

COHESION: Composite Graph Convolutional Network with Dual-Stage Fusion for Multimodal Recommendation

Jinfeng Xu

The University of Hong Kong
HongKong SAR, China
jinfeng@connect.hku.hk

Zheyu Chen

The Hong Kong Polytechnic
University
HongKong SAR, China
zheyu.chen@connect.polyu.hk

Wei Wang

Shenzhen MSU-BIT University
Shenzhen, China
ehomewang@ieee.org

Xiping Hu

Beijing Institute of Technology
Beijing, China
huxp@bit.edu.cn

Sang-Wook Kim

Hanyang University
South Korea
wook@hanyang.ac.kr

Edith C. H. Ngai*

The University of Hong Kong
HongKong SAR, China
chngai@eee.hku.hk

Abstract

Recent works in multimodal recommendations, which leverage diverse modal information to address data sparsity and enhance recommendation accuracy, have garnered considerable interest. Two key processes in multimodal recommendations are modality fusion and representation learning. Previous approaches in modality fusion often employ simplistic attentive or pre-defined strategies at early or late stages, failing to effectively handle irrelevant information among modalities. In representation learning, prior research has constructed heterogeneous and homogeneous graph structures encapsulating user-item, user-user, and item-item relationships to better capture user interests and item profiles. Modality fusion and representation learning were considered as two independent processes in previous work. This paper reveals that these two processes are complementary and can support each other. Specifically, powerful representation learning enhances modality fusion, while effective fusion improves representation quality. Stemming from these two processes, we introduce a **CO**mposite graph **H**convolutional **nE**twork with dual-stage **fUSION** for the multimodal recommendation, named COHESION. Specifically, it introduces a dual-stage fusion strategy to reduce the impact of irrelevant information, refining all modalities using behavior modality in the early stage and fusing their representations at the late stage. It also proposes a composite graph convolutional network that utilizes user-item, user-user, and item-item graphs to extract heterogeneous and homogeneous latent relationships within users and items. Besides, it introduces a novel adaptive optimization to ensure balanced and reasonable representations across modalities. Extensive experiments

on three public datasets demonstrate the significant superiority of COHESION over various competitive baselines.

CCS Concepts

• **Information systems** → **Recommender systems**;

Keywords

Multimodal, Recommender System, Dual-Stage Fusion

ACM Reference Format:

Jinfeng Xu, Zheyu Chen, Wei Wang, Xiping Hu, Sang-Wook Kim, and Edith C. H. Ngai. 2025. COHESION: Composite Graph Convolutional Network with Dual-Stage Fusion for Multimodal Recommendation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3726302.3729927>

1 Introduction

The rapid development of the Internet has led to information explosion, making recommender systems an indispensable tool in human society. Traditional recommender systems rely on modeling user preferences through historical user-item interactions [3, 14, 27, 29]. However, the data sparsity problem always frustrates the accuracy of recommendations. As the types of information on social media become increasingly diverse, multimodal information has recently been used to alleviate the data sparsity problem [15, 17, 30]. A line of work [2, 8] integrates multimodal information as side information to enhance item representations. To better learn the representation, many works [23, 25, 28] construct the user-item bipartite graph and utilize the graph convolution network (GCN) to enhance the representation learning in multimodal recommendations. Despite some previous works achieving notable recommendation performance via constructing the heterogeneous user-item graph, the relationships between user-user and item-item remain unexplored. Recently, DualGNN [23] constructs the user-user graph to explore the hidden preference pattern by similar users. LATTICE [36] builds the item-item graph to capture semantically correlated signals. Both of these two homogeneous graphs improve the performance of recommendations. In this paper, we pose two comprehensive questions: **Would exploiting both user-user and item-item graphs achieve even better performance? And how?** To answer this question, we select the most standard

*Corresponding authors

This work was supported by the Hong Kong UGC General Research Fund no. 17203320 and 17209822, and the project grants from the HKU-SCF FinTech Academy.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '25, Padua, Italy

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1592-1/2025/07
<https://doi.org/10.1145/3726302.3729927>

GCN-based multimodal recommendation model MMGCN¹ [25] and carefully design its three variants²: MMGCN^u, MMGCNⁱ, and MMGCN^{ui}, which incorporate only user-user, only item-item, and both user-user and item-item graphs, respectively. To assess the impact of fusion strategies, we implement the early fusion and late fusion strategies for each MMGCN variant, namely MMGCN_e and MMGCN_l, respectively. Table 1 shows the recommendation performance of eight MMGCN variants on the Baby dataset. In Table 1, our findings include: 1) The late fusion strategy outperforms the early fusion strategy when integrating only user-user or item-item graphs. Conversely, the early fusion strategy is more effective than late fusion when both graph types are integrated, possibly because late fusion permits deeper modality-specific learning, but it is more susceptible to the adverse effects of irrelevant inter-modality information. 2) Among the early fusion strategies, MMGCN^{ui} shows performance superior to MMGCN^u and MMGCNⁱ, suggesting that combining both user and item graphs facilitates more effective representation learning. These observations suggest that suboptimal fusion strategies in prior studies may have constrained the representation learning capabilities of the otherwise well-designed models. This hypothesis is empirically supported by the results presented in Table 6.

To liberate the representation learning ability of composite graph convolutional network, we introduce a **COM**posite **gRAPH** convolutional **NE**twork with dual-stage **FUSION** for a multimodal recommendation, named COHESION. Specifically, we propose a novel dual-stage fusion strategy to mitigate the negative effects of irrelevant information among different modalities by refining all modalities using behavior modality in the early stage and fusing all modalities' representations in the late stage. We believe that refining all modalities using behavior modality can enable each modality to map to a better semantic space. This is because the goal of multimodal recommendation is to predict user preferences instead of understanding multimodal content. Moreover, we propose the composite graph convolutional network to learn the representations for all modalities, including homogeneous and heterogeneous graphs. More specifically, we utilize a user-item graph to mine the relations between users and items and utilize user-user and item-item graphs to extract latent relations within users and items, respectively. Thanks to our dual-stage fusion strategy, our composite graph convolutional network liberates outstanding representation learning ability. Finally, we fuse all modality representations in the late stage and propose an adaptive optimization, which can make the representations learned from different modalities balanced. Extensive experiments on three widely used datasets demonstrate the superiority of COHESION over various competitive baselines.

2 Related work

2.1 Multimodal Recommendation

Multimodal recommendation systems endeavor to harness multimodal information to enhance the accuracy of recommendations and to address the issue of data sparsity. Prior works [2, 8] alleviate

Table 1: Performance of MMGCN variants in Baby dataset.

