A Multi-Interaction Graph Attention Network for Bundle Recommendation

Zhenhua Huang[1,2,3], Chao Guan[1], Pengwei Han[1], and Zhaohong Jia*[1]

[1] Anhui University, Hefei, Anhui, China
[2] University of Science and Technology of China, Hefei, China
[3] Dataspace Institute, Hefei Comprehensive National Science Center, Hefei, China

Abstract. Bundle recommendation aims to provide personalized product bundles, yet existing methods suffer from redundant learning (overlapping user-item and user-bundle interactions) and undifferentiated interaction modeling (uniform aggregation of nodes). To address these challenges, we propose a novel method named **M**ulti-interaction **G**raph attention network for **B**undle **R**ecommendation (**MGBR**), which integrates multi-interaction separation, topology-aware structural learning, and cross-task preference transfer. Specifically, MGBR constructs three homogeneous graphs (user-bundle, user-item, bundle-item) to isolate interaction redundancy and employs a hierarchical graph attention mechanism to dynamically assign weights to interactions (e.g., distinguishing core vs. peripheral items in bundles). To enhance structural discriminability, we propose a structural hint learning module with dual objectives: (1) node degree prediction to preserve user activity and bundle popularity patterns; (2) neighbor degree sum prediction to capture local topological dependencies. Additionally, multi-task learning transfers knowledge between item-level and bundle-level preferences through shared user embeddings. Extensive experiments on two real public datasets demonstrate that MGBR outperforms previous bundle recommendation methods by 1.30% - 4.65% on NetEase and Youshu.

Keywords: Bundle recommendation · Graph attention network · Structural hint learning · Multi-task learning.

1 Introduction

The recommendation system provides personalized preferences for users and has become the most widely applied information system [1]. In recent years, graph neural networks (GNNs)[3] have emerged as a cutting-edge technology for making personalized and context-aware recommendations. GNNs can effectively model complex relationships and dependencies among items or users in a graph structure. They open up possibilities for designing recommendation systems that leverage rich interactions and connections among items or users, which leads to improved recommendation accuracy and diversity[2].

Bundle recommendation [5, 6] is a specific topic in recommendation systems. It provides customers with tailored product recommendations while increasing revenue for businesses. Early works addressed the bundle recommendation problem using parameter sharing or joint loss functions to learn user-item interactions and user-bundle interactions [4, 7]. A GNN-based model was first proposed by Chang et al. [8], which unifies the two interactions and bundle-item affiliation into a heterogeneous graph and employs a sampling method for training. Tan et al. [9] considered the co-purchase and co-occurrence information of the items to model intention-oriented hierarchical representations based on graph convolutional propagation. CrossCBR [24] further adopts contrastive learning to model the cooperative association between user-level and bundle-level graph views to achieve mutual enhancement. Zhu et al.[25] proposed a novel graph learning paradigm called counterfactual learning for bundle recommendation to mitigate the impact of data sparsity problems and improve bundle recommendation. Ma et al.[15] proposed an "early fusion and late contrast" strategy, which effectively captures user preferences and mitigates the sparse bundle-item association problem, ultimately enhancing recommendation performance. Despite the effectiveness of the existing bundle recommendation methods, they still have the following limitations:

Redundancy in learning hinders preference modeling. Previous studies have addressed the interplay between user-bundle interaction information and user-item interaction information in an integrated fashion, which presents a challenge. Specifically, as the user-item interaction information is being learned, the implicit item details are once again included in the learning of the user-bundle interaction information, given that the bundle is influenced by its constituent items. This phenomenon significantly shapes the modeling of the user's bundle preference.

Inability to distinguish the different interactions. Not all items in the bundle are necessarily influential in characterizing the bundle. As shown in Figure 1(a), most of the existing works [8–10] aggregated the information of neighboring nodes based on the graph topology, treated all nodes equally, and did not consider the different importance of nodes in multiple interactions. In our works, as shown in Figure 1(b), where the Solid lines represent existing interactions, the lighter the color, the lower the impact. The dashed lines represent possible interactions. We can visualize that user u is more likely to interact with bundle b_1 rather than b_2.

To address these shortcomings, we propose a model MGBR for bundle recommendation. Specifically, we first construct three bipartite graphs based on interaction and dependency relationships, such as the user-bundle interaction graph and bundle-item affiliation graph. Since not all items in a bundle necessarily contribute to its characteristic features, we incorporate a set of structural hint features (e.g., node degree distribution and neighborhood degree distribution) into the interaction graph construction, which are designed to capture inherent topological properties. This structural learning mechanism enables the model to generate more discriminative node embeddings and graph representations through systematic acquisition of these structural patterns. The enhanced discriminative power ultimately leads to improved accuracy in downstream prediction tasks.

