724.
- [47] Shijie Zhang, Hongzhi Yin, Tong Chen, Zi Huang, Lizhen Cui, and Xiangliang Zhang. 2021. Graph embedding for recommendation against attribute inference attacks. In *The Web Conference*. 3002–3014.
- [48] Yujia Zhou, Zhicheng Dou, Yutao Zhu, and Ji-Rong Wen. 2021. PSSL: Self-supervised Learning for Personalized Search with Contrastive Sampling. In *CIKM*. 2749–2758.