

LGMRec: Local and Global Graph Learning for Multimodal Recommendation

Zhiqiang Guo¹, Jianjun Li^{1*}, Guohui Li^{2*}, Chaoyang Wang³, Si Shi⁴, Bin Ruan¹

¹ School of Computer Science and Technology, Huazhong University of Science and Technology, Wuhan, China

² School of Software Engineering, Huazhong University of Science and Technology, Wuhan, China

³ Wuhan Digital Engineering Institute, Wuhan, China

⁴ Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ), Shenzhen, China

{zhiqiangguo, jianjunli, guohuili}@hust.edu.cn, sunwardtree@outlook.com, shisi@gml.ac.cn, binruan0227@gmail.com

Abstract

The multimodal recommendation has gradually become the infrastructure of online media platforms, enabling them to provide personalized service to users through a joint modeling of user historical behaviors (e.g., purchases, clicks) and item various modalities (e.g., visual and textual). The majority of existing studies typically focus on utilizing modal features or modal-related graph structure to learn user local interests. Nevertheless, these approaches encounter two limitations: (1) Shared updates of user ID embeddings result in the consequential coupling between collaboration and multimodal signals; (2) Lack of exploration into robust global user interests to alleviate the sparse interaction problems faced by local interest modeling. To address these issues, we propose a novel Local and Global Graph Learning-guided Multimodal Recommender (LGMRec), which jointly models local and global user interests. Specifically, we present a local graph embedding module to independently learn collaborative-related and modality-related embeddings of users and items with local topological relations. Moreover, a global hypergraph embedding module is designed to capture global user and item embeddings by modeling insightful global dependency relations. The global embeddings acquired within the hypergraph embedding space can then be combined with two decoupled local embeddings to improve the accuracy and robustness of recommendations. Extensive experiments conducted on three benchmark datasets demonstrate the superiority of our LGMRec over various state-of-the-art recommendation baselines, showcasing its effectiveness in modeling both local and global user interests.

Introduction

With the explosive growth of massive multimedia information (e.g., images, texts, and videos) on online media platforms, such as YouTube and Tiktok, a lot of efforts have been devoted to multimodal recommender systems (MRSs) to assist these platforms in providing personalized services to users. Nowadays, the primary task of MRSs is to design an effective way to integrate item multimodal information into traditional user-item interaction modeling frameworks to capture comprehensive user interests.

Some early studies on MRSs adopt either the linear fusion between item modal features and their ID embeddings (He

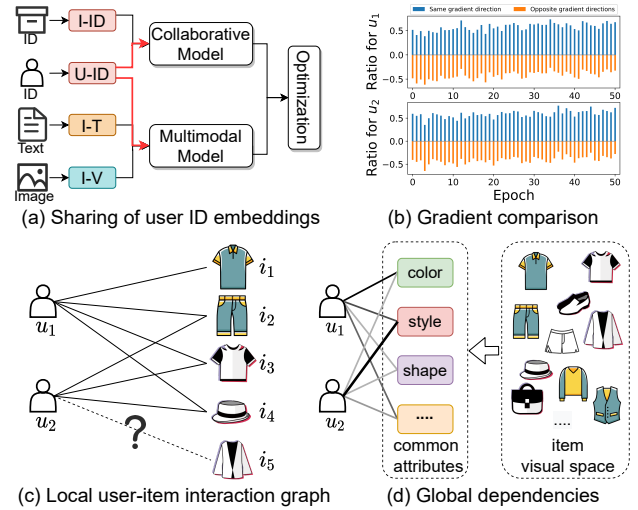


Figure 1: Illustrations of (a) sharing of user ID embeddings, (b) the gradient comparison of user ID embeddings updated from different models during training, (c) local user-item interaction graph, and (d) global dependencies between users and attributions. Darker lines indicate greater user interest.

and McAuley 2016; Liu, Wu, and Wang 2017; Wei et al. 2021) or the attention mechanism on item modalities (Chen et al. 2017, 2019; Liu et al. 2019) to model representations of users and items. However, The efficacy of these models is somewhat constrained as they only model low-order user-item interactions. The surge of research on graph-based recommendations (Wang et al. 2019; He et al. 2020; Mao et al. 2021; Wu et al. 2021) has sparked a wave of explorations in using graph neural networks (GNN) to enhance multimodal recommendations. These works typically capture higher-order user interests from the user-item graph that integrates multimodal contents (Wei et al. 2019, 2020; Wang et al. 2021; Yi et al. 2022; Tao et al. 2022; Wei et al. 2023), or construct modality-aware auxiliary graph structures to transfer multimodal knowledge into item and user embeddings (Zhang et al. 2021, 2022a; Zhou et al. 2023).

Though achieving remarkable progress, existing studies on MRSs still suffer from the following two limitations in modeling user interests. (1) *Coupling*. **Firstly**, collabora-

*Corresponding author.

tion and multimodal information provide different avenues for exploring user interests. In general, collaborative signals emphasize similar user behavior patterns, while modal knowledge is reflected through content similarity. However, prior works (Wei et al. 2019; Yi et al. 2022) often overlook this matter and share user ID embeddings in both collaborative and multimodal modeling modules (red line in Figure 1 (a)) to learn user interests that couple collaborative and multimodal signals. Experimentally, we randomly select two users from the Baby dataset and exhibit the gradient comparison of their ID embeddings (with 64 dimensions) from the collaborative and multimodal modeling modules in Figure 1 (b). In the early stages of training, the ratio of gradients with opposite directions (orange bar) from the two modules in all dimensions exceeds 50% for each user, which demonstrates that collaborative and multimodal signals generally have different guidance for user embedding learning¹. Though this ratio slightly decreases as the training continues, the coupling design still restricts stable updates of user embeddings. (2) *Locality*. **Secondly**, most existing methods (Tao et al. 2022; Zhou et al. 2023) only learn local user interests from the interaction graph (Figure 1 (c)), lacking the exploration of user global interests. Sparse user-item interactions limit their modeling of robust user interests. As shown in Figure 1 (d), user global (general) interests are usually related to item attribute labels that do not rely on the local interactions. Specifically, items usually have multiple common attributions from visual space, such as *color*, *style*, *shape*. Users have different interests in various attributes. For example, u_1 may like clothes with bright colors, while u_2 prefers a simple style. A method that modeling only local interests may recommend the shirt i_1 to u_2 based on similar behaviors, i.e., same purchases (i_2, i_3, i_4) between u_1 and u_2 . But, the global interests of u_2 can provide additional guidance, making it more likely to recommend the outerwear i_5 with simple style that match u_2 's true interests.

To address the aforementioned issues, we propose a novel Local and Global Graph Learning-guided Multimodal Recommender (LGMRec), which explores capturing and exploiting both local and global representations of users and items to facilitate multimodal recommendation. Specifically, to address the first limitation, we present the local graph embedding module to independently capture collaborative-related and modality-related local user interests by performing message propagation on user-item interaction graphs with ID embeddings and modal features, respectively. In view of the many-to-many dependency relationship between attributes and items is similar to that between hyperedges and nodes in hypergraphs, **we further consider each implicit attribute as a hyperedge, and present a global hypergraph embedding module to model hypergraph structure dependencies, so as to address the second limitation**. Extensive experimental results on three real-world datasets demonstrate that LGMRec surpasses various recommendation baselines significantly, and verify its effectiveness and robustness in

¹In fact, approximately 94.26% of users in Baby dataset present such a situation, that is, more than 50% of the embedding dimensions have opposite gradient directions during the training process.

modeling local and global user interests.