Metrics	Recall@10	NDCG@10	Recall@20	NDCG@20
MMGCN _e	0.0351	0.0589	0.0192	0.0249
MMGCN _e ^u	0.0454	0.0710	0.0242	0.0311
MMGCN _e ⁱ	0.0592	0.0924	0.0303	0.0399
MMGCN _e ^{ui}	0.0630	0.0981	0.0330	0.0426
MMGCN _l	0.0378	0.0615	0.0200	0.0261
MMGCN _l ^u	0.0493	0.0747	0.0266	0.0287
MMGCN _l ⁱ	0.0627	0.0992	0.0330	0.0424
MMGCN _l ^{ui}	0.0601	0.0940	0.0313	0.0401

the data sparsity challenge by extracting and utilizing visual content to enrich the representation of items based on matrix factorization techniques. Tran et al. [21] extend this approach by incorporating textual content to further enrich item representations. More recently, several studies [23, 41] integrate information from multiple modalities concurrently. Drawing inspiration from traditional recommendation systems, Wei et al. [24, 25] enhance the quality of user-item interactions through the implementation of a bipartite graph structure. To more effectively extract and analyze relationships between users and between items, DualGNN [23] introduces a user-user homogeneous graph to uncover latent relationships among users. Similarly, FREEDOM [41] and LATTICE [36] incorporate item-item homogeneous graphs to augment modality-specific information. MONET [12] and MARIO [11] leverage modality-aware attention and tailored GCN to explore items' multimodal features. In the realm of exploring advanced structural possibilities, LGMRec [7] and DiffMM [10] investigate the potential of hyper-graph structures and diffusion models, respectively, to significantly enhance the effectiveness of multimodal recommendation systems. To better alleviate the data sparsity problem, numerous studies [20, 30, 42] integrate self-supervised learning tasks within the multimodal recommendation framework to discern hidden preference patterns across different modalities. However, it is important to note that self-supervised learning can also introduce irrelevant noisy information, which may include augmented self-supervised signals derived from user misclick behaviors or popularity biases, potentially reducing the model accuracy.

2.2 Multimodal Fusion

Multimodal information can effectively alleviate the data sparsity problem, and improve the accuracy of recommendations. However, multimodal information may lead to modal noise due to the feature of different modalities existing in different semantic spaces. Therefore, multimodal fusion plays a crucial role in noise mitigation. According to the timing of the modality fusion, most existing methods could be divided into two types: early fusion and late fusion. The early fusion methods [8, 36] fuse the modality features with ID embedding before constructing the heterogeneous graph. However, the modality-specific features will be lost during training. Differently, the late fusion methods [23, 25, 41] learn the representation of each modality, and fuse all the predicted ratings in the end.

With the rapid development of representation learning ability in recent multimodal recommendation models, predefined fusion strategies are more fragile to the negative impact of irrelevant information among modalities and are difficult to fuse different

¹MMGCN is the earliest and simplest graph-based multimodal recommendation model, which only contains GCNs without extra components.

²All of the user-user graph, item-item graph, early fusion strategy, and late fusion strategy in this investigation are designed following a recent survey paper [37].

modalities effectively. To this end, we propose a tailored dual-stage fusion strategy, which effectively liberates the representation learning ability for multimodal recommendation models.

3 Methodology

In this section, we present our proposed COHESION¹ in detail. The overall architecture is shown in Fig. 1.

3.1 Preliminary

Let a set of users $u \in \mathcal{U}$ and a set of items $i = \{i_m|i_t, i_v\} \in \mathcal{I}$, where $m \in \mathcal{M}$ is the modality, \mathcal{M} is the set of modalities, and i_t, i_v and i_v represent the behavior, textual, and visual modality features, respectively. $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a given graph with a node set \mathcal{V} and an edge set \mathcal{E} , where $|\mathcal{V}| = |\mathcal{U}| + |\mathcal{I}|$. The user-item interaction matrix is denoted as $\mathcal{R} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$. Specifically, each entry $\mathcal{R}_{u,i}$ indicates whether the user u is connected to item i , with a value of 1 representing a connection and 0 otherwise. The graph structure of \mathcal{G} can be denoted as the adjacency matrix $\mathcal{A} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{I}|) \times (|\mathcal{U}|+|\mathcal{I}|)}$:

$$\mathcal{A} = \begin{bmatrix} 0^{|\mathcal{U}| \times |\mathcal{U}|} & \mathcal{R} \\ \mathcal{R}^T & 0^{|\mathcal{I}| \times |\mathcal{I}|} \end{bmatrix}. \quad (1)$$

The symmetrically normalized matrix is $\tilde{\mathcal{A}} = \mathcal{D}^{-\frac{1}{2}} \mathcal{A} \mathcal{D}^{-\frac{1}{2}}$, where \mathcal{D} represents a diagonal degree matrix. It is worth noting that we only consider three types of modalities in this paper, including behavior, textual, and visual modalities. However, we note it can be extended to more modalities easily.

3.2 Early Fusion

Multimodal information can effectively alleviate the data-sparsity problem for traditional recommendations. However, it also inevitably contains irrelevant information among modalities [19, 33]. To this end, we propose an early fusion strategy, which refines all modalities by behavior modality before constructing heterogeneous graphs. Our early fusion strategy can effectively build a bridge to mitigate the gap between different modalities by behavior modality. In particular, we first transform the feature embeddings $E_{i_m} = \{E_{i_t}, E_{i_v}\}$ into the same dimension:

$$\dagger E_{i_m} = W_m^{\triangleright} \sigma(W_m^{\triangleleft} E_{i_m} + b_1) + b_2, \quad (2)$$

where $W_m^{\triangleleft} \in \mathbb{R}^{d_m \times 4d_i}$, $W_m^{\triangleright} \in \mathbb{R}^{4d_i \times d_i}$, b_1 , and b_2 are trainable parameters. d_m and d_i denote embedding dimensionality of modality m and behavior modality i , respectively. $\sigma(\cdot)$ denotes the Leaky-ReLU function [32]. Then we refine each modality embedding $\dagger E_{i_m}$ by behavior modality $\dagger E_{i_t}$ respectively:

$$\dagger E_{i_m} = \sqrt{0.5 \times ((\dagger E_{i_m})^2 + (\dagger E_{i_t})^2) + \epsilon}, \quad (3)$$

where $\dagger E_{i_m}$ denotes refined modality embedding for modality m , and ϵ is a positive infinitesimal to prevent zero vector. Through this tailored yet simple early fusion strategy, we enhance the representation learning capability of multimodal recommendation models by refining different modalities via the behavior modality. This early fusion strategy mitigates the gap between modalities, leading to more accurate and robust recommendations.

¹The code will be released at: <https://github.com/Jinfeng-Xu/COHESION> after the SIGIR'25 conference.

3.3 Heterogeneous Graph

To better learn the modality-specific user and item representations, we construct a user-item graph for each modality, which is denoted as $\mathcal{G}_m = \{\mathcal{G}_i, \mathcal{G}_t, \mathcal{G}_v\}$. Following previous works [25, 41], we maintain the same graph structure \mathcal{G} for different \mathcal{G}_m , but only retain the node features associated with a specific modality m . To alleviate the over-smoothing problem, we design a residual-based GCN for each user-item graph. The message propagation stage at the l -th graph convolution layer can be formulated as:

$$\dagger E_m^{(l)} = (\text{Sim}(\tilde{\mathcal{A}} \dagger E_m^{(l-1)}, \dagger E_m^{(0)}) + \epsilon) \times \tilde{\mathcal{A}} \dagger E_m^{(l-1)}, \quad (4)$$

where $\dagger E_m^{(l)}$ denotes the refined representations of users and items in l -th graph convolution layer for the modality m , and $\dagger E_m^{(0)} = \text{Con}(E_{u_m}, \dagger E_{i_m}^{(0)})$ is the initial layer embedding, which concatenated by random initialed user embedding E_{u_m} and refined modality embedding $\dagger E_{i_m}^{(0)}$. $\tilde{\mathcal{A}}$ is a symmetrically normalized matrix, $\text{Sim}(\cdot)$ denotes the cosine similarity function, and ϵ is a small positive bias to avoid a zero vector. The final embedding for each modality is updated by summation:

$$\dagger \bar{E}_m = \sum_{l=0}^L \dagger E_m^{(l)}, \quad (5)$$

where L is the total number of layers.