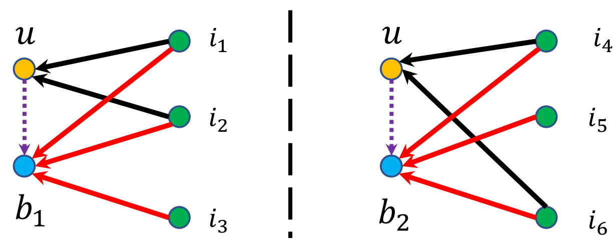

(a) The same type of interaction has the same influences

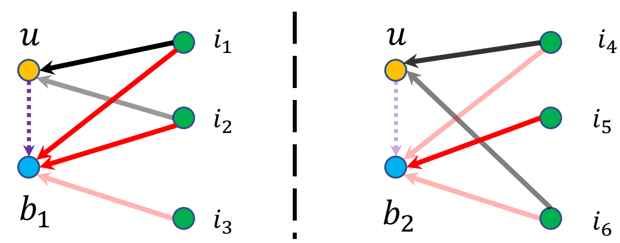

(b) The same type of interaction has the different influences

Fig. 1: Illustration of the different influences among users, bundles, and items.

Then we utilize a graph attention mechanism to consider the relationships between different entities from the constructed graphs and obtain the representations of various entities. Finally, to enhance bundle preference prediction, we leverage multi-task learning by sharing user representations and transferring knowledge of user preferences at the item level. This approach allows us to improve the accuracy of bundle preference prediction by incorporating relevant user information into the model's training process. Comprehensive experiments on two real public datasets show that our proposed approach outperforms various representative bundle recommendation methods. In general, the main contributions of this paper are summarized as follows:

Contributions: (i) We propose a multi-interaction graph attention network to solve the problem of bundle recommendation by incorporating different influences of multi-interactions among users, items, and bundles. (ii) We propose integrating structural hint learning into bundle recommendation through a graph structural loss function that incorporates node and neighborhood degree distributions. This formulation enables dynamic adaptation to heterogeneous bundle structures. (Iii) Extensive experiments on two real-world datasets have demonstrated the effectiveness of the proposed methods. Our proposed method outperforms existing baselines by 4.84% - 11.24%.

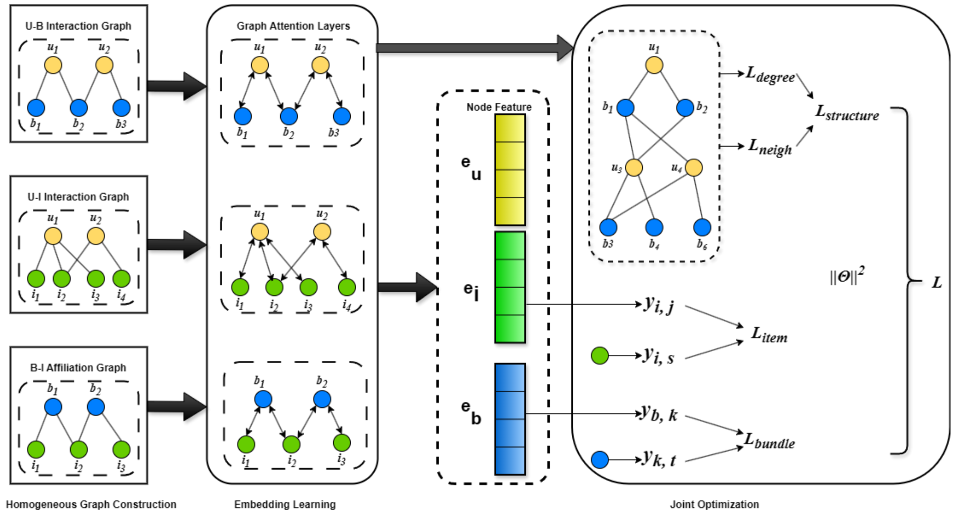

Fig. 2: Schematic illustration of the model MGBR. Yellow nodes are users, blue nodes are bundles, and green nodes are items.

2 Problem Definitions

Suppose we have a set U of N users, a set I of M items, and a set B of K bundles. We define the three matrices, user-bundle interaction matrix, user-item interaction matrix, and bundle-item affiliation matrix as $X^{N\times K} = \{x_{ub} \mid u \in U, b \in B\}$, $Y^{N\times M} = \{y_{ui} \mid u \in U, i \in I\}$, $Z^{K\times M} = \{z_{bi} \mid b \in B, i \in I\}$, respectively. Where $x_{ub} \in \{0, 1\}$ denotes that user u has interacted with bundle b or not. Similarly, $y_{ui} = 1$ denotes that user u has interacted with item i, and $z_{bi} = 1$ means bundle b contains item i. Based on the above definitions, the problem of bundled recommendation is as follows:

Input: Users U, Bundles B, Items I, user-item interaction matrix X, user-bundle interaction matrix Y, and bundle-item affiliation matrix Z.

Output: The probability that a user u will interact with a bundle b.