Related work

Graph-based Recommendation The powerful ability of graph neural networks (Kipf and Welling 2016; Hu et al. 2019) in modeling high-order connectivity has greatly promoted the development of recommender systems. Specifically, graph-based recommendation methods model user and item representations by naturally converting the user history interactions into a user-item bipartite graph. Early studies directly inherit the message propagation mechanism of vanilla graph neural network to aggregate high-order neighbor information to represent users and items (Berg, Kipf, and Welling 2017; Ying et al. 2018; Wang et al. 2019). Later, by simplifying the message propagation process, some graph-based recommendation methods further improve recommendation performance (Chen et al. 2020; He et al. 2020; Mao et al. 2021). Additionally, some other methods explore more node dependencies to enhance the representations of users and items (Ma et al. 2019; Sun et al. 2019, 2020a; Li et al. 2022). Later, contrastive learning is also adopted to enhance graph-based recommendations (Lee et al. 2021; Wu et al. 2021; Yu et al. 2022; Lin et al. 2022; Yang et al. 2021; Cai et al. 2023) to construct contrastive views. However, since no modality features are considered, their modeling abilities are limited by sparse interactions.

Hypergraph learning for Recommendation By constructing the hyperedge structure containing more than two nodes, hypergraph learning (Feng et al. 2019; Gao et al. 2020) can enhance the generalization ability of the model via capturing complex node dependencies. Some recommendation methods (Ji et al. 2020; Wang et al. 2020; He et al. 2021; Yu et al. 2021; Xia et al. 2021; Zhang et al. 2022b) try to build hypergraph structures and node-hyperedge connections to capture high-order interaction patterns and achieve substantial performance improvements. To further improve performance, several recently developed methods (Xia et al. 2022; Xia, Huang, and Zhang 2022) combine self-supervised learning and hypergraph learning to model robust user and item representations. For example, HCCF (Xia et al. 2022) enhances collaborative filtering with the hypergraph-guided self-supervised learning paradigm. Different from these works that generate hypergraph dependencies via only collaborative embeddings, our work achieves hypergraph structure learning with the modeling of modality-aware global relations.

Multi-modal Recommendation The multi-modal recommendation has become the basic application on online media platforms to provide personalized services to users by analyzing the massive multi-modal information (e.g., images and textual descriptions) and user historical behaviors (e.g., reviews, clicks). Early studies on MRSs usually incorporate multi-modal contents as side information to extend the vanilla CF framework (He and McAuley 2016; Chen, He, and Kan 2016; Gao, Zhang, and Xu 2017; Du et al. 2020) or utilize deep autoencoder to model modal features (Guo et al. 2022; Liu et al. 2022). Inspired by the great success of graph-based recommendation methods (He et al. 2020;

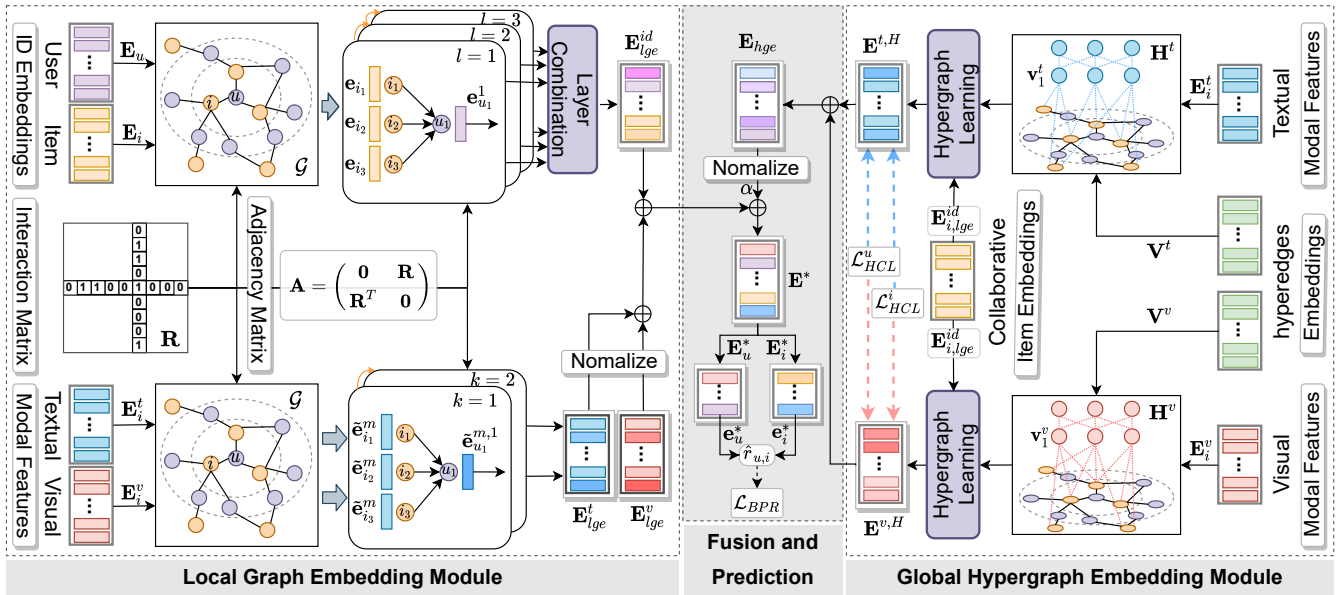


Figure 2: The framework of the proposed LGMRec with visual and textual modalities of items (i.e., $m \in \{v, t\}$).

Ma et al. 2019; Mao et al. 2021), some studies directly model user high-order interests on modality-specific interaction graphs (Wei et al. 2019, 2020; Sun et al. 2020b; Du et al. 2022; Kim et al. 2022). For instance, MMGCN (Wei et al. 2019) incorporates modality information into the graph message passing to infer modality-related user preferences. Another line utilizes auxiliary semantic graph structures learned from multimodal features to enhance user or item representations (Wang et al. 2021; Zhang et al. 2021). For example, LATTICE (Zhang et al. 2021) is a representative method that exploits modal content similarity to generate auxiliary latent item semantic relations to promote recommendation. Recently, Some works (Wei et al. 2021; Yi et al. 2022; Tao et al. 2022; Zhang et al. 2022a; Zhou et al. 2023; Wei et al. 2023) introduce contrastive learning into MRSs to model robust user and item representations. However, these methods usually perform message passing along the edges of user-item interactions to obtain local user interests, failing to explore modality-aware comprehensive user interests.

Methodology

In this section, we first formulate the problem of multimodal recommendation and present the overall framework of our LGMRec, and then introduce each component in detail.

Problem Statement and Overview

We set the user set as $\mathcal{U} = \{u\}$ and item set as $\mathcal{I} = \{i\}$. The ID embeddings of each user $u \in \mathcal{U}$ and item $i \in \mathcal{I}$ are denoted as $\mathbf{e}_u \in \mathbb{R}^d$ and $\mathbf{e}_i \in \mathbb{R}^d$, respectively, where d is the embedding dimension. The user-item interactions can be represented as a matrix $\mathbf{R} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$, in which the element $r_{u,i} = 1$ if user u interacts with item i , and $r_{u,i} = 0$ otherwise. Based on interaction matrix \mathbf{R} , we can construct the user-item interaction graph $\mathcal{G} = \{\mathcal{U} \cup \mathcal{I}, \mathcal{E}\}$, where \mathcal{E} is

edge set build on observed interactions, i.e., a nonzero $r_{u,i}$ corresponds to an edge between user u and item i on the graph \mathcal{G} . Further, we incorporate item multimodal contents and denote the original modality feature of item i generated from pre-trained models as $\mathbf{e}_i^m \in \mathbb{R}^{d_m}$ under modality $m \in \mathcal{M}$, where \mathcal{M} is the set of modalities and d_m denotes the dimension of modal features. In this work, we consider two mainstream modalities, vision v and text t , i.e., $\mathcal{M} = \{v, t\}$. Given the above settings, the multimodal recommendation aims to learn a prediction function to forecast the score $\hat{r}_{u,i}$ of an item i adopted by a user u via joint modeling user behaviors and multimodal contents. Formally,

$$\hat{r}_{u,i} = \text{PREDICTION}(\mathbf{R}, \mathbf{E}^{id}, \{\mathbf{E}_i^m\}_{m \in \mathcal{M}}) \quad (1)$$

where $\text{PREDICTION}(\cdot)$ is the prediction function, $\mathbf{E}^{id} = [\mathbf{e}_{u_1}, \dots, \mathbf{e}_{u_{|\mathcal{U}|}}, \mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_{|\mathcal{I}|}}] \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{I}|) \times d}$ denotes the ID embedding matrix by stacking all the ID embeddings of users and items, $\mathbf{E}_i^m = [\mathbf{e}_{i_1}^m, \dots, \mathbf{e}_{i_{|\mathcal{I}|}}^m] \in \mathbb{R}^{|\mathcal{I}| \times d_m}$ is the item modal feature matrix under modality m .