3.4 Late Fusion

In Section 1, we explored the impact of fusion strategies on representation learning. Previous work [38] also pointed out that existing fusion strategies may fail to capture modality-specific features and even corrupt the learned single-modality representation. To improve the quality of the modality fusion, we propose a tailored early fusion strategy to refine all modalities via behavior modality, which effectively alleviates the gap between modalities. After utilizing heterogeneous graphs to extract the hidden relationships between users and items, we fuse all the refined modality representations using attentive weights. Notably, our work is the first to effectively combine both early fusion and late fusion strategies to enhance the models' representation learning capabilities.

$$\dagger E_f = \text{Con}(\alpha_m \times \dagger \bar{E}_m | m = i, t, v), \quad (6)$$

where $\dagger E_f$ is the final fused representation for all refined modalities. $\text{Con}(\cdot)$ denotes the adaptive attention concatenation, and α_m is a learnable attention weight for each modality.

Conducting late fusion before enhancing the representations of users and items by homogeneous graphs can avoid the $|\mathcal{M}|$ times computation and storage cost of constructing homogeneous graphs for each modality separately and allow homogeneous graphs to capture features with modal synergy⁴ [5, 41].

3.5 Homogeneous Graph

To improve the recommendation performance by enhancing the representations for users and items, we construct the user-user and item-item homogeneous graphs.

⁴Partial information decomposition [1, 26] states the multivariate mutual information is decomposed into three forms of interaction (Uniqueness, Redundancy, and Synergy). Recent study [5] finds that interaction across different modalities allows homogeneous graphs to capture features with modal synergy.

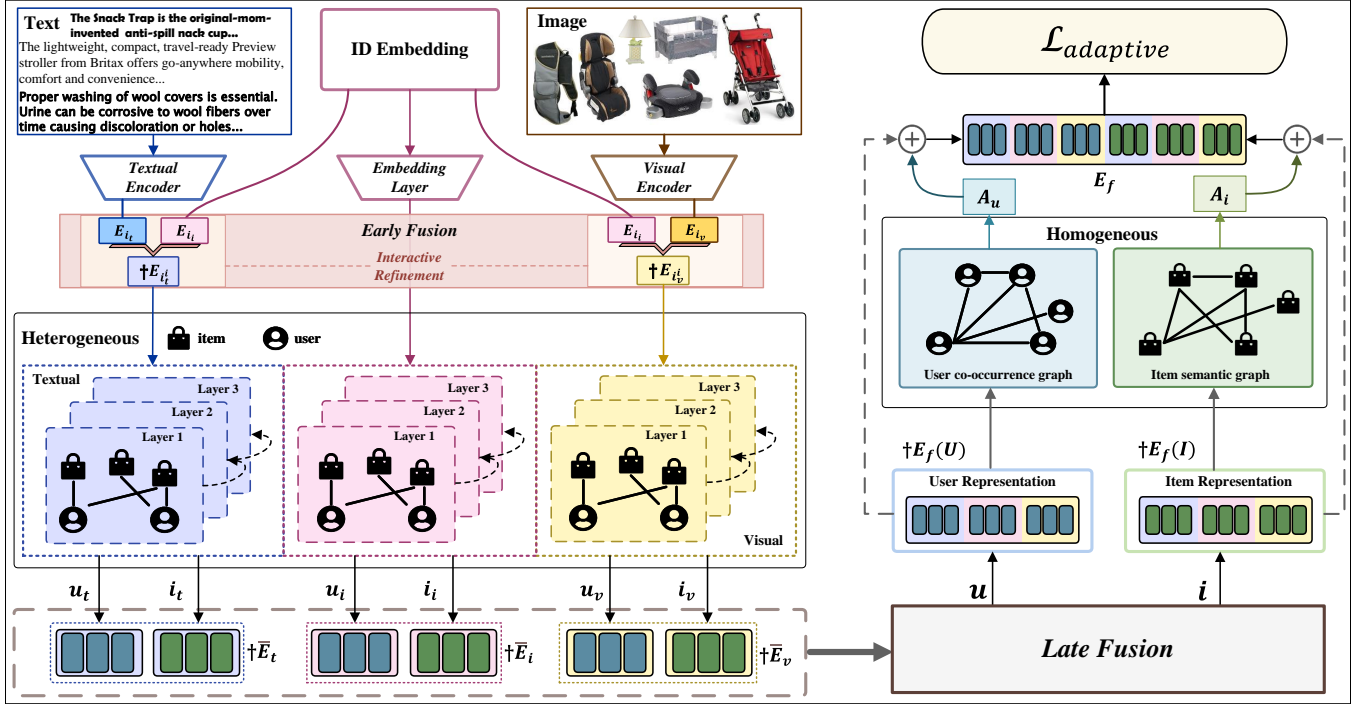


Figure 1: The overview structure of COHESION.

3.5.1 User-User Graph. We construct the user-user graph to explore relations between users. Considering the computation overhead, we only capture a subset of users with the most similar preferences. To this end, the pairwise cosine similarity between all users is calculated as:

$$s_{u,\tilde{u}} = \frac{(e_u)^\top e_{\tilde{u}}}{\|e_u\| \|e_{\tilde{u}}\|}. \quad (7)$$

For each user u , we retain the edges with top- k users $S_u^k = \{s_{u,\tilde{u}} | u, \tilde{u} \in \mathcal{U}\}$ among all user-user pairs:

$$s_{u,\tilde{u}} = \begin{cases} s_{u,\tilde{u}} & \text{if } s_{u,\tilde{u}} \in \text{top-}k(s_{u,v} | v \in \mathcal{U}) \\ 0 & \text{otherwise} \end{cases}, \quad (8)$$

where $s_{u,v} | v \in \mathcal{U}$ represents the neighbor scores of v for the user u . For the constructed user-user graph $\mathcal{G}^u = \{\mathcal{U}, S_u^k\}$, we apply the Softmax function for propagation:

$$A_u^{(l)} = A_u^{(l-1)} + \sum_{s_{u,\tilde{u}} \in S_u^k} \frac{\exp(s_{u,\tilde{u}})}{\sum_{s_{u,\tilde{u}} \in S_u} \exp(s_{u,\tilde{u}})} A_{\tilde{u}}^{(l-1)}, \quad (9)$$

where $A_u^{(l)}$ is l -th layer representation of user u learned from \mathcal{G}^u . $A_u^{(0)}$ is the user side $E_f(U)$ of $\dagger E_f$, which is obtained by $\dagger E_f[|\mathcal{U}|]$.