3 Methodology

The overall architecture of our proposed MGBR is illustrated in Figure 2. The model contains the following three components:

3.1 Homogeneous Graph Construction

Three homogeneous graphs are constructed to capture the interactions among different entities (i.e., users, items, and bundles) in a bundle recommendation scenario. The three graphs are used to model user-item interactions, user-bundle interactions, and bundle-item affiliations, respectively. The three homogeneous graphs represent the relationships and dependencies between these entities in a unified and structured manner. By leveraging graph-based representations, MGBR captures the complex interactions and dependencies among users, items, and bundles, which are crucial for understanding user preferences.

We construct the graphs G_1, G_2, and G_3 based on the existing user-item interaction information, user-bundle interaction information, and bundle-item affiliation information with their corresponding adjacency matrices X, Y, and Z, respectively.

The adjacency matrix of the three homogeneous graphs is defined as:

$$A_o = \begin{pmatrix} I_u & Y \\ Y^T & I_b \end{pmatrix} \tag{1}$$

$$A_p = \begin{pmatrix} I_u & X \\ X^T & I_i \end{pmatrix} \tag{2}$$

$$A_q = \begin{pmatrix} I_b & Z \\ Z^T & I_i \end{pmatrix} \tag{3}$$

where I_u, I_b, and I_i are identity matrices for users, items, and bundles. It is assumed that every node is self-connected.

3.2 Embedding Learning

$E_u \in \mathbb{R}^{N \times d}, E_b \in \mathbb{R}^{K \times d}, E_i \in \mathbb{R}^{M \times d}$ denote the initialized user embedding, bundle embedding, and item embedding, respectively. d is the embedding dimension.

To distinguish the different influences of nodes in multiple interactions, we utilize a graph attention mechanism to learn representations of different nodes in the constructed graphs. The learned embedding capture the complex interactions and dependencies among users, items, and bundles, and serve as compact and informative representations of the entities in the recommendation process. As shown in Figure 2, MGBR calculates the attention value between different nodes and focuses on those interactions with high weights to accurately grasp user preferences. Formally, the model is formulated as:

$$X^{(l+1)} = f(X^{(l)}, A_i), \tag{4}$$

where $f(\cdot)$ means the graph attention mechanism. $X^{(l)}$ are the feature of different nodes at the l-th layer and A_i denotes the adjacent matrix of the homogeneous graphs.

User-Bundle Interaction Graph Learning. By employing the graph attention layer to process the user-bundle interaction matrix, MGBR is able to calculate attention values that capture the relevance between users and bundles. These attention values are then used to update the representations of users and bundles, enabling the model to capture the important interactions between them. Accordingly, we reformulate Eq.4 as:

$$E_{u,o}^{(l+1)}, E_{b,o}^{(l+1)} = f(E_u^{(l)}, E_b^{(l)}, A_o) \tag{5}$$

where $E_u^{(0)}, E_b^{(0)}$ is set as E_u, E_b at initial iteration. The specifically embedding updating rules for user u and bundle b are formulated as follows:

$$\begin{aligned} e_{u,o}^{l+1} &= \sigma(W_1(\alpha_{uu} e_{u,o}^l + \sum_{b \in \mathcal{N}_u} \alpha_{ub} e_{b,o}^l)), \\ e_{b,o}^{l+1} &= \sigma(W_1(\alpha_{bb} e_{b,o}^l + \sum_{u \in \mathcal{N}_b} \alpha_{bu} e_{u,o}^l)). \end{aligned} \tag{6}$$

User-Item Interaction Graph Learning. The users' preference for an individual item in a bundle can significantly influence their overall interest in that bundle. For instance, a user may reject a bundle due to the presence of a disliked item.

We leverage the attention factor to capture not only the user preferences for items but also the characteristics of the items themselves. By incorporating the attention factor, MGBR can effectively capture the nuanced preferences of users towards individual items within bundles, which allows for more accurate modeling of user preferences. Similar to Eq. 6, the embedding updating rules in the user-item interaction graph are defined as follows:

$$\begin{aligned} e_{u,p}^{l+1} &= \sigma(W_2(\alpha_{uu} e_{u,p}^l + \sum_{i \in \mathcal{N}_u} \alpha_{ui} e_{i,p}^l)), \\ e_{i,p}^{l+1} &= \sigma(W_2(\alpha_{ii} e_{i,p}^l + \sum_{u \in \mathcal{N}_i} \alpha_{iu} e_{u,p}^l)). \end{aligned} \tag{7}$$

Bundle-Item affiliation Graph Learning. The items included in a bundle can have varying degrees of importance in determining the quality of the bundle. To capture these dependencies between items and learn accurate representations of bundles, we utilize the existing bundle-item affiliation information. This allows the model to understand the interdependencies among items within a bundle, and incorporate this information into the learned representations of bundles. Similar to Eq. 7, the embedding updating rules for bundle b and item i can be formulated as follows:

$$\begin{aligned} e_{b,q}^{l+1} &= \sigma(W_3(\alpha_{bb} e_{b,q}^l + \sum_{i \in \mathcal{N}_b} \alpha_{bi} e_{i,q}^l)), \\ e_{i,q}^{l+1} &= \sigma(W_3(\alpha_{ii} e_{i,q}^l + \sum_{b \in \mathcal{N}_i} \alpha_{ib} e_{b,q}^l)). \end{aligned} \tag{8}$$

In Equation 6-8, where W_1, W_2, W_3 are trainable matrices, σ is non-linear activation function. N_u, N_b, N_i represent neighbors of user u, bundle b and item i, respectively. α_{ij} denotes the attention value between node i and node j.