Overview. As illustrated in Figure 2, the framework of LGMRec consists of three major components: (i) Local graph embedding (LGE) module, which adopts GNN to capture collaborative-related and modality-related user local interests on user-item interaction graph with ID embeddings and modal features, respectively; (ii) Global hypergraph embedding (GHE) module, which learns the global user and item representations by capturing the global hypergraph structure dependencies from different item modal feature spaces; and (iii) Fusion and prediction module, which fuses both local and global embeddings to predict final user preference scores for items.

Local Graph Embedding (LGE) Module

The LGE module is designed to independently learn the collaborative-related and modality-related user and item rep-

representations with local topology structure for avoiding unstable updates of user embeddings and promoting decoupled user interest learning.

Collaborative Graph Embedding (CGE) We first capture the high-order connectivity via the message propagation on the user-item interaction graph with ID embeddings. In particular, the collaborative graph propagation function $\text{CGPROG}(\cdot)$ in the $(l+1)$ -th layer can be formatted as,

$$\mathbf{E}^{l+1} = \text{CGPROG}(\mathbf{E}^l) = \left(\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \right) \mathbf{E}^l, \quad (2)$$

where $\text{CGPROG}(\cdot)$ function inherits the lightweight form of the simplified graph convolutional network (Chen et al. 2020; He et al. 2020), $\mathbf{A} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{I}|) \times (|\mathcal{U}|+|\mathcal{I}|)}$ is the adjacency matrix constructed from interaction matrix \mathbf{R} , and \mathbf{D} is the diagonal matrix of \mathbf{A} . Each diagonal element $\mathbf{D}_{j,j}$ in \mathbf{D} denotes the number of nonzero entries in the j -th row vector of matrix \mathbf{A} . The initial embeddings matrix is set as $\mathbf{E}^0 = \mathbf{E}^{id}$. Then, we adopt the layer combination (He et al. 2020) to integrate all embeddings from hidden layers,

$$\mathbf{E}_{lge}^{id} = \text{LAYERCOMB}(\mathbf{E}^0, \mathbf{E}^1, \mathbf{E}^2, \dots, \mathbf{E}^L), \quad (3)$$

where $\mathbf{E}_{lge}^{id} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{I}|) \times d}$ is collaborative-related embeddings of users and items with local neighborhood information. We use the mean function to achieve $\text{LAYERCOMB}(\cdot)$ for embedding integration.

Modality Graph Embedding (MGE) Considering the semantic differences between modalities, we further independently infer the modality-related embeddings of users and items on the interaction graphs with modal features. **The original modal features of items are usually generated from different pre-trained models, e.g., ResNet (He et al. 2016), BERT (Kenton and Toutanova 2019),** they have different dimensions in different feature spaces. We require the projection of high-dimensional modal feature \mathbf{e}_i^m of each item into a unified embedding space \mathbb{R}^d as,

$$\tilde{\mathbf{e}}_i^m = \text{TRANSFORM}(\mathbf{e}_i^m) = \mathbf{e}_i^m \cdot \mathbf{W}_m, \quad (4)$$

where $\tilde{\mathbf{e}}_i^m$ is item i 's transformed modal feature, $\text{TRANSFORM}(\cdot)$ is a projection function parameterized by a transformation matrix $\mathbf{W}_m \in \mathbb{R}^{d_m \times d}$. Due to the difficulty in obtaining user modal information, existing methods often reuse user ID embedding as input for modality-specific graphs, resulting in coupling of collaborative and modal signals. Different from them, we initialize the user modal features by aggregating item modal features,

$$\tilde{\mathbf{e}}_u^m = \frac{1}{|\mathcal{N}_u|} \sum_{i \in \mathcal{N}_u} \tilde{\mathbf{e}}_i^m, \quad (5)$$

where \mathcal{N}_u denotes the neighbor set of user $u \in \mathcal{U}$ on user-item interaction graph \mathcal{G} . This operation ensures the separate updates of ID embedding and modal features. Thereafter, we can construct the modal feature matrix $\tilde{\mathbf{E}}^m = [\tilde{\mathbf{e}}_{u_1}^m, \dots, \tilde{\mathbf{e}}_{u_{|\mathcal{U}|}}^m, \tilde{\mathbf{e}}_{i_1}^m, \dots, \tilde{\mathbf{e}}_{i_{|\mathcal{I}|}}^m] \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{I}|) \times d}$ as initial input $\tilde{\mathbf{E}}^{m,0}$ to learn modality-related embeddings via implementing a light graph propagation function $\text{MGPROG}(\cdot)$,

$$\tilde{\mathbf{E}}^{m,k+1} = \text{MGPROG}(\tilde{\mathbf{E}}^{m,k}) = \left(\mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}} \right) \tilde{\mathbf{E}}^{m,k}. \quad (6)$$

Here, we choose high-order modal embeddings $\tilde{\mathbf{E}}^{m,K}$ in the last K -th layer as the modality-related embeddings (i.e., $\mathbf{E}_{lge}^m = \tilde{\mathbf{E}}^{m,K}$) with local modal information.

Global Hypergraph Embedding (GHE) Module

The GHE module is designed to capture the modality-aware global representations of users and items against sparse and noisy user behaviors.

Hypergraph Dependency Constructing Explicit attribute information of item modalities is often unavailable, especially for visual modalities. Hence, we define learnable implicit attribute vectors $\{\mathbf{v}_a^m\}_{a=1}^A$ ($\mathbf{v}_a^m \in \mathbb{R}^{d_m}$) as hyperedge embeddings under modality m to adaptively learn the dependencies between implicit attributes and items/users, where A is the number of hyperedges. Specifically, We obtain hypergraph dependency matrices in low-dimensional embedding space by,

$$\mathbf{H}_i^m = \mathbf{E}_i^m \cdot \mathbf{V}^{m\top}, \quad \mathbf{H}_u^m = \mathbf{A}_u \cdot \mathbf{H}_i^{m\top}, \quad (7)$$

where $\mathbf{H}_i^m \in \mathbb{R}^{|\mathcal{I}| \times A}$ and $\mathbf{H}_u^m \in \mathbb{R}^{|\mathcal{U}| \times A}$ are the item-hyperedge and user-hyperedge dependency matrices, respectively. \mathbf{E}_i^m is the raw item modal feature matrix, $\mathbf{V}^m = [\mathbf{v}_1^m, \dots, \mathbf{v}_A^m] \in \mathbb{R}^{A \times d_m}$ is the hyperedge vector matrix, and $\mathbf{A}_u \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$ is the user-related adjacency matrix extracted from \mathbf{A} . **Intuitively, items with similar modal features are more likely to be connected to the same hyperedge. The user-hyperedge dependencies are indirectly derived through the user-item interactions, which implies the user behavior intention, i.e., the more frequently users interact with items under a certain attribute, the more they may prefer the attribute.**

To further avoid the negative impact of meaningless relationships, we employ the Gumbel-Softmax reparameterization (Jang, Gu, and Poole 2017) to ensure that an item is attached to only one hyperedge as much as possible,

$$\tilde{\mathbf{h}}_{i,*}^m = \text{SOFTMAX} \left(\frac{\log \delta - \log(1 - \delta) + \mathbf{h}_{i,*}^m}{\tau} \right), \quad (8)$$

where $\mathbf{h}_{i,*}^m \in \mathbb{R}^A$ is the i -th row vector of \mathbf{H}_i^m that reflects the relations between item i and all hyperedges. $\delta \in \mathbb{R}^A$ is a noise vector, where each value $\delta_j \sim \text{Uniform}(0, 1)$, and τ is a temperature hyperparameter. Afterwards, we can get the augmented item-attribute hypergraph dependency matrix $\tilde{\tilde{\mathbf{H}}}_i^m$. By performing similar operations on \mathbf{H}_u^m , we can obtain the augmented user-attribute relation matrix $\tilde{\tilde{\mathbf{H}}}_u^m$.