3.5.2 Item-Item Graph. To explore relations between items, we construct an item-item graph to capture semantically correlative signals. However, the dense graph introduces huge noise, and the sparse graph leads to modality features under-exploited [23]. Therefore, we only consider the top- k most similar items $S_i^k = \{s_{i,\tilde{i}} | i, \tilde{i} \in \mathcal{I}\}$ for each item $i \in \mathcal{I}$. To this end, the pairwise cosine

similarity between all items is calculated as follows:

$$s_{i,\tilde{i}} = \frac{(s_i)^\top s_{\tilde{i}}}{\|s_i\| \|s_{\tilde{i}}\|}. \quad (10)$$

Same as the user-user graph, for each item i , we retain the edges with top- k items $S_i^k = \{s_{i,\tilde{i}} | i, \tilde{i} \in \mathcal{I}\}$ among all item-item pairs:

$$s_{i,\tilde{i}} = \begin{cases} s_{i,\tilde{i}} & \text{if } s_{i,\tilde{i}} \in \text{top-}k(s_{i,j} | j \in \mathcal{I}) \\ 0 & \text{otherwise} \end{cases}, \quad (11)$$

where $s_{i,j} | j \in \mathcal{I}$ represents the neighbor scores for the item i . We further construct a item semantic graph $\mathcal{G}^i = \{\mathcal{I}, S_i^k\}$, and apply weighted aggregation function for propagation:

$$A_i^{(l)} = \sum_{s_{i,\tilde{i}} \in S_i^k} s_{i,\tilde{i}} \times A_{\tilde{i}}^{(l-1)}, \quad (12)$$

where $A_i^{(l)}$ is l -th layer representation of item i learned from \mathcal{G}^i . $A_i^{(0)}$ is the item side $E_f(I)$ of $\dagger E_f$, which can be obtained by $\dagger E_f[|\mathcal{I}|]$. Then, we enhance the user and item representations by user-user and item-item graphs, respectively:

$$E_f(U) = \dagger E_f(U) + A_u^{L_u}, \quad E_f(I) = \dagger E_f(I) + A_i^{L_i}, \quad (13)$$

$$E_f = \text{Con}(E_f(U) | E_f(I)), \quad (14)$$

where $A_u^{L_u}$ and $A_i^{L_i}$ are the final layer results for the user-user graph and the item-item graph, respectively. $\text{Con}(\cdot)$ denotes a concatenation operation. $E_f(U)$ denotes the final user representation, which can be obtained by $\dagger E_f[|\mathcal{U}|]$, and $E_f(I)$ denotes the final item representation, which can be obtained by $\dagger E_f[|\mathcal{I}|]$.

3.6 Adaptive Optimization

Bayesian Personalized Ranking (BPR) loss [18] are widely adopted to optimize the trainable parameters. In particular, BPR increases the gap between user u 's predicted ratings of the positive item p and negative item n for each triplet $(u, p, n) \in \mathcal{D}$, where \mathcal{D} denotes the training set. Positive item p is the item that has interaction with user u , and negative item n is the item randomly sampled from items without any interaction with user u .

Previous work [4] pointed out that the BPR loss is limited to optimizing the overall user preferences and does not directly optimize the learning of modality-specific preferences, i.e., user interests specific to each modality. As a result, there could be some modality-specific interest that is not well-learned, potentially leading to unreliable final user rating predictions. To better learn the features of all modalities, we propose an adaptive BPR loss function, which pays more attention to the weak modality. The adaptive BPR loss function is defined as follows:

$$w_{u,p,n}^m = 1 - \frac{\exp(y_{u,p}^m - y_{u,n}^m)}{\sum_{m \in M} \exp(y_{u,p}^m - y_{u,n}^m)}, \quad (15)$$

$$\mathcal{L} = - \sum_{(u,p,n) \in \mathcal{D}} \log(\sigma(\sum_{m \in M} w_{u,p,n}^m (y_{u,p}^m - y_{u,n}^m))) + \lambda(\|\Theta\|_2^2), \quad (16)$$

where $w_{u,p,n}^m$ is the weight for each modality calculated by reversing the rating gap, and the rating gap is the distance between user u 's predicted ratings of the positive item $y_{u,p}^m$ and the negative item $y_{u,n}^m$ for each modality. $\sigma(\cdot)$ is the Sigmoid function, λ is a hyper-parameter for regularization term, and Θ denotes model parameters.

To provide a clearer overview of our COHESION, we summarize the learning process of COHESION in Algorithm 1.

4 Experiments

In this section, we conduct extensive experiments on some widely used real-world datasets. Experiment results can evaluate the following questions: **RQ1**: How does COHESION perform compared with various state-of-the-art models? **RQ2**: How do the key components in COHESION influence the performance of recommendation accuracy? **RQ3**: Can COHESION effectively mitigate the irrelevant information among modalities? **RQ4**: How does COHESION perform compared with various state-of-the-art models in different sparse data scenarios? **RQ5**: How efficient is COHESION compared with various state-of-the-art models? **RQ6**: How sensitive is COHESION under the perturbation of hyper-parameters?

Table 2: Statistics of three experimented datasets with modalities item Visual(V) and Textual(T).

Dataset	Baby		Sports		Clothing	
Modality	V	T	V	T	V	T
Embed Dim	4,096	384	4,096	384	4,096	384
User	19,445		35,598		39,387	
Item	7,050		18,357		23,033	
Interaction	160,792		296,337		278,677	
Sparsity	99.88%		99.95%		99.97%	

Algorithm 1 Learning Process of COHESION

- 1: **Input:** $\mathcal{U}, \mathcal{I}, \mathcal{M}, \mathcal{G}$, layer number L of heterogeneous graph \mathcal{G} , layer number L_u of user-user graph \mathcal{G}^u , and layer number L_i of item-item graph \mathcal{G}^i
- 2: **Output:** Optimization loss \mathcal{L}
- 3: Initialize E_u, E_{i_m} ;
- 4: Transform item embeddings of all modalities into the same dimension $\dagger E_{i_m} \leftarrow E_{i_m}$ with Eq.2;
- 5: (Early Fusion) Refine item embeddings of each modality by behavior modality $\dagger E_{i_m} \leftarrow \dagger E_{i_m}$ with Eq.3;
- 6: **for** $l = 1 \dots L$ **do**
- 7: Conduct message propagation within the residual-based heterogeneous graph $\dagger E_m^{(l)} \leftarrow \dagger E_m^{(l-1)}$ with Eq.4;
- 8: **end for**
- 9: Get embedding $\dagger \bar{E}_m$ for each modality with Eq.5;
- 10: (Late Fusion) Attentively fuse all modality embeddings $\dagger E_f \leftarrow \dagger \bar{E}_m$ with Eq.6;
- 11: Construct the user-user graph \mathcal{G}^u with Eq.7 and Eq.8.
- 12: Construct the item-item graph \mathcal{G}^i with Eq.10 and Eq.11.
- 13: **for** $l = 1 \dots L_u$ **do**
- 14: Conduct message propagation within the user-user homogeneous graph $A_u^{(l)} \leftarrow A_u^{(l-1)}$ with Eq.9;
- 15: **end for**
- 16: **for** $l = 1 \dots L_i$ **do**
- 17: Conduct message propagation within the item-item homogeneous graph $A_i^{(l)} \leftarrow A_i^{(l-1)}$ with Eq.12;
- 18: **end for**
- 19: Get final enhanced embedding E_f by both user-user and item-item graphs with Eq.13 and Eq.14.
- 20: Calculate adaptive BPR loss \mathcal{L} with Eq.15 and Eq.16.