3.3 Interaction Prediction and Optimization

Interaction Prediction. After we iterated the above propagation several times, we connect different embeddings of the same entities to combine the information received from different interactions for prediction:

$$\begin{aligned} e_u^{(l)} &= e_{u,o}^{(l)} || e_{u,p}^{(l)}, \\ e_b^{(l)} &= e_{b,o}^{(l)} || e_{b,q}^{(l)}, \\ e_i^{(l)} &= e_{i,p}^{(l)} || e_{i,q}^{(l)}. \end{aligned} \tag{9}$$

Then, we use the inner product to calculate the probability of the interaction occurs:

$$\begin{aligned} \hat{y}_{ub} &= e_u^{(l)} \odot e_b^{(l)}, \\ \hat{y}_{ui} &= e_u^{(l)} \odot e_i^{(l)}. \end{aligned} \tag{10}$$

Structural Hint Learning. To enable the graph neural network to effectively capture the topological structure of the user-bundle graph G_{ub}, we propose a structure loss $\mathcal{L}_{\text{structure}}$, which comprises two components: the node degree prediction loss $\mathcal{L}_{\text{degree}}$ and the neighbor degree sum prediction loss $\mathcal{L}_{\text{neigh}}$. Specifically, $\mathcal{L}_{\text{degree}}$ employs the mean squared error (MSE) to supervise the model in predicting the degree $\hat{d}_i$ of each node, ensuring alignment with its ground-truth degree d_i. Meanwhile, $\mathcal{L}_{\text{neigh}}$ focuses on predicting the sum of the degrees of each node's neighbors $\hat{s}_i$, measuring its deviation from the true value s_i. Ultimately, the structure loss is defined as a weighted combination of these two terms, balancing the representation of local and neighborhood structural properties. The node degree prediction loss is formulated as:

$$\mathcal{L}_{\text{degree}} = \frac{1}{N_u + N_b} \sum_{i=1}^{N_u+N_b} \left(\hat{d}_i - d_i\right)^2 \tag{11}$$

where N_u and N_b denote the number of user nodes and bundle nodes in the user-bundle graph G_{ub}, $\hat{d}_i$ denotes the predicted degree of node i, and d_i denotes its ground-truth degree. α and β are weight coefficients controlling the contributions of $\mathcal{L}_{\text{degree}}$ and $\mathcal{L}_{\text{neigh}}$ to the total structure loss $\mathcal{L}_{\text{structure}}$. Similarly, the neighbor degree sum prediction loss is given by:

$$\mathcal{L}_{\text{neigh}} = \frac{1}{N_u + N_b} \sum_{i=1}^{N_u+N_b} \left(\hat{s}_i - s_i\right)^2 \tag{12}$$

where $\hat{s}_i$ denotes the predicted sum of the degrees of the neighbors of node i, and s_i denotes its ground-truth sum. The total structure loss combines these two terms:

$$\mathcal{L}_{\text{structure}} = \alpha\mathcal{L}_{\text{degree}} + \beta\mathcal{L}_{\text{neigh}} \tag{13}$$

where α and β are weight coefficients controlling the contributions of $\mathcal{L}_{\text{degree}}$ and $\mathcal{L}_{\text{neigh}}$ to the total structure loss $\mathcal{L}_{\text{structure}}$.

Optimization. As mentioned earlier, the user's preference for a bundle is closely tied to the items it contains. Building upon this idea, we employ multi-task learning in the optimization process to enhance the recommendation capability of our MGBR model. Specifically, we construct two triples: one comprises of a user u_i, an observed item v_j, and an unobserved item v_s, while the other consists of a user u_i, an observed bundle b_k, and an unobserved item b_t, formally as:

$$\mathcal{T}_1 = \{(u_i, v_j, v_s)|R_{i,j} = 1, R_{i,s} = 0\} \tag{14}$$

$$\mathcal{T}_2 = \{(u_i, b_k, b_t)|R_{i,k} = 1, R_{i,t} = 0\} \tag{15}$$

where $R_{i,j} = 1$ denotes that u_i has interacted with v_j; otherwise, $R_{i,j} = 0$. Then, we adopt the Bayesian Personalized Ranking (BPR) loss as the loss function [11].

$$\mathcal{L}_{item} = \sum_{(u_i,v_j,v_s)\in\mathcal{T}_1} -ln\sigma(\hat{y}_{ij} - \hat{y}_{is}) \tag{16}$$

$$\mathcal{L}_{bundle} = \sum_{(u_i,b_k,b_t)\in\mathcal{T}_2} -ln\sigma(\hat{y}_{ik} - \hat{y}_{it}) \tag{17}$$

Combining with the item loss and bundle loss and structure loss, we reach the objective function as follows:

$$\mathcal{L} = \lambda_1\mathcal{L}_{item} + \lambda_2\mathcal{L}_{bundle} + \lambda_3\mathcal{L}_{structure} + \lambda_4 \left\|\Theta\right\|_2, \tag{18}$$

where λ , Θ represent the regularization weights and the parameters of the model, respectively.