Hypergraph Message Passing By taking the attribute hyperedge as an intermediate hub, we achieve hypergraph message passing to deliver global information to users and items without being limited by hop distances. Formally,

$$\mathbf{E}_i^{m,h+1} = \text{DROP}(\tilde{\tilde{\mathbf{H}}}_i^m) \cdot \text{DROP}(\tilde{\tilde{\mathbf{H}}}_i^{m\top}) \cdot \mathbf{E}_i^{m,h}, \quad (9)$$

where $\mathbf{E}_i^{m,h}$ is the global embedding matrix of items in the h -th hypergraph layer, and $\text{DROP}(\cdot)$ denotes a dropout function. We take collaborative embedding matrix $\mathbf{E}_{i,lge}^{id}$ of items

as the initial global embedding matrix when $h = 0$. Further, we can calculate the global user embedding matrix as,

$$\mathbf{E}_u^{m,h+1} = \text{DROP}(\tilde{\mathbf{H}}_u^m) \cdot \text{DROP}(\tilde{\mathbf{H}}_i^{m\top}) \cdot \mathbf{E}_i^{m,h}. \quad (10)$$

Apparently, the hypergraph passing explicitly enables global information transfer by taking the item collaborative embedding and modality-aware hypergraph dependencies as input. Then, we can obtain the global embeddings matrix \mathbf{E}_{ghe} by aggregating global embeddings from all modalities,

$$\mathbf{E}_{ghe} = \sum_{m \in \mathcal{M}} \mathbf{E}^{m,H}, \quad \mathbf{E}^{m,H} = [\mathbf{E}_u^{m,H}, \mathbf{E}_i^{m,H}], \quad (11)$$

where $\mathbf{E}_u^{m,H} \in \mathbb{R}^{|\mathcal{U}| \times d}$ and $\mathbf{E}_i^{m,H} \in \mathbb{R}^{|\mathcal{I}| \times d}$ are global embedding matrices of user u and item i obtained in the H -th hypergraph layer under modality m , respectively.

To further achieve the robust fusion of global embeddings among different modalities, we develop cross-modal hypergraph contrastive learning to distill the self-supervision signals for global interest consistency. Specifically, we take the global embeddings of users acquired in different modalities as positive pairs and different users as negative pairs, and then employ the InfoNCE (Gutmann and Hyvärinen 2010) to formally define user-side hypergraph contrastive loss as,

$$\mathcal{L}_{\text{HCL}}^u = \sum_{u \in \mathcal{U}} -\log \frac{\exp(s(\mathbf{E}_u^{v,H}, \mathbf{E}_u^{t,H})/\tau)}{\sum_{u' \in \mathcal{U}} \exp(s(\mathbf{E}_u^{v,H}, \mathbf{E}_{u'}^{t,H})/\tau)}, \quad (12)$$

where $s(\cdot)$ is the cosine function, and τ is the temperature factor, generally set to 0.2. Note here we only consider visual and textual modalities, i.e., $m \in \{v, t\}$. Similarly, we can define item-side cross-modal contrastive loss $\mathcal{L}_{\text{HCL}}^i$.

Fusion and Prediction

We acquire the final representations \mathbf{E}^* of users and items by fusing their two types of local embeddings \mathbf{E}_{lge}^{id} , \mathbf{E}_{lge}^m and global embeddings \mathbf{E}_{ghe} ,

$$\mathbf{E}^* = \mathbf{E}_{lge}^{id} + \sum_{m \in \mathcal{M}} \text{NORM}(\mathbf{E}_{lge}^m) + \alpha \cdot \text{NORM}(\mathbf{E}_{ghe}), \quad (13)$$

where $\text{NORM}(\cdot)$ is a normalization function to alleviate the value scale difference among embeddings, α is an adjustable factor to control the integration of global embeddings.

We then use inner product to calculate the preference score $\hat{r}_{u,i}$ of user u towards item i , i.e., $\hat{r}_{u,i} = \mathbf{e}_u^* \cdot \mathbf{e}_i^{*\top}$. The Bayesian personalized ranking (BPR) loss (Rendle et al. 2012) is employed to optimize model parameters,

$$\mathcal{L}_{\text{BPR}} = - \sum_{(u,i^+,i^-) \in \mathcal{R}} \ln \sigma(\hat{r}_{u,i^+} - \hat{r}_{u,i^-}) + \lambda_1 \|\Theta\|_2^2, \quad (14)$$

where $\mathcal{R} = \{(u, i^+, i^-) | (u, i^+) \in \mathcal{G}, (u, i^-) \notin \mathcal{G}\}$ is a set of triples for training, $\sigma(\cdot)$ is the sigmoid function, and λ_1 and Θ represent the regularization coefficient and model parameters, respectively.

Finally, we integrate hypergraph contrastive loss with the BPR (Rendle et al. 2012) loss into a unified objective as,

$$\mathcal{L} = \mathcal{L}_{\text{BPR}} + \lambda_2 \cdot (\mathcal{L}_{\text{HCL}}^u + \mathcal{L}_{\text{HCL}}^i) \quad (15)$$

where λ_2 is a hyperparameter for loss term weighting. We minimize the joint objective \mathcal{L} by using Adam optimizer (Kingma and Ba 2014). The weight-decay regularization term is applied over model parameters Θ .

Dataset	#User	#Item	#Interaction	Sparsity
Baby	19,445	7,050	160,792	99.883%
Sports	35,598	18,357	296,337	99.955%
Clothing	39,387	23,033	278,677	99.969%

Table 1: Statistics of the three evaluation datasets

Experiment

Experimental Settings

Datasets To evaluate our proposed model, we conduct comprehensive experiments on three widely used Amazon datasets (McAuley et al. 2015): Baby, Sports and Outdoors, Clothing Shoes and Jewelry. We refer to them as **Baby**, **Sports**, **Clothing** for brief. We adopt the 5-core setting to filter users and items for each dataset. The three datasets include both visual and textual modal features. In this work, we use the 4096-dimensional original visual features and 384-dimensional original textual features that have been extracted and published in prior work (Zhou et al. 2023). The statistics of the three datasets are summarized in Table 1.

Evaluation Protocols For each dataset, we randomly split historical interactions into training, validation, and testing sets with 8 : 1 : 1 ratio. Two widely used protocols are used to evaluate the performance of top- n recommendation: Recall ($\text{R}@n$) and Normalized Discounted Cumulative Gain (He et al. 2015) ($\text{N}@n$). We tune n in $\{10, 20\}$ and report the average results for all users in the testing set.

Parameter Settings For a fair comparison, we optimize all models with the default batch size 2048, learning rate 0.001, and embedding size $d = 64$. For all graph-based methods, the number L of collaborative graph proration layers is set to 2. In addition, we initialize the model parameters with the Xavier method (Glorot and Bengio 2010). For our model, the optimal hyper-parameters are determined via grid search on the validation set. Specifically, the number of modal graph embedding layers and hypergraph embedding layers (K and H) are tuned in $\{1, 2, 3, 4\}$. The number A of hyperedge is searched in $\{1, 2, 4, 8, 16, 32, 64, 128, 256\}$. The dropout ratio ρ and the adjust factor α are tuned in $\{0.1, 0.2, \dots, 1.0\}$. We search both the adjust weight λ_2 of contrastive loss and the regularization coefficient λ_1 in $\{1e^{-6}, 1e^{-5}, \dots, 0.1\}$. The early stop mechanism is adopted, i.e., the training will stop when $\text{R}@20$ on the verification set does not increase for 20 successive epochs. We implement LGMRec² with MMRec (Zhou 2023).