4.1 Settings

4.1.1 Datasets. Our experiments are conducted on three real-world datasets: Baby, Sports, and Clothing from the Amazon dataset [16], which include visual and textual modalities for each item. Table 2 shows the statistics of these datasets. Following most previous works [31, 34, 42], we adopt the 5-core setting to filter users and items for each dataset. We use the pre-trained sentence-transformer to extract textual features with 384 dimensions and use the published 4096-dimensional visual features as in MMRec [39]. For each dataset, we use the ratio 8:1:1 to split the historical interactions randomly for training, validation, and testing.

4.1.2 Baseline. To evaluate the effectiveness of our COHESION, we compare it with the following methods, which can be divided into two groups. We detailed all of them as follows:

1) Traditional recommendation models:

- **MF-BPR** [18] utilizes the BPR loss to optimize the classic collaborative filtering method, which learns user and item representations by the matrix factorization method.
- **LightGCN** [9] excludes some unnecessary components from GCN-based collaborative filtering to make it more appropriate for recommendation.
- **SimGCL** [35] proposes a graph contrastive learning method, which directly injects random noises into the representation.

Table 3: Performance comparison of baselines and COHESION in terms of Recall@K (R@K) and NDCG@K (N@K). The superscript * indicates the improvement is statistically significant where the p-value is less than 0.01.

Datasets	Baby				Sports				Clothing			
Model	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20
MF-BPR	0.0357	0.0575	0.0192	0.0249	0.0432	0.0653	0.0241	0.0298	0.0187	0.0279	0.0103	0.0126
LightGCN	0.0479	0.0754	0.0257	0.0328	0.0569	0.0864	0.0311	0.0387	0.0340	0.0526	0.0188	0.0236
SimGCL	0.0513	0.0804	0.0273	0.0350	0.0601	0.0919	0.0327	0.0414	0.0356	0.0549	0.0195	0.0244
LayerGCN	0.0529	0.0820	0.0281	0.0355	0.0594	0.0916	0.0323	0.0406	0.0371	0.0566	0.0200	0.0247
VBPR	0.0423	0.0663	0.0223	0.0284	0.0558	0.0856	0.0307	0.0384	0.0281	0.0415	0.0158	0.0192
MMGCN	0.0378	0.0615	0.0200	0.0261	0.0370	0.0605	0.0193	0.0254	0.0218	0.0345	0.0110	0.0142
DualGNN	0.0448	0.0716	0.0240	0.0309	0.0568	0.0859	0.0310	0.0385	0.0454	0.0683	0.0241	0.0299
SLMRec	0.0529	0.0775	0.0290	0.0353	0.0663	0.0990	0.0365	0.0450	0.0452	0.0675	0.0247	0.0303
LATTICE	0.0547	0.0850	0.0292	0.0370	0.0620	0.0953	0.0335	0.0421	0.0492	0.0733	0.0268	0.0330
FREEDOM	0.0627	<u>0.0992</u>	0.0330	0.0424	0.0717	<u>0.1089</u>	0.0385	<u>0.0481</u>	<u>0.0628</u>	<u>0.0941</u>	<u>0.0341</u>	<u>0.0420</u>
BM3	0.0564	0.0883	0.0301	0.0383	0.0656	0.0980	0.0355	0.0438	0.0422	0.0621	0.0231	0.0281
LGMRec	<u>0.0639</u>	0.0989	<u>0.0337</u>	<u>0.0430</u>	<u>0.0719</u>	0.1068	<u>0.0387</u>	0.0477	0.0555	0.0828	0.0302	0.0371
COHESION	0.0680*	0.1052*	0.0354*	0.0454*	0.0752*	0.1137*	0.0409*	0.0503*	0.0665*	0.0983*	0.0358*	0.0438*
Improv.	6.42%	6.05%	5.04%	5.58%	4.59%	4.41%	5.68%	4.57%	5.89%	4.46%	4.99%	4.28%

- **LayerGCN** [40] utilizes the residual connection to build a layer-refined GCN to alleviate LightGCN’s over-smoothing problem.

2) Multimodal recommendation models:

- **VBPR** [8] integrates the visual and textual features with ID embeddings as side information for each item, which can be seen as multimodal matrix factorization.
- **MMGCN** [25] utilizes GCN for each modality to learn the modality-specific features, and then fuses all the user predicted ratings for all modalities to get the final predicted rating.
- **DualGNN** [23] proposes a user-user graph to explore the hidden preference pattern.
- **LATTICE** [36] proposes an item-item graph to capture semantically correlated signals among items.
- **FREEDOM** [41] optimizes LATTICE by freezing the item-item graph and denoising the user-item graph.
- **SLMRec** [20] proposes a self-supervised framework for the multimodal recommendation, which builds a node self-discrimination task to uncover the hidden multimodal patterns of items.
- **BM3** [42] simplifies the self-supervised framework by directly perturbing the representation through a dropout mechanism.
- **LGMRec** [7] utilizes local embeddings with local topological information and global embeddings with hypergraphs.

4.1.3 Evaluation Metrics. For a fair comparison, we follow the settings of previous works [23, 41, 42] to adopt two widely-used evaluation metrics for top-K recommendation: Recall@K and NDCG@K. We report the average scores for all users in the test dataset under both $K = 10$ and $K = 20$, respectively.

4.1.4 Implementation Details. Following the existing work [39], we fix the user and item embedding size to 64 for all models, and then initialize them with the Xavier method [6]. Meanwhile, we optimize all models with Adam optimizer [13]. For each baselines, we perform a grid search for all the hyper-parameters following its published paper to find the optimal setting. All models are implemented by PyTorch and evaluated on an RTX 4070Ti GPU card. For our proposed COHESION, we perform a grid search on the learning

rate in $\{1e-1, 1e-2, 1e-3, 1e-4\}$, the regularization weight λ in $\{1e-1, 1e-2, 1e-3, 1e-4\}$, and residual-based GCN layers L in $\{1, 2, 3, 4\}$. We empirically fix the number of GCN layers in the homogeneous graph (both user-user graph and item-item graph) with $L_u/L_i = 1$ [23, 36]. The k of top- k in the user-user graph is also empirically set as 10 [23, 36, 41]. To avoid over-fitting, the early stopping and total epochs are fixed at 20 and 1000, respectively. Following MMRRec [39], we use Recall@20 on the validation dataset as the indicator for best record updating.

4.2 Performance Comparison (RQ1)

Table 3 shows the performance comparison of our proposed COHESION with other baselines on three widely used datasets. The optimal and optimal results are indicated in **bold** and underlined, respectively. We have the following observations:

- Most multimodal models achieve better performance than traditional recommendation models, which verifies that multimodal information can effectively alleviate the data sparsity problem in traditional recommendation systems.
- Some multimodal recommendation methods result in inconsistent performance across various scenarios. For example, LATTICE outperforms SLMRec on the Baby dataset, but is less effective on the Sports dataset. The self-supervised learning works better in the Sports dataset than in the Baby dataset.
- COHESION makes a huge improvement over all multimodal recommendation models. This validates the effectiveness of our COHESION. Our dual-stage fusion strategy effectively bridges different modalities by refining each modality via behavior modality. Our well-designed composite graph convolutional network leverages both heterogeneous and homogeneous graphs to improve the learning representation ability. Moreover, we introduce an adaptive optimization, which is designed to enhance the comprehensive learning of under-learned modalities. As a result, COHESION outperforms existing multimodal methods. For example, COHESION improves the strongest baseline (FREEDOM

or LGMRec) in terms of R@20 on the Baby and Sports datasets by 6.05% and 4.41%, respectively.