4 Experiments

In this section, we conduct comprehensive experiments on two real-world public datasets to evaluate our proposed method for bundle recommendation.

4.1 Experimental Settings

Datasets. We evaluate all the models on two public datasets, NetEase[4] and Youshu[5]. The statistics of the datasets are summarized in Table 1.

[4] https://music.163.com
[5] http://www.yousuu.com

NetEase: This dataset is obtained from NetEase Cloud Music [4], which is a music app with a song list as the core structure, the largest and best quality song list library in China. In this dataset, songs are items and song lists are bundles.

Youshu: This dataset is collected from Youshu [7], which is a Chinese book review site. Similarly to Netease, every bundle is a list of books, and the items are books.

Evaluation Metrics. For each user in the training and testing sets, we treat bundles and items with which the user has not interacted as negative samples. The interactions between user and bundle pairs are scored by the trained model and ranked in descending order. In addition to the Recall@K, we also adopted normalized discounted cumulative gain (NDCG@K for short) to evaluate the effectiveness of the top-K bundle recommendation.

Table 1: The statistics of the datasets.

Dataset	NetEase	Youshu
#Users	18,528	8,039
#Bundles	22,864	4,771
#Items	123,628	32,770
#U-B interactions	302,303	51,377
#U-I interactions	1,128,065	138,515
#B-I affiliations	1,778,838	176,667

Baselines. We compared our model with the following representative methods:

MFBPR [11]: BPR is a pairwise ranking algorithm based on implicit feedback. This model is a standard baseline for top-N recommendation tasks by taking advantage of BPR loss.

DAM [7]: This is a deep attention-based multi-task model that uses the representation of items in a factorized attention network aggregation bundle.

BGCN [8]: This is the first to propose a graph neural network model to solve the problem of bundle recommendation.

RNBR [10]: This model considers the interactions of neighbors and uses a relational graph neural network to inject different relations into the representation of bundles and items.

MIDGN [12]: it formulates user preference into local and global views and leverages the intent disentanglement to learn user's intents for bundle recommendation.

LightGCL [13]: this method utilizes a novel framework of SVD to generate augmented views for self-supervised contrastive learning.

Table 2: Performance comparison of MGBR and the baselines.

Model	NetEase						Youshu					
	Recall@20	NDCG@20	Recall@40	NDCG@40	Recall@80	NDCG@80	Recall@20	NDCG@20	Recall@40	NDCG@40	Recall@80	NDCG@80
DAM	0.0411	0.0210	0.0690	0.0281	0.1090	0.0372	0.2082	0.1198	0.2890	0.1418	0.3915	0.1658
BGCN	0.0491	0.0258	0.0829	0.0346	0.1304	0.0453	0.2347	0.1345	0.3248	0.1593	0.4355	0.1851
RNBR	0.0644	0.0327	0.1023	0.0431	0.1602	0.0551	0.2732	0.1571	0.3814	0.1857	0.4937	0.2116
MIDGN	0.0678	0.0343	0.1085	0.0451	0.1623	0.0564	0.2682	0.1527	0.3712	0.1808	0.5024	0.2126
LightGCL	0.0705	0.0371	0.1122	0.0481	0.1712	0.0594	0.2710	0.1556	0.3693	0.1827	0.4995	0.2178
CrossCBR	0.0844	0.0458	0.1255	0.0567	0.1794	0.0679	0.2842	0.1670	0.3787	0.1939	0.5002	0.2208
MGBR	**0.0855**	**0.0471**	**0.1288**	**0.0584**	**0.1821**	**0.0696**	**0.2900**	**0.1705**	**0.3963**	**0.1998**	**0.5134**	**0.2253**
%Improve	**1.30%**	**2.84%**	**2.63%**	**3.00%**	**1.51%**	**2.50%**	**2.04%**	**2.10%**	**4.65%**	**3.04%**	**2.64%**	**2.04%**

CrossCBR [14]: it utilizes a cross-view contrastive learning to achieve cross-view cooperative association based on the BGCN view construction method.

Parameter Settings. Our model is implemented in Pytorch with DGL[6] package. We initialized the model parameter by the Xavier [16] initializer and took the Adam [17] as the optimizer. The dimension of the embedding vector is set to 64. The batch size is set to 2048. In terms of the hyperparameters, the learning rate is searched in {0.0001,0.0003,0.001,0.01} and regularization weight is tuned in {0.0001, 0.001}.