Baselines We compare our proposed LGMRec with the following four groups of recommendation baselines, including (1) General CF Models: **BPR** (Rendle et al. 2012); (2) Graph-based Recommenders: **LightGCN** (He et al. 2020), **SGL** (Wu et al. 2021), **NCL** (Lin et al. 2022); (3) Hypergraph-based Recommenders: **HCCF** (Xia et al. 2022), **SHT** (Xia, Huang, and Zhang 2022); and (4) Multi-Modal Recommenders: **VBPR** (He and McAuley 2016), **MMGCN** (Wei et al. 2019), **GRCN** (Wei et al. 2020),

²<https://github.com/georgeguo-cn/LGMRec>

Datasets	Baby				Sports				Clothing			
	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20
BPR	0.0379	0.0607	0.0202	0.0261	0.0452	0.0690	0.0252	0.0314	0.0211	0.0315	0.0118	0.0144
LightGCN	0.0464	0.0732	0.0251	0.0320	0.0553	0.0829	0.0307	0.0379	0.0331	0.0514	0.0181	0.0227
SGL	0.0532	0.0820	0.0289	0.0363	0.0620	0.0945	0.0339	0.0423	0.0392	0.0586	0.0216	0.0266
NCL	0.0538	0.0836	0.0292	0.0369	0.0616	0.0940	0.0339	0.0421	0.0410	0.0607	0.0228	0.0275
HCCF	0.0480	0.0756	0.0259	0.0332	0.0573	0.0857	0.0317	0.0394	0.0342	0.0533	0.0187	0.0235
SHT	0.0470	0.0748	0.0256	0.0329	0.0564	0.0838	0.0306	0.0384	0.0345	0.0541	0.0192	0.0243
VBPR	0.0424	0.0663	0.0223	0.0284	0.0556	0.0854	0.0301	0.0378	0.0281	0.0412	0.0158	0.0191
MMGCN	0.0498	0.0749	0.0261	0.0315	0.0582	0.0825	0.0305	0.0382	0.0329	0.0564	0.0219	0.0253
GRCN	0.0531	0.0835	0.0291	0.0370	0.0600	0.0921	0.0324	0.0407	0.0431	0.0664	0.0230	0.0289
LATTICE	0.0536	0.0858	0.0287	0.0370	0.0618	0.0950	0.0337	0.0423	0.0459	0.0702	0.0253	0.0306
MMGCL	0.0522	0.0778	0.0289	0.0355	0.0660	0.0994	0.0362	0.0448	0.0438	0.0669	0.0239	0.0297
MICRO	0.0570	0.0905	0.0310	0.0406	0.0675	0.1026	0.0365	0.0463	0.0496	0.0743	0.0264	0.0332
SLMRec	0.0540	0.0810	0.0296	0.0361	0.0676	0.1007	0.0374	0.0462	0.0452	0.0675	0.0247	0.0303
BM3	0.0538	0.0857	0.0301	0.0378	0.0659	0.0979	0.0354	0.0437	0.0450	0.0669	0.0243	0.0295
LGMRec	0.0644*	0.1002*	0.0349*	0.0440*	0.0720*	0.1068*	0.0390*	0.0480*	0.0555*	0.0828*	0.0302*	0.0371*
Improv.	12.98%	10.72%	12.58%	8.37%	6.51%	4.09%	4.28%	3.67%	11.90%	11.44%	14.39%	1.75%

Table 2: Overall performances of LGMRec and other baselines on three datasets. The best result is in boldface and the second best is underlined. The t-tests validate the significance of performance improvements with p -value ≤ 0.05 .

LATTICE (Zhang et al. 2021), MMGCL (Yi et al. 2022), MICRO (Zhang et al. 2022a) SLMRec (Tao et al. 2022), BM3 (Zhou et al. 2023).

Performance Comparison

The performance comparison for all methods on the three datasets is summarized in Table 2, from which we have the following key observations: (1) **The superiority of LGM-Rec.** LGMRec substantially outperforms all other baselines and achieves promising performance across different datasets. We attribute such significant improvements to: i) The modeling of separated local embeddings that excavates user decoupled interests; ii) The hypergraph learning injects the modality-related global dependencies to local graph embeddings to mitigate interactive sparsity. (2) **The effectiveness of modal features.** Introducing knowledge-rich modality information is beneficial for boosting performance. Experimentally, though only linearly fusing the ID embeddings and modal features of items, the performance of VBPR still outperforms its counterpart (i.e., BPR). By effectively modeling the modal information, the multimodal recommenders (e.g., MMGCN, LATTICE, SLMRec, BM3) with LightGCN as the backbone network basically achieve better results than LightGCN. (3) **The effectiveness of hypergraph learning.** Hypergraph-based recommenders (i.e., HCCF and SHT) outperform the graph-based CF model LightGCN, suggesting the effectiveness of modeling global dependencies under hypergraph architecture. Besides, the significant improvement of LGMRec over competitive baselines further demonstrates the potential of hypergraph networks in modeling modality-aware global dependencies.

Ablation Study

We conduct ablation studies to explore the compositional effects of LGMRec. From the results reported in Table 3, we

Components	Baby		Sports		Clothing	
	R@20	N@20	R@20	N@20	R@20	N@20
<i>w/o MM</i>	0.0732	0.0320	0.0829	0.0379	0.0514	0.0227
<i>w/o LGE</i>	0.0806	0.0351	0.0851	0.0392	0.0741	0.0327
<i>w/o CGE</i>	0.0947	0.0423	0.0997	0.0448	0.0807	0.0360
<i>w/o MGE</i>	0.0929	0.0417	0.0988	0.0440	0.0804	0.0357
<i>w/o GHE</i>	0.0972	0.0430	0.1032	0.0468	0.0803	0.0364
<i>w/o HCL</i>	0.0992	0.0434	0.1051	0.0474	0.0812	0.0368
<i>w/ SUID</i>	0.0869	0.0379	0.0895	0.0395	0.0713	0.0307
LGMRec	0.1002	0.0440	0.1068	0.0480	0.0828	0.0371

Table 3: Ablation of different components on LGMRec.

can find: (1) The variant *w/o MM* without multimodal contents degenerates into LightGCN and achieves the worst performance, indicating that introducing modality features can greatly improve accuracy. (2) Removing either *LGE* or *GHE* can cause performance drops of LGMRec, demonstrating the benefits of modeling both local and global user interests. Notably, the variant *w/o LGE* performs worse than *w/o GHE*, which indicates that local interests directly related to user behavior are more important, and global interests can serve as a supplement. (3) In local graph embeddings, the variant *w/o CGE* (with *MGE* only) achieves better performance than *w/o MGE* (with *CGE* only) on all datasets, which reveals the importance of integrating multimodal features into user-item interaction modeling. (4) The variant *w/o HCL* removes hypergraph contrastive learning and only linearly adds all global embeddings. Its performances indicate that contrastive fusion of global embeddings of different modalities can improve performance by modeling the inter-modal global semantic consistency. (5) The variant *w/ SUID* that still shares user ID embeddings in both *MGE* and *CGE* modules performs worse than LGMRec, verifying the benefits of




Figure 3: Performance w.r.t. different user interaction sparsity degrees in terms of R@20 on Baby and Sports datasets.




Figure 4: Performances under different settings of two key hyperparameters (A and α) on Clothing datasets.

independently modeling user decoupled interests.