- DualGNN and LATTICE both achieve competitive performance by adding a user-user and item-item graph, respectively. Our COHESION significantly outperforms these two methods. We owe our superiority to our novel dual-stage fusion strategy liberates the representational capabilities of a well-designed composite graph convolutional network.

4.3 Ablation Study (RQ2)

In this section, we empirically evaluate the effectiveness of different components of COHESION.

4.3.1 Effect of Early Fusion. To investigate the effects of the early fusion component in COHESION, we set up the following model variants:

- **COHESION-itv**: Removing the entire early fusion process.
- **COHESION-tv**: Removing refinement operation of textual and visual modalities in the early fusion process.
- **COHESION-iv**: Removing the refinement operation of behavior and visual modalities in the early fusion process.
- **COHESION-it**: Removing the refinement operation of behavior and textual modalities in the early fusion process.
- **COHESION-v**: Removing the refinement operation of visual modality in the early fusion process.
- **COHESION-t**: Removing the refinement operation of textual modality in the early fusion process.
- **COHESION-i**: Removing the refinement operation of behavior modality in the early fusion process.




Figure 2: Effect of early fusion.

Fig. 2 shows the results (Recall@20) of these variants and COHESION on all three datasets. We have the following observations:

- 1) Results of **COHESION-itv** < **COHESION-tv**, **COHESION-iv** < **COHESION-v**, and **COHESION-it** < **COHESION-t** showing that refining behavior modality also leads to a slight improvement in recommendation performance.
- 2) COHESION performs better than **COHESION-i**, **COHESION-v**, and **COHESION-t**, which shows that refining any modality can lead to improve the recommendation accuracy.
- 3) COHESION outperforms all seven variants, verifying the effectiveness of our early fusion strategy.

4.3.2 Effect of Homogeneous Graphs. To identify the effectiveness of homogeneous graphs in COHESION, we set up the following variants: 1) **COHESION-ui**, which removes both user-user and

Table 4: Effect of homogeneous graphs in terms of Recall@K (R@K) and NDCG@K (N@K).

Dataset	Variants	R@10	R@20	N@10	N@20
Baby	COHESION	0.0680	0.1052	0.0354	0.0454
	COHESION-u	0.0660	0.1024	0.0349	0.0443
	COHESION-i	0.0564	0.0884	0.0307	0.0389
	COHESION-ui	0.0547	0.0870	0.0291	0.0375
Sports	COHESION	0.0752	0.1137	0.0409	0.0503
	COHESION-u	0.0705	0.1086	0.0384	0.0482
	COHESION-i	0.0622	0.0963	0.0337	0.0425
	COHESION-ui	0.0603	0.0937	0.0327	0.0413
Clothing	COHESION	0.0665	0.0983	0.0358	0.0438
	COHESION-u	0.0630	0.0942	0.0339	0.0417
	COHESION-i	0.0499	0.0744	0.0271	0.0331
	COHESION-ui	0.0478	0.0701	0.0262	0.0319

item-item graphs. 2) **COHESION-i**, which removes only the item-item graph. 3) **COHESION-u**, which removes only the user-user graph. As illustrated in Table 4, we observe that both user-user and item-item graphs could improve the performance of recommendations. The item-item graph can better represent the user's preference than the user-user graph, which is identified by the results **COHESION-ui** < **COHESION-i** < **COHESION-u** < **COHESION**. Note that user-user and item-item graphs can improve recommendation accuracy together to achieve higher performance.

Table 5: Effect of adaptive optimization in terms of Recall@K (R@K) and NDCG@K (N@K).

Dataset	Variants	R@10	R@20	N@10	N@20
Baby	DualGNN	0.0448	0.0716	0.0240	0.0309
	DualGNN+	0.0492	0.0753	0.0261	0.0328
	LATTICE	0.0547	0.0850	0.0292	0.0370
	LATTICE+	0.0573	0.0881	0.0303	0.0388
Sports	DualGNN	0.0568	0.0859	0.0310	0.0385
	DualGNN+	0.0587	0.0894	0.0316	0.0394
	LATTICE	0.0620	0.0953	0.0335	0.0421
	LATTICE+	0.0634	0.0974	0.0343	0.0430
Clothing	DualGNN	0.0454	0.0683	0.0241	0.0299
	DualGNN+	0.0499	0.0737	0.0270	0.0334
	LATTICE	0.0492	0.0733	0.0268	0.0330
	LATTICE+	0.0524	0.0787	0.0285	0.0354

4.3.3 Effect of Adaptive Optimization. Our adaptive optimization could serve as a play-and-plug component for other models to improve their performance. To evaluate the effectiveness of adaptive optimization in other models, we design the corresponding model variants, named **DualGNN+** and **LATTICE+**, where symbol + means using adaptive optimization. As shown in Table 5, our adaptive optimization achieves significant improvement for all models on all datasets.

Table 6: Performance comparison on all three datasets (Recall@20). The notation Visual means only adding visual modality, Textual means only adding textual modality, and Multimodal means adding both visual and textual modalities. Improv. denotes the improvement of adding multimodal compared to adding the optimal single modality.

Dataset	Modality	VBPR	MMGCN	DualGNN	SLMRec	LATTICE	FREEDOM	BM3	LGMRec	COHESION
Baby	Visual	0.0665	0.0634	0.0830	0.0789	0.0779	0.0816	0.0806	0.0954	0.0878
	Textual	0.0638	0.0623	0.0943	0.0797	0.0880	0.0980	0.0842	0.0997	0.1002
	Multimodal Improv.	0.0663	0.0615	0.0716	0.0775	0.0848	0.0992	0.0833	0.0989	0.1052
Sports	Visual	0.0849	0.0574	0.0896	0.0928	0.0902	0.0879	0.0973	0.1059	0.1085
	Textual	0.0714	0.0623	0.0958	0.1011	0.0978	0.1083	0.0989	0.1093	0.1099
	Multimodal Improv.	0.0856	0.0605	0.0859	0.0990	0.0953	0.1089	0.0980	0.1068	0.1137
Clothing	Visual	0.0433	0.0363	0.0730	0.0662	0.0705	0.0816	0.0608	0.0805	0.0928
	Textual	0.0397	0.0358	0.0784	0.0698	0.0767	0.0937	0.0650	0.0847	0.0956
	Multimodal Improv.	0.0415	0.0345	0.0683	0.0675	0.0733	0.0941	0.0621	0.0828	0.0983




Figure 3: Visualization of distribution through t-SNE. We randomly sample 500 items from Baby dataset. Red, green, and blue represent the ID, textual, and visual modalities.

4.4 Effectiveness of Mitigating Irrelevant Information among Modalities (RQ3)

The negative effect of irrelevant information among modalities will decrease modal representation learning ability and even lead to a weaker performance for incorporating multiple modalities than incorporating a single modality. Table 6 shows that most existing models suffer from this dilemma, we observe that only FREEDOM and our COHESION achieve higher performance for incorporating multiple modalities and that COHESION made a much more significant improvement of 4.99%. It verifies that our early fusion strategy bridges different modalities to mitigate the negative effect of irrelevant information among modalities.