4.2 Performance Comparison

Table 2 shows the experimental results of all methods and the improvements, which are calculated between our proposed method and the basic baselines. The optimal results are marked in bold and the underline values are the suboptimal results. We can observe that our model outperforms all baselines in terms of the Recall@K and NDCG@k metric. Specifically, on the Recall@k metric, MGBR outperforms the best baseline by 1.30% - 2.63% on the NetEase dataset and 2.04% - 4.65% on the Youshu dataset. As for the NDCG@K metric, the improvements over the best baseline by 2.84%, 3.00%, and 2.50% for NDCG@20, NDCG@40, and NDCG@80 on the NetEase dataset and improve by 2.10%, 3.04%, and 2.04% respectively on the Youshu dataset.

The improvement mainly comes from the following reasons:

1) Instead of building a unified heterogeneous graph, we construct three bipartite graphs based on multi-interaction. This approach avoids redundancy in learning and hinders preference modeling, allowing for a more effective representation of the various complex relationships among users, items, and bundles.

2) Based on a graph attention network, MGBR effectively leverages the higher-order relationships between different entities. The attention mechanism in MGBR allows for the learning of influence values for different interactions, recognizing that not all interactions have the same impact on users.

3) The incorporation of two graph structure losses, namely the node degree prediction loss and the neighbor degree sum prediction loss , enhances the model's

[6] https://www.dgl.ai.

Table 3: Performance comparison of Structural Hint Learning on MGBR.

Model	NetEase				Youshu			
	Recall@20	NDCG@20	Recall@40	NDCG@40	Recall@20	NDCG@20	Recall@40	NDCG@40
w/o Both	0.0794	0.0395	0.1187	0.0483	0.2784	0.1583	0.3712	0.1842
w/o $\mathcal{L}_{\mathbf{neigh}}$	0.0834	0.0434	0.1245	0.0553	0.2871	0.1667	0.3851	0.1926
w/o $\mathcal{L}_{\mathbf{degree}}$	0.0821	0.0426	0.1219	0.0523	0.2848	0.1627	0.3795	0.1874
MGBR	**0.0855**	**0.0471**	**0.1288**	**0.0584**	**0.2900**	**0.1705**	**0.3963**	**0.1998**

capability to preserve the topological properties of the graph. These losses ensure that the learned representations align with both local node connectivity and broader neighborhood structures, further improving the robustness and accuracy of the model.

4.3 Ablation Study

Effects of graph structural loss. To assess the contributions of the structural losses in the user-bundle graph G_{ub}, an ablation study is conducted by comparing the full model, which incorporates both the node degree prediction loss $\mathcal{L}_{\text{degree}}$ and the neighbor degree sum prediction loss $\mathcal{L}_{\text{neigh}}$, against variants where each loss is individually removed—denoted as w/o $\mathcal{L}_{\text{degree}}$ and w/o $\mathcal{L}_{\text{neigh}}$—as well as a baseline with both losses excluded w/o Both. Table 3 results reveal that removing $\mathcal{L}_{\text{degree}}$ leads to a more pronounced performance decline compared to the removal of $\mathcal{L}_{\text{neigh}}$,When $\mathcal{L}_{\text{degree}}$ is removed, the model loses its explicit supervision on predicting the local connectivity of nodes in G_{ub}. As $\mathcal{L}_{\text{degree}}$ ensures that the embeddings reflect the ground-truth degrees of users and bundles (e.g., user activity or bundle popularity), its absence may weaken the model's ability to distinguish highly active users or popular bundles.

Effects of graph attention mechanism. In this part, we verified the validity of the graph attention mechanism in our model. We compared our model with the model based on graph convolution network (GCN) [18]. Figure.3 shows the results of the comparison, where GCN-HG is applied to the three homogeneous graphs with GCN. From the results, we can observe that the model MGBR based on the graph attention mechanism is superior to the GCN-based model in all metrics. The results show that the model performs better than GCN-HG by 6.51% - 34.43% and 5.15% - 7.34% on the NetEase dataset and Youshu dataset, respectively. The graph convolution network aggregates information from connected neighbors equally. In contrast, our model better aggregates information by calculating the attention values between neighbors. This demonstrates that the graph attention mechanism is essential to capture the varying degrees of influence that different interactions may have on users' preferences.

Effects of multi-task learning. We compare the performance of MGBR under training with no item loss, and item loss. As shown in Fig. 4, the model added item-level loss performs better than only trained with bundle-level loss by 2.44% - 21.93% and 6.03% - 9.21% on the Netease dataset and Youshu dataset, respectively.

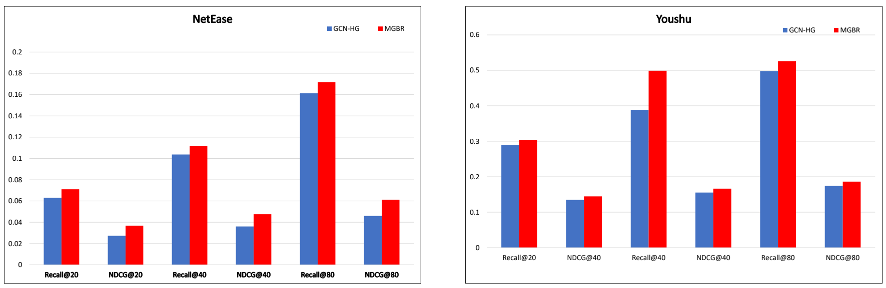

(a) (b)

Fig. 3: Performance with different graph neural network.