In-Depth Analysis

Performance with Different Data Sparsity We further study the influence of sparse user interactions by comparing LGMRec with five representative multimodal recommendation baselines: MMGCN, LATTICE, MMGCL, SLMRec, and BM3, on Baby and Sports datasets. Multiple user groups are constructed according to the number of their interactions. For example, the first user group contains users interacting with 0–5 items. From the results in Figure 3, we can observe that: (1) The superior performance of LGMRec is consistent across user groups with different sparsity degrees, revealing the effectiveness of LGMRec in alleviating interaction sparsity by modeling local and global representations. (2) LGMRec achieves more performance gains on sparser user groups. Specifically, LGMRec realizes 19.95% and 10.83% improvements over the best baseline for the sparsest and densest group on Baby, respectively, verifying the robustness of LGMRec in dealing with sparser user interactions.

Hyperparameter Analysis Figure 4 reports the impact of two key hyperparameters of LGMRec on Clothing dataset:

Hyperedge number A . From the left figure in Figure 4, we can observe that LGMRec presents performance promotion as the number of hyperedges increases, demonstrating the effectiveness of capturing multi-hyperedge global structures, especially for sparser Clothing datasets.

Adjustable weight α . Impact of weight α of fusing global embeddings is also investigated in Figure 4. We can find that the performance first rises to an optimal value ($\alpha = 0.2$) and then declines, which suggests that an appropriate α can improve accuracy by properly supplementing global embeddings, but a too large α may negatively affect performance.

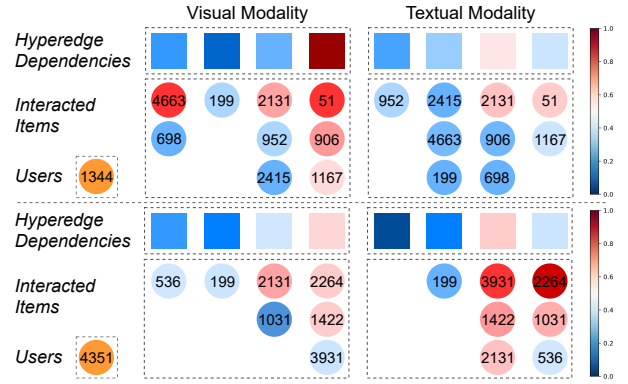


Figure 5: Case study of learned global dependencies of two users u_{1344} and u_{4351} with four hyperedges on Baby dataset.

Case Study

We qualitatively study the global hypergraph dependencies. Specifically, we randomly select two users u_{1344} , u_{4351} with similar global embeddings learned on Baby dataset. Hypergraph dependencies under visual and textual modalities for the two users and the items they interact with are presented in Figure 5. The four hyperedges (squares) are shaded depending on the user-hyperedge dependency score. Moreover, the interacted items (circles) are arranged below the corresponding hyperedges in order, according to the maximum item-hyperedge dependency score. From Figure 5, we can observe that: (1) The user-hyperedge dependencies differ in different modalities. For example, the global interests of user u_{1344} in the visual modality are mainly related to the 4-th attribute hyperedge. Under the textual modality, user u_{1344} has larger dependency scores with the 3-rd hyperedges. Thus, we guess that the four items (i_{51} , i_{906} , i_{1167} , and i_{2131}) closely related to head hyperedges can reflect user u_{1344} 's true preferences, while item i_{4663} attached to the 1-st hyperedge may be a noise interaction. (2) Although the interacted items are largely non-overlapping, user u_{4351} and user u_{1344} still have similar hyperedge dependencies, demonstrating why their global embeddings are similar. The results further reveal that LGMRec can exploit global hypergraph learning to distill similar knowledge of item modal features for performance improvement.

Conclusion

In this work, we proposed a novel model LGMRec for MRSs, which captures and utilizes local embeddings with local topological information and global embeddings with hypergraph dependencies. Specifically, we adopted a local graph embedding module to independently learn collaborative-related and modality-related local user interests. A global hypergraph embedding module is further designed to mine global user interests. Extensive experiments on three datasets demonstrated the superiority of our model over various baselines. For future work, we intend to seek better means of modeling the differences and commonalities among modalities for further performance improvement.

Acknowledgements

We would like to thank all anonymous reviewers for their valuable comments. The work was partially supported by the National Key R&D Program of China under Grant No. 2022YFC3802101 and the National Natural Science Foundation of China under Grant No. 62272176.

References

- Berg, R. v. d.; Kipf, T. N.; and Welling, M. 2017. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*.
- Cai, X.; Huang, C.; Xia, L.; and Ren, X. 2023. LightGCL: Simple Yet Effective Graph Contrastive Learning for Recommendation. *arXiv preprint arXiv:2302.08191*.
- Chen, J.; Zhang, H.; He, X.; Nie, L.; Liu, W.; and Chua, T.-S. 2017. Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention. In *Proceedings of SIGIR*, 335–344.
- Chen, L.; Wu, L.; Hong, R.; Zhang, K.; and Wang, M. 2020. Revisiting graph based collaborative filtering: A linear residual graph convolutional network approach. In *Proceedings of AAAI*, volume 34, 27–34.
- Chen, T.; He, X.; and Kan, M.-Y. 2016. Context-aware image tweet modelling and recommendation. In *Proceedings of ACM MM*, 1018–1027.
- Chen, X.; Chen, H.; Xu, H.; Zhang, Y.; Cao, Y.; Qin, Z.; and Zha, H. 2019. Personalized fashion recommendation with visual explanations based on multimodal attention network: Towards visually explainable recommendation. In *Proceedings of SIGIR*, 765–774.
- Du, X.; Wang, X.; He, X.; Li, Z.; Tang, J.; and Chua, T.-S. 2020. How to learn item representation for cold-start multimedia recommendation? In *Proceedings of ACM MM*, 3469–3477.
- Du, X.; Wu, Z.; Feng, F.; He, X.; and Tang, J. 2022. Invariant Representation Learning for Multimedia Recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 619–628.
- Feng, Y.; You, H.; Zhang, Z.; Ji, R.; and Gao, Y. 2019. Hypergraph neural networks. In *Proceedings of AAAI*, volume 33, 3558–3565.
- Gao, J.; Zhang, T.; and Xu, C. 2017. A unified personalized video recommendation via dynamic recurrent neural networks. In *Proceedings of ACM MM*, 127–135.
- Gao, Y.; Zhang, Z.; Lin, H.; Zhao, X.; Du, S.; and Zou, C. 2020. Hypergraph learning: Methods and practices. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5): 2548–2566.
- Glorot, X.; and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of AISTATS*, 249–256.
- Guo, Z.; Li, G.; Li, J.; and Chen, H. 2022. TopicVAE: Topic-aware Disentanglement Representation Learning for Enhanced Recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 511–520.
- Gutmann, M.; and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of AISTATS*, 297–304.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of CVPR*, 770–778.
- He, L.; Chen, H.; Wang, D.; Jameel, S.; Yu, P.; and Xu, G. 2021. Click-through rate prediction with multi-modal hypergraphs. In *Proceedings of CIKM*, 690–699.
- He, R.; and McAuley, J. 2016. VBPR: visual bayesian personalized ranking from implicit feedback. In *Proceedings of AAAI*, volume 30.
- He, X.; Chen, T.; Kan, M.-Y.; and Chen, X. 2015. Trirank: Review-aware explainable recommendation by modeling aspects. In *Proceedings of CIKM*, 1661–1670.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of SIGIR*, 639–648.
- Hu, F.; Zhu, Y.; Wu, S.; Wang, L.; and Tan, T. 2019. Hierarchical graph convolutional networks for semi-supervised node classification. *arXiv preprint arXiv:1902.06667*.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *Proceedings of ICLR*.
- Ji, S.; Feng, Y.; Ji, R.; Zhao, X.; Tang, W.; and Gao, Y. 2020. Dual channel hypergraph collaborative filtering. In *Proceedings of SIGKDD*, 2020–2029.
- Kenton, J. D. M.-W. C.; and Toutanova, L. K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*, 4171–4186.
- Kim, T.; Lee, Y.-C.; Shin, K.; and Kim, S.-W. 2022. MARIO: Modality-Aware Attention and Modality-Preserving Decoders for Multimedia Recommendation. In *Proceedings of CIKM*, 993–1002.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Lee, D.; Kang, S.; Ju, H.; Park, C.; and Yu, H. 2021. Bootstrapping user and item representations for one-class collaborative filtering. In *Proceedings of SIGIR*, 317–326.
- Li, G.; Guo, Z.; Li, J.; and Wang, C. 2022. MDGCF: Multi-Dependency Graph Collaborative Filtering with Neighborhood-and Homogeneous-level Dependencies. In *Proceedings of CIKM*, 1094–1103.
- Lin, Z.; Tian, C.; Hou, Y.; and Zhao, W. X. 2022. Improving graph collaborative filtering with neighborhood-enriched contrastive learning. In *Proceedings of WWW*, 2320–2329.
- Liu, F.; Cheng, Z.; Sun, C.; Wang, Y.; Nie, L.; and Kankanhalli, M. 2019. User diverse preference modeling by multimodal attentive metric learning. In *Proceedings of the 27th ACM International Conference on Multimedia*, 1526–1534.