To gain a better understanding of the advantages of our dual-stage strategy for mitigating irrelevant information among modalities, we take a step: randomly pick 500 items from the Baby dataset and use the t-SNE [22] algorithm to map the representations of late fusion strategy and our dual-stage strategy into a two-dimensional space, respectively. Fig. 3 visualizes the distributions of all modalities. The distributions of all modalities from our dual-stage fusion strategy exhibit significantly greater similarity compared to those from the late fusion strategy, further validating the effectiveness




Figure 4: Sparsity degree analysis on three datasets.

of our dual-stage fusion approach in aligning modality representations.

4.5 Different Sparse Data Scenarios (RQ4)

In this section, we verify the effectiveness of COHESION in different sparse data scenarios. We conduct experiments on sub-datasets with varying levels of data sparsity using all three datasets. We compare the performance of COHESION against three outstanding baselines (LATTICE, FREEDOM, LGMRec). We split user groups by the number of interactions in the training set (e.g., the first group consists of users with 0-5 interacted items). Fig. 4 shows that COHESION consistently outperforms all baselines on all datasets with different degrees of sparsity, which demonstrates the effectiveness in different sparse data scenarios.

4.6 Efficiency Analysis (RQ5)

4.6.1 Time Complexity. We provide a theoretical time complexity analysis for our COHESION. Firstly, the heterogeneous graph cost is $O(2 \times L \times |\mathcal{G}| \times d/B)$, where $|\mathcal{G}|$ is the number of edges, d is the embedding dimension, B is the batch size, and L is the GCN layers number. Secondly, the homogeneous graph cost is $O((L_u |\mathcal{U}|^2 + L_i |\mathcal{I}|^2) \times d)$, where L_u and L_i are the number of a user-user graph and an item-item graph, respectively. Lastly, the BPR loss cost is $O(2 \times d \times B)$. Note that COHESION doesn't adopt self-supervised tasks, which makes it more efficient than SSL-based models. Since we conduct our late fusion strategy before enhancing representation by the user-user and item-item graphs, we save $|\mathcal{M}|$ times of computation


Table 7: Comparison of COHESION against competitive baselines on efficiency. Parameter denotes the learnable parameter for the model, and Time denotes the total duration of each iteration, including both the training and evaluation phases.

Dataset	Metrics	VBPR	MMGCN	DualGNN	SLMRec	LATTICE	FREEDOM	BM3	LGMRec	COHESION
Baby	Time (s/epoch)	2.12	3.96	5.48	3.75	4.13	2.77	2.55	5.93	4.47
	Parameter (M)	3.2	1.4	3.8	2.0	33.6	33.6	33.6	38.6	37.0
Sports	Time (s/epoch)	3.55	11.01	15.90	6.14	11.90	5.86	5.39	8.98	7.91
	Parameter (M)	6.0	1.4	5.9	3.8	86.0	86.0	86.0	99.9	91.5
Clothing	Time (s/epoch)	3.90	13.04	16.93	6.94	14.10	6.29	6.09	10.02	9.05
	Parameter (M)	6.8	1.4	6.4	4.3	107.5	107.5	107.5	126.1	113.4

and storage resources compared to DualGNN, LATTICE, and FREEDOM. This allows us to maintain similar computation and storage resource requirements as competitive baselines when using both the user-user and item-item graphs.

4.6.2 Computational Resource. Moreover, we report the training time (including both computation time and evaluating time) and model parameters for our COHESION and baselines in Table 7. We observe that our COHESION achieves a significant improvement over competitive baselines (LATTICE, FREEDOM, BM3, and LGMRec) while requiring comparable computation resources and training time. This observation underscores the efficiency of our COHESION model, highlighting its ability to achieve superior recommendation performance without necessitating an increase in resource consumption or training time, thereby demonstrating its practical applicability in multimodal recommendation scenarios.



4.6.3 Convergency. In Figure 5, we show the training curves of our COHESION and competitive baselines (LATTICE, FREEDOM, and LGMRec) on all three datasets, as the numbers of iterations and epochs increase. The faster convergence of our COHESION method is obviously observed, which suggests the advantage of COHESION in training efficiency while maintaining superior recommendation accuracy. This enhanced performance can be attributed to our innovative dual-stage fusion strategy, which effectively mitigates the influence of irrelevant information across modalities and brings a positive effect for model optimization to achieve fast convergence. Our dual-stage fusion strategy not only expedites the convergence speed but also enhances recommendation performance.

**Figure 5: Convergence study in terms of Recall@20.**

4.7 Hyper-parameter Sensitivity Study (RQ6)

4.7.1 Number of Layer L of Heterogeneous Graph. As shown in Fig. 6, we find that the optimal number of layer L for Baby, Sports, and Clothing datasets are 1, 2, and 2, respectively. The reason is that the degree of association between users and items is smaller in the Baby dataset than in the Sports and Clothing datasets. It is

important to emphasize that for all other multimodal recommendation methods using GCN, the best number of layers of GCN is 1 for all these three datasets. It verifies that residual-based GCN can alleviate the over-smoothing problem to some extent.

**Figure 6: Effect of the number of heterogeneous graph layers.****Figure 7: Performance of COHESION with respect to different regularization weight λ and learning rate.**

4.7.2 The Learning Rate and Regularization Weight. Fig. 7 shows the results of COHESION under different learning rates and regularization weights on Baby, Sports, and Clothing datasets. We can observe that $1e-3$, $1e-4$, and $1e-3$ are the suggested learning rates for the Baby, Sports, and Clothing datasets, respectively. Moreover, the suggested regularization weights are $1e-4$, $1e-3$, and $1e-3$ for the Baby, Sports, and Clothing datasets, respectively.

5 Conclusion

In this paper, we examine the limitations of the fusion strategy in previous works and further reveal the complementary capability between fusion strategy and representation learning. To this end, we propose a composite graph convolutional network with a dual-stage fusion strategy for multimodal recommendation, named COHESION. More specifically, our well-designed composite graph convolutional network achieves outstanding representation learning ability by leveraging both heterogeneous and homogeneous graphs. Moreover, our tailored dual-stage fusion strategy further liberates the representation learning ability of our composite graph convolutional network by mitigating the negative effect of irrelevant information among modalities. Finally, we propose an adaptive optimization to control the balance of attention among all modalities. We conduct comprehensive experiments and ablation studies to demonstrate the effectiveness of our COHESION.