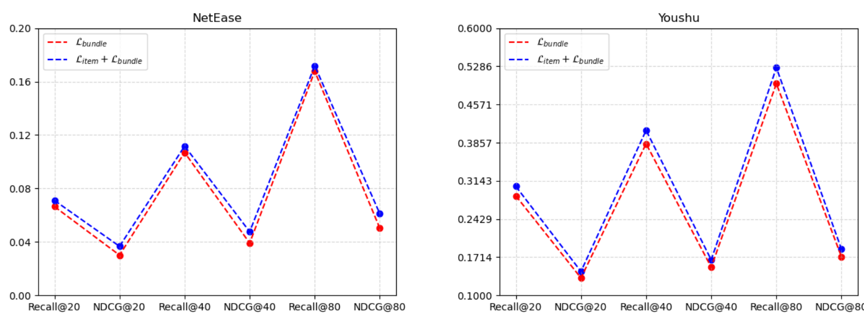

(a) (b)

Fig. 4: Performance with different loss functions.

5 Related Work

Product bundling is a widely-utilized marketing tactic employed by brick-and-mortar retailers as well as e-commerce platforms[19]. It typically involves grouping together a curated collection of related products that users consume as a whole under circumstances, e.g., limited total price[20, 21], or specific intents[22, 23].

Bundle recommendation focuses on capturing the relationship between users, items and bundles, and recommending bundles to users. LIRE [6] automatically recommends a list of relevant items for each user based on the user's history of interaction with lists and items. Cao *et al.* [4] propose a joint learning framework for the item and bundle recommendation using embedding factorization models. Chen *et al.* [7] propose a factorized attention multi-task model. Empowered by the success of the graph neural networks in the node representing, it is widely used in bundle recommendation. BGCN [8] constructs the two kinds of interaction and affiliation into the graph and utilizes the GCN to learn the representation of user and bundle. IHBR [9] considers the co-purchase and co-occurrence information within items for modeling intention-oriented hierarchical representations. Wang *et al.* [10] present a relational graph neural network with neighbor interactions for bundle recommendation.Recent research has significantly advanced bundle recommendation by incorporating contrastive learning techniques. MIDGN [12] partitions user-bundle preferences into local and global views and employs an intent decomposition strategy to model multiple latent intents, with its dual-view contrastive learning approach effectively enhancing recommendation performance. CrossCBR [14], built upon the BGCN [8] framework, innovatively constructs bundle and item views, leveraging cross-view contrastive learning to improve representation affinity among similar users while demonstrating that: (1) multi-view formulations are critical for capturing hierarchical user preferences, and (2) cross-view contrastive mechanisms, by explicitly modeling collaborative associations, substantially boost performance. Building on this, Ma et al. [15] enhance the multi-view architecture and propose MultiCBR, an optimized contrastive learning framework tailored for multi-view scenarios. By adopting an "early fusion, late contrast" strategy, MultiCBR effectively captures user preferences and addresses the sparsity issue in bundle-item associations.

6 Conclusions

We propose a graph neural network-based model called MGBR for bundle recommendation. To avoid information overlap, we differentiate the representation of interactions and affiliations between users, items, and bundles. MGBR utilizes an attention mechanism with residual networks to learn representations from a homogeneous graph, considering the dependencies among different entities. Moreover, we leverage user-item and user-bundle information simultaneously to enhance the performance of bundle recommendations. Meanwhile, we design dual structural loss components to jointly optimize node degree distribution and neighborhood topology in the user-bundle graph. Our experimental results on two real-world datasets demonstrate the effectiveness of MGBR in providing accurate and personalized bundle recommendations. In the future, we will explore user intents among the multiple interactions and further consider how to generate high-quality bundles.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 71971002), Natural Science Research Project of Anhui Educational Committee (Grant No. 2024AH040011), and the Anhui Postdoctoral Scientific Research Program Foundation (Grant No. 2024B828).