- Liu, Q.; Wu, S.; and Wang, L. 2017. Deepstyle: Learning user preferences for visual recommendation. In *Proceedings of SIGIR*, 841–844.
- Liu, X.; Tao, Z.; Shao, J.; Yang, L.; and Huang, X. 2022. EliMRec: Eliminating Single-modal Bias in Multimedia Recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 687–695.
- Ma, J.; Cui, P.; Kuang, K.; Wang, X.; and Zhu, W. 2019. Disentangled graph convolutional networks. In *Proceedings of ICML*, 4212–4221.
- Mao, K.; Zhu, J.; Xiao, X.; Lu, B.; Wang, Z.; and He, X. 2021. UltraGCN: ultra simplification of graph convolutional networks for recommendation. In *Proceedings of CIKM*, 1253–1262.
- McAuley, J.; Targett, C.; Shi, Q.; and Van Den Hengel, A. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of SIGIR*, 43–52.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*.
- Sun, J.; Zhang, Y.; Guo, W.; Guo, H.; Tang, R.; He, X.; Ma, C.; and Coates, M. 2020a. Neighbor interaction aware graph convolution networks for recommendation. In *Proceedings of SIGIR*, 1289–1298.
- Sun, J.; Zhang, Y.; Ma, C.; Coates, M.; Guo, H.; Tang, R.; and He, X. 2019. Multi-graph convolution collaborative filtering. In *Proceedings of ICDM*, 1306–1311.
- Sun, R.; Cao, X.; Zhao, Y.; Wan, J.; Zhou, K.; Zhang, F.; Wang, Z.; and Zheng, K. 2020b. Multi-modal knowledge graphs for recommender systems. In *Proceedings of CIKM*, 1405–1414.
- Tao, Z.; Liu, X.; Xia, Y.; Wang, X.; Yang, L.; Huang, X.; and Chua, T.-S. 2022. Self-supervised learning for multimedia recommendation. *IEEE Transactions on Multimedia*.
- Wang, J.; Ding, K.; Hong, L.; Liu, H.; and Caverlee, J. 2020. Next-item recommendation with sequential hypergraphs. In *Proceedings of SIGIR*, 1101–1110.
- Wang, Q.; Wei, Y.; Yin, J.; Wu, J.; Song, X.; and Nie, L. 2021. Dualgnn: Dual graph neural network for multimedia recommendation. *IEEE Transactions on Multimedia*.
- Wang, X.; He, X.; Wang, M.; Feng, F.; and Chua, T.-S. 2019. Neural graph collaborative filtering. In *Proceedings of SIGIR*, 165–174.
- Wei, W.; Huang, C.; Xia, L.; and Zhang, C. 2023. Multi-Modal Self-Supervised Learning for Recommendation. In *Proceedings of WWW*.
- Wei, Y.; Wang, X.; Li, Q.; Nie, L.; Li, Y.; Li, X.; and Chua, T.-S. 2021. Contrastive learning for cold-start recommendation. In *Proceedings of ACM MM*, 5382–5390.
- Wei, Y.; Wang, X.; Nie, L.; He, X.; and Chua, T.-S. 2020. Graph-refined convolutional network for multimedia recommendation with implicit feedback. In *Proceedings of the 28th ACM International Conference on Multimedia*, 3541–3549.
- Wei, Y.; Wang, X.; Nie, L.; He, X.; Hong, R.; and Chua, T.-S. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In *Proceedings of the 27th ACM International Conference on Multimedia*, 1437–1445.
- Wu, J.; Wang, X.; Feng, F.; He, X.; Chen, L.; Lian, J.; and Xie, X. 2021. Self-supervised graph learning for recommendation. In *Proceedings of SIGIR*, 726–735.
- Xia, L.; Huang, C.; Xu, Y.; Zhao, J.; Yin, D.; and Huang, J. 2022. Hypergraph contrastive collaborative filtering. In *Proceedings of SIGIR*, 70–79.
- Xia, L.; Huang, C.; and Zhang, C. 2022. Self-supervised hypergraph transformer for recommender systems. In *Proceedings of SIGKDD*, 2100–2109.
- Xia, X.; Yin, H.; Yu, J.; Wang, Q.; Cui, L.; and Zhang, X. 2021. Self-supervised hypergraph convolutional networks for session-based recommendation. In *Proceedings of AAAI*, volume 35, 4503–4511.
- Yang, Y.; Wu, L.; Hong, R.; Zhang, K.; and Wang, M. 2021. Enhanced graph learning for collaborative filtering via mutual information maximization. In *Proceedings of SIGIR*, 71–80.
- Yi, Z.; Wang, X.; Ounis, I.; and Macdonald, C. 2022. Multi-modal graph contrastive learning for micro-video recommendation. In *Proceedings of SIGIR*, 1807–1811.
- Ying, R.; He, R.; Chen, K.; Eksombatchai, P.; Hamilton, W. L.; and Leskovec, J. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of SIGKDD*, 974–983.
- Yu, J.; Yin, H.; Li, J.; Wang, Q.; Hung, N. Q. V.; and Zhang, X. 2021. Self-supervised multi-channel hypergraph convolutional network for social recommendation. In *Proceedings of WWW*, 413–424.
- Yu, J.; Yin, H.; Xia, X.; Chen, T.; Cui, L.; and Nguyen, Q. V. H. 2022. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *Proceedings of SIGIR*, 1294–1303.
- Zhang, J.; Zhu, Y.; Liu, Q.; Wu, S.; Wang, S.; and Wang, L. 2021. Mining latent structures for multimedia recommendation. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3872–3880.
- Zhang, J.; Zhu, Y.; Liu, Q.; Zhang, M.; Wu, S.; and Wang, L. 2022a. Latent Structure Mining with Contrastive Modality Fusion for Multimedia Recommendation. *IEEE Transactions on Knowledge and Data Engineering*.
- Zhang, X.; Xu, B.; Yang, L.; Li, C.; Ma, F.; Liu, H.; and Lin, H. 2022b. Price does matter! modeling price and interest preferences in session-based recommendation. In *Proceedings of SIGIR*, 1684–1693.
- Zhou, X. 2023. MMRec: Simplifying Multimodal Recommendation. *arXiv preprint arXiv:2302.03497*.
- Zhou, X.; Zhou, H.; Liu, Y.; Zeng, Z.; Miao, C.; Wang, P.; You, Y.; and Jiang, F. 2023. Bootstrap Latent Representations for Multi-Modal Recommendation. In *Proceedings of WWW*, 845–854.