References

- [1] Nils Bertschinger, Johannes Rauh, Eckehard Olbrich, Jürgen Jost, and Nihat Ay. 2014. Quantifying unique information. *Entropy* 16, 4 (2014), 2161–2183.
- [2] Jingyuan Chen, Hanwang Zhang, Xiangnan He, Liqiang Nie, Wei Liu, and Tat-Seng Chua. 2017. Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*. 335–344.
- [3] Zheyu Chen, Jinfeng Xu, and Haibo Hu. 2025. Don't Lose Yourself: Boosting Multimodal Recommendation via Reducing Node-neighbor Discrepancy in Graph Convolutional Network. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1–5.
- [4] Xue Dong, Xuemeng Song, Minghui Tian, and Linmei Hu. 2024. Prompt-based and weak-modality enhanced multimodal recommendation. *Information Fusion* 101 (2024), 101989.
- [5] Benoit Dufumier, Javiera Castillo-Navarro, Devis Tuia, and Jean-Philippe Thiran. 2024. What to align in multimodal contrastive learning? *arXiv preprint arXiv:2409.07402* (2024).
- [6] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, 249–256.
- [7] Zhiqiang Guo, Jianjun Li, Guohui Li, Chaoyang Wang, Si Shi, and Bin Ruan. 2024. LGMRec: Local and Global Graph Learning for Multimodal Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 8454–8462.
- [8] Ruining He and Julian McAuley. 2016. VBPR: visual bayesian personalized ranking from implicit feedback. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 30.
- [9] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [10] Yangqin Jiang, Lianghao Xia, Wei Wei, Da Luo, Kangyi Lin, and Chao Huang. 2024. DiffMM: Multi-Modal Diffusion Model for Recommendation. (2024).
- [11] Taeri Kim, Yeon-Chang Lee, Kijung Shin, and Sang-Wook Kim. 2022. MARIO: modality-aware attention and modality-preserving decoders for multimedia recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 993–1002.
- [12] Yungi Kim, Taeri Kim, Won-Yong Shin, and Sang-Wook Kim. 2024. MONET: Modality-Embracing Graph Convolutional Network and Target-Aware Attention for Multimedia Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 332–340.
- [13] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [14] Yehuda Koren, Steffen Rendle, and Robert Bell. 2021. Advances in collaborative filtering. *Recommender systems handbook* (2021), 91–142.
- [15] Fei Lei, Zhongqi Cao, Yuning Yang, Yibo Ding, and Cong Zhang. 2023. Learning the User's Deeper Preferences for Multi-modal Recommendation Systems. *ACM Transactions on Multimedia Computing, Communications and Applications* 19, 3s (2023), 1–18.
- [16] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*. 43–52.
- [17] Zongshen Mu, Yueting Zhuang, Jie Tan, Jun Xiao, and Siliang Tang. 2022. Learning hybrid behavior patterns for multimedia recommendation. In *Proceedings of the 30th ACM International Conference on Multimedia*. 376–384.
- [18] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [19] Jinhui Tang, Xiaoyu Du, Xiangnan He, Fajie Yuan, Qi Tian, and Tat-Seng Chua. 2019. Adversarial training towards robust multimedia recommender system. *IEEE Transactions on Knowledge and Data Engineering* 32, 5 (2019), 855–867.
- [20] Zhulin Tao, Xiaohao Liu, Yewei Xia, Xiang Wang, Lifang Yang, Xianglin Huang, and Tat-Seng Chua. 2022. Self-supervised learning for multimedia recommendation. *IEEE Transactions on Multimedia* (2022).
- [21] Nhu-Thuat Tran and Hady W Lauw. 2022. Aligning Dual Disentangled User Representations from Ratings and Textual Content. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1798–1806.
- [22] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [23] Qifan Wang, Yinwei Wei, Jianhua Yin, Jianlong Wu, Xuemeng Song, and Liqiang Nie. 2021. Dualgnn: Dual graph neural network for multimedia recommendation. *IEEE Transactions on Multimedia* (2021).
- [24] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, and Tat-Seng Chua. 2020. Graph-refined convolutional network for multimedia recommendation with implicit feedback. In *Proceedings of the 28th ACM international conference on multimedia*. 3541–3549.
- [25] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 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.
- [26] Paul I. Williams and Randall D Beer. 2010. Nonnegative decomposition of multivariate information. *arXiv preprint arXiv:1004.2515* (2010).
- [27] Jinfeng Xu, Zheyu Chen, Jinze Li, Shuo Yang, Hewei Wang, and Edith CH Ngai. 2024. AlignGroup: Learning and Aligning Group Consensus with Member Preferences for Group Recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. 2682–2691.
- [28] Jinfeng Xu, Zheyu Chen, Jinze Li, Shuo Yang, Wei Wang, Xiping Hu, and Edith C-H Ngai. 2024. FourierKAN-GCF: Fourier Kolmogorov-Arnold Network—An Effective and Efficient Feature Transformation for Graph Collaborative Filtering. *arXiv preprint arXiv:2406.01034* (2024).
- [29] Jinfeng Xu, Zheyu Chen, Zixiao Ma, Jiyi Liu, and Edith CH Ngai. 2024. Improving Consumer Experience With Pre-Purify Temporal-Decay Memory-Based Collaborative Filtering Recommendation for Graduate School Application. *IEEE Transactions on Consumer Electronics* (2024).
- [30] Jinfeng Xu, Zheyu Chen, Shuo Yang, Jinze Li, Hewei Wang, and Edith C-H Ngai. 2024. MENTOR: Multi-level Self-supervised Learning for Multimodal Recommendation. *arXiv preprint arXiv:2402.19407* (2024).
- [31] Jinfeng Xu, Zheyu Chen, Shuo Yang, Jinze Li, Wei Wang, Xiping Hu, Steven Hoi, and Edith Ngai. 2025. A Survey on Multimodal Recommender Systems: Recent Advances and Future Directions. *arXiv preprint arXiv:2502.15711* (2025).
- [32] Jin Xu, Zishan Li, Bowen Du, Miaomiao Zhang, and Jing Liu. 2020. Reluplex made more practical: Leaky ReLU. In *2020 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 1–7.
- [33] Guipeng Xu, Xinyu Li, Ruobing Xie, Chen Lin, Chong Liu, Feng Xia, Zhanhui Kang, and Leyu Lin. 2024. Improving Multi-modal Recommender Systems by Denoising and Aligning Multi-modal Content and User Feedback. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3645–3656.
- [34] Jing Yi and Zhenzhong Chen. 2021. Multi-modal variational graph auto-encoder for recommendation systems. *IEEE Transactions on Multimedia* 24 (2021), 1067–1079.
- [35] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. 2022. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 1294–1303.
- [36] Jinghao Zhang, Yanqiao Zhu, Qiang Liu, Shu Wu, Shuhui Wang, and Liang Wang. 2021. Mining latent structures for multimedia recommendation. In *Proceedings of the 29th ACM International Conference on Multimedia*. 3872–3880.
- [37] Hongyu Zhou, Xin Zhou, Zhiwei Zeng, Lingzi Zhang, and Zhiqi Shen. 2023. A comprehensive survey on multimodal recommender systems: Taxonomy, evaluation, and future directions. *arXiv preprint arXiv:2302.04473* (2023).
- [38] Hongyu Zhou, Xin Zhou, Lingzi Zhang, and Zhiqi Shen. 2023. Enhancing dyadic relations with homogeneous graphs for multimodal recommendation. In *ECAI 2023*. IOS Press, 3123–3130.
- [39] Xin Zhou. 2023. MMRec: Simplifying Multimodal Recommendation. *arXiv preprint arXiv:2302.03497* (2023).
- [40] Xin Zhou, Donghui Lin, Yong Liu, and Chunyan Miao. 2023. Layer-refined graph convolutional networks for recommendation. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 1247–1259.
- [41] Xin Zhou and Zhiqi Shen. 2023. A tale of two graphs: Freezing and denoising graph structures for multimodal recommendation. In *Proceedings of the 31st ACM International Conference on Multimedia*. 935–943.
- [42] Xin Zhou, Hongyu Zhou, Yong Liu, Zhiwei Zeng, Chunyan Miao, Pengwei Wang, Yuan You, and Feijun Jiang. 2023. Bootstrap latent representations for multi-modal recommendation. In *Proceedings of the ACM Web Conference 2023*. 845–854.