References

1. Zhang, Shuai and Yao, Lina and Sun, Aixin and Tay, Yi.: Deep learning based recommender system: A survey and new perspectives. ACM computing surveys (CSUR), 1–38 (2019)
2. Wu, Shiwen and Sun, Fei and Zhang, Wentao and Xie, Xu and Cui, Bin.: Graph neural networks in recommender systems: a survey. ACM computing surveys, pp. 1–37. ACM New York, NY (2022)
3. Zhou, Jie and Cui, Ganqu and Hu, Shengding and Zhang, Zhengyan and Yang, Cheng and Liu, Zhiyuan and Wang, Lifeng and Li, Changcheng and Sun, Maosong.: Graph neural networks: A review of methods and applications. AI open, pp. 57–81. Elsevier (2022)
4. Cao, Da and Nie, Liqiang and He, Xiangnan and Wei, Xiaochi and Zhu, Shunzhi and Chua, Tat-Seng.: Embedding factorization models for jointly recommending items and user generated lists. Proceedings of the 40th international ACM SIGIR conference on research and development in information retrieval, 99–110 (2017)
5. Pathak, Apurva and Gupta, Kshitiz and McAuley, Julian.: Generating and personalizing bundle recommendations on steam. Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, 1073–1076 (2017)
6. Liu, Yidan and Xie, Min and Lakshmanan, Laks VS.: Recommending user generated item lists. Proceedings of the 8th ACM Conference on Recommender systems, 185–192 (2014)
7. Chen, Liang and Liu, Yang and He, Xiangnan and Gao, Lianli and Zheng, Zibin.: Matching User with Item Set: Collaborative Bundle Recommendation with Deep Attention Network. In: IJCAI, pp. 2095–2101. (2019)
8. Chang, Jianxin and Gao, Chen and He, Xiangnan and Jin, Depeng and Li, Yong.: Bundle recommendation with graph convolutional networks. Proceedings of the 43rd international ACM SIGIR conference on Research and development in Information Retrieval, 1673–1676 (2020)
9. Tan, Meng and Chen, Wei and Wang, Weiqing and Liu, An and Zhao, Lei.: Intention-oriented Hierarchical Bundle Recommendation with Preference Transfer. 2021 IEEE International Conference on Web Services (ICWS), 107–116 (2021)
10. Wang, Xin and Liu, Xiao and Liu, Jin and Wu, Hao.: Relational graph neural network with neighbor interactions for bundle recommendation service. 2021 IEEE International Conference on Web Services (ICWS), 167–172 (2021)
11. Rendle, Steffen and Freudenthaler, Christoph and Gantner, Zeno and Schmidt-Thieme, Lars.: BPR: Bayesian personalized ranking from implicit feedback. Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, 452–461 (2009)
12. Sen Zhao, Wei Wei, Ding Zou, and Xianling Mao. 2022. Multi-view intent disentangle graph networks for bundle recommendation. In AAAI. AAAI Press, 4379–4387

13. Xuheng Cai, Chao Huang, Lianghao Xia, and Xubin Ren. 2023. LightGCL: Simple yet effective graph contrastive learning for recommendation. CoRR abs/2302.08191 (2023)
14. Yunshan Ma, Yingzhi He, An Zhang, Xiang Wang, and Tat-Seng Chua. 2022. CrossCBR: Cross-view contrastive learning for bundle recommendation. In KDD. ACM, 1233–1241
15. Ma, Y., He, Y., Wang, X., Wei, Y., Du, X., Fu, Y., Chua, T. S. (2024). MultiCBR: Multi-view Contrastive Learning for Bundle Recommendation. ACM Transactions on Information Systems, 42(4), 1–23
16. Glorot, Xavier and Bengio, Yoshua.: Understanding the difficulty of training deep feedforward neural networks. Proceedings of the thirteenth international conference on artificial intelligence and statistics, 249–256 (2010)
17. Kingma, Diederik P and Ba, Jimmy.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, (2014)
18. Kipf, Thomas N and Welling, Max.: Semi-supervised classification with graph convolutional networks. International Conference on Learning Representations(ICLR), (2017)
19. Sun, Zhu and Yang, Jie and Feng, Kaidong and Fang, Hui and Qu, Xinghua and Ong, Yew Soon.: Revisiting Bundle Recommendation: Datasets, Tasks, Challenges and Opportunities for Intent-aware Product Bundling. Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, 2900–2911 (2022)
20. Robert Garfinkel, Ram Gopal, Arvind Tripathi, and Fang Yin. 2006. Design of a Shopbot and Recommender System for Bundle Purchases. DSS 42(3) pp.1974–1986 (2006).
21. Yong Ge, Hui Xiong, Alexander Tuzhilin, and Qi Liu. 2014. Cost-aware Collaborative Filtering for Travel Tour Recommendations. TOIS 32, 1 (2014), 1–31.
22. Jinze Bai, Chang Zhou, Junshuai Song, Xiaoru Qu, Weiting An, Zhao Li, and Jun Gao. 2019. Personalized Bundle List Recommendation. In WWW. 60–71.
23. Yun He, Yin Zhang, Weiwen Liu, and James Caverlee. 2020. Consistency-Aware Recommendation for User-Generated Item List Continuation. In WSDM. 250–258.
24. Ma, Yunshan and He, Yingzhi and Zhang, An and Wang, Xiang and Chua, Tat-Seng.: CrossCBR: Cross-view Contrastive Learning for Bundle Recommendation. Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 1233–1241 (2022).
25. Zhu, Shixuan and Shen, Qi and Zhang, Yiming and Pang, Yitong and Wei, Zhihua. 2022. Data-Augmented Counterfactual Learning for Bundle Recommendation. arXiv preprint arXiv:2210.10555.