Appendix

Complexity Analysis of LGMRec

We conduct the complexity analysis of LGMRec. The computational cost of LGMRec mainly comes from two parts. In the local graph embedding module, the collaborative graph embedding has a $\mathcal{O}(L \times |\mathcal{E}| \times d)$ complexity, where L is the number of graph message passing layers, $|\mathcal{E}|$ is the number of edges in user-item interaction graph \mathcal{G} . For modality graph embedding, the computational cost of modal feature initialization is $\mathcal{O}(|\mathcal{M}| \times d_m \times d)$, where $|\mathcal{M}|$ is the number of modalities. The modality graph propagation has the same computational complexity as the collaborative graph embedding. Thus, the overall time complexity of the local graph embedding module is $\mathcal{O}(((L + K) \times |\mathcal{E}| + |\mathcal{M}| \times d_m) \times d)$.

For the global hypergraph embedding module, the time complexity of hypergraph dependency constructing is $\mathcal{O}(|\mathcal{M}| \times A \times |\mathcal{I}| \times (|\mathcal{U}| + d_m))$, where A is the number of hyperedges. The hypergraph message passing schema takes $\mathcal{O}(|\mathcal{M}| \times (|\mathcal{I}| \times H + |\mathcal{U}|) \times A \times d)$ complexity with the global information propagation, where H is the number of hypergraph layers. The cost of the hypergraph contrastive learning is $\mathcal{O}(B \times (|\mathcal{U}| + |\mathcal{I}|) \times d)$ with only two modalities v and t , where B is the batch size. The overall time complexity of the global hypergraph embedding module is $\mathcal{O}(|\mathcal{M}| \times A \times (|\mathcal{I}| \times ((|\mathcal{U}| + d_m) + H \times d)) + |\mathcal{U}|)$. In practice, our two modules can be executed in parallel, which makes LGMGN quite efficient in actual execution. During the training process, its actual running time is comparable to existing methods, such as MMGCL (Yi et al. 2022), and SLMRec (Tao et al. 2022). Compared to existing graph-based multimodal recommenders, LGMRec only involves $\mathcal{O}(|\mathcal{M}| \times A \times d_m)$ extra parameters for the memory cost.

Baselines

(i) General CF Models

- **BPR** (Rendle et al. 2012) maps the user and item in a low-dimensional latent embedding space and utilizes Bayesian pairwise ranking loss to optimize model parameters.

(ii) Graph-based Recommendations

- **LightGCN** (He et al. 2020) is a typical graph-based CF method that utilizes light graph convolutional networks to learn high-order connectivity of users and items.
- **SGL** (Wu et al. 2021) introduces contrastive learning to enhance graph collaborative filtering. We implement this method by data augmentation with random edge dropout.
- **NCL** (Lin et al. 2022) enhances the graph-based CF model by identifying structural and semantic neighboring nodes as positive samples to construct contrastive views.

(ii) Hypergraph-based Recommendations

- **HCCF** (Xia et al. 2022) leverages the hypergraph neural network to inject the global collaborative relations into the graph-based recommendation.
- **SHT** (Xia, Huang, and Zhang 2022) captures the global collaborative embeddings for contrastive learning by joint utilizing hypergraph encoder and multi-head attention mechanism.

Parameters	L	K	H	A	α	ρ	λ_1	λ_2
Baby	2	2	1	4	0.3	0.5	$1e^{-6}$	$1e^{-4}$
Sports	4	2	1	4	0.6	0.4	$1e^{-6}$	$1e^{-4}$
Clothing	3	2	2	64	0.2	0.2	$1e^{-6}$	$1e^{-4}$

Table 4: Parameter setting for three datasets.

(iii) Multi-Modal Recommendations

- **VBPR** (He and McAuley 2016) integrates modal features with ID embeddings to extend the traditional CF paradigm.
- **MMGCN** (Wei et al. 2019) learns fine-grained modality-specific user preferences by achieving message-passing on the user-item bipartite graph of each modality.
- **GRCN** (Wei et al. 2020) is a structure-refined graph multimedia recommender, in which modality contents are used to adjust the structure of interaction graph by identifying the noisy edges.
- **LATTICE** (Zhang et al. 2021) exploits multi-modal features to mine the latent semantic structure between items to improve multi-modal recommendation.
- **MMGCL** (Yi et al. 2022) includes contrastive learning into multimodal recommendation via graph augmentation with modality-related edge dropout and masking.
- **MICRO** (Zhang et al. 2022a) extends the LATTICE (Zhang et al. 2021) to fuse multimodal features by introducing contrastive learning to capture modality-shared and modality-specific information.
- **SLMRec** (Tao et al. 2022) devises three types of data augmentation at different granularity to achieve multi-modal self-supervised tasks.
- **BM3** (Zhou et al. 2023) utilizes a simple latent embedding dropout mechanism to generate contrastive view in self-supervised learning for multimodal recommendation.

Parameter Setting

For a fair comparison, we optimize all models with the default batch size 2048, learning rate 0.001, and embedding size $d = 64$. Table 4 presents other optimal parameter settings for the three datasets. In addition, all experiments in this paper are performed in the same experimental environment with Intel(R) Xeon(R) Silver 4210R CPU @ 2.40GHz and GeForce RTX 3090.

Hyperparameter Analysis

We explore the impacts of all hyperparameters for performance and report them in Figure 6.

- **Collaborative graph layers L** . From Figure 6 (a), the results of the single-layer model are slightly inferior to that of the multi-layer model. The outcomes suggest that combining ID embeddings with sufficient multi-layer local structure information can obtain higher-quality user and item representations. In addition, the sparser datasets may require a deeper and larger local structure receptive field to facilitate recommendation, e.g., $L = 4$ on Sports, and $L = 3$ on Clothing.




Figure 6: Performances under different settings of all hyperparameters of LGMRec on Baby, Sports, and Clothing datasets.

- **Modality graph layers K .** The results in Figure 6 (b) demonstrate that the two-layer model achieves better performance and the increasing of layers does not bring a performance improvement, which indicates that the discrimination of nodes is decreasing as the layer number increases. The reason may be that aggregating deeper neighbors may lead to knowledge redundancy of node modality features.
- **Hypergraph layers H .** The impact of hypergraph layer H is shown in Figure 6 (c). From the results, we can see that shallow global embedding performs better than multi-layer, possibly because multi-layer hypergraph propagation can lead to the excessive smoothness of node representations and reduce performance. In practice, we can take $H = 1$ for all datasets.
- **Hyperedge number A .** Figure 6 (d) shows the performance of LGMRec with different settings of hyperedge number A . As mentioned in our experiments, the performance promotes as the number of hyperedges increases on the sparser Clothing dataset. For Baby and Sports datasets, the performance usually reaches its optimum at $A = 4$. The results demonstrate the effectiveness of capturing multi-hyperedge global structures, especially for sparser datasets.
- **Adjustable factor α .** The performance of LGMRec with different settings of weight α is reported in Figure 6 (e).

The results on the three datasets show a consistent trend, that is, the performance first increases to the optimal value and then decreases. The results suggest that properly supplementing global embeddings is suitable for modeling robust user interests. So, we can set $\alpha = 0.3, 0.6, 0.2$ for Baby, Sports, and Clothing datasets, respectively.

- **Drop ratio ρ .** We tune the dropout ratio ρ from $\{0.1, 0.2, \dots, 1.0\}$ to control the retention of hypergraph structure dependencies. The results in Figure 6 (f) indicate that proper dropout (compared to no dropout, i.e., $\rho = 1.0$) is suitable to suppress the global noise and improve the robustness of the representation.
- **Coefficient λ_2 .** Coefficient λ_2 determines the influence of hypergraph contrastive loss, and the performance of LGMRec under different λ_2 is shown in Figure 6 (g). Similar to adjustable factor α , the performance also first improves to reach optimal and then declines as λ_2 increases. The results illustrate that an appropriate λ_2 can mitigate value scale differences between HCL loss and BPR loss. In practice, we can uniformly set $\lambda_2 = 1e^{-4}$ on the three datasets.
- **Regularization coefficient λ_1 .** We perform a grid search for parameter λ_1 to verify the effect of regularization. From Figure 6 (h), we can see that the effect of different coefficient λ_1 is negligible on the three datasets. Therefore, a small $\lambda_1 = 1e^{-6}$ is desirable.