

CrossCBR: Cross-view Contrastive Learning for Bundle Recommendation

Yunshan Ma
Sea-NExT Joint Lab & National
University of Singapore
yunshan.ma@u.nus.edu

Yingzhi He
Sea-NExT Joint Lab & National
University of Singapore
heyingzhi@u.nus.edu

An Zhang*
Sea-NExT Joint Lab & National
University of Singapore
an_zhang@nus.edu.sg

Xiang Wang
University of Science and Technology
of China
xiangwang1223@gmail.com

Tat-Seng Chua
Sea-NExT Joint Lab & National
University of Singapore
dcscs@nus.edu.sg

ABSTRACT

Bundle recommendation aims to recommend a bundle of related items to users, which can satisfy the users' various needs with one-stop convenience. **Recent methods usually take advantage of both user-bundle and user-item interactions information to obtain informative representations for users and bundles, corresponding to bundle view and item view, respectively.** However, they either use a unified view without differentiation or loosely combine the predictions of two separate views, while the crucial cooperative association between the two views' representations is overlooked.

In this work, we propose to model the cooperative association between the two different views through cross-view contrastive learning. By encouraging the alignment of the two separately learned views, **each view can distill complementary information from the other view, achieving mutual enhancement.** Moreover, by enlarging the dispersion of different users/bundles, the self-discrimination of representations is enhanced. Extensive experiments on three public datasets demonstrate that our method outperforms SOTA baselines by a large margin. Meanwhile, our method requires minimal parameters of three set of embeddings (user, bundle, and item) and the computational costs are largely reduced due to more concise graph structure and graph learning module. In addition, various ablation and model studies demystify the working mechanism and justify our hypothesis. Codes and datasets are available at <https://github.com/mysbupt/CrossCBR>.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Bundle Recommendation, Contrastive Learning, Graph Neural Network

*Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

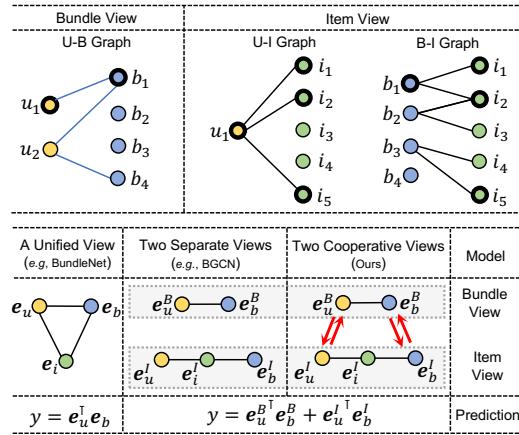


Figure 1: Top: The bundle and item views presented in the U-B, U-I and B-I graphs. Bottom: Our work models the cooperative association between views, where the superscripts B and I denote the bundle and item view, and the subscripts u , b , and i stand for the user, bundle, and item.

ACM Reference Format:

Yunshan Ma, Yingzhi He, An Zhang, Xiang Wang, and Tat-Seng Chua. 2022. CrossCBR: Cross-view Contrastive Learning for Bundle Recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22)*, August 14–18, 2022, Washington, DC, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3534678.3539229>

1 INTRODUCTION

Bundle recommendation aims to recommend a set of items related with the same theme to users. In a variety of online applications, such as music platforms and fashion shopping sites, serving bundles instead of individual items can boost the users' experience in a one-stop manner. More importantly, platforms taking bundles as the marketing strategy can increase sales revenue and attract customers fond of bundle discounts. Consequently, both the users and platforms would prefer bundles (e.g., music playlist and fashion outfit) instead of single items (e.g., single song and piece-of-clothing). Therefore, developing effective bundle recommender systems is attracting a surge of interest in both academia and industry.

Scrutinizing prior studies on bundle recommendation [2, 3, 5, 9, 26], we can systematize the sources of user preferences as two views:

(1) **bundle view**, which depicts user preferences through the **user-bundle** interactions and can be reorganized as an U-B graph; and (2) **item view**, which delineates user behaviors and bundle knowledge at the granularity of items — *i.e.*, **user-item** interactions in the form of an U-I graph and **bundle-item** affiliation in the form of a B-I graph, respectively. These two views allow us to understand user interests and construct the recommender models from different perspectives. **However, there exist clear discrepancies between these two views which have not been modeled in prior studies.** Consider the running example on the top part of Figure 1, where u_1 is the target user who has interacted with bundle b_1 and individual items i_1 , i_2 , and i_5 before, and the historical bundle b_1 consists of items i_1 and i_2 . By taking the bundle view solely, a recommender model is highly likely to route b_4 to u_1 , if the behavioral similarity between users u_1 and u_2 is well captured. On the other hand, by taking the item view, a recommender model is prone to yield b_2 and b_3 as the interested item of u_1 , since b_2 holds items (*i.e.*, i_2) shared with the historical bundle b_1 and b_3 contains items (*i.e.*, i_5) individually preferred by u_1 . Clearly, the bundle view emphasizes the behavioral similarity among users, **while the item view highlights the content relatedness among bundles and users' item-level preference.** Hence they are complementary but different, and the cooperation of these two views is the key to accurate bundle recommendation.

However, the cooperative association between these two views has been loosely modeled or even overlooked in existing works [2, 3, 5, 9]. One research line like BundleNet [9] blindly merges the two views into a unified tripartite graph and employs graph neural networks (GNNs) to aggregate the neighboring information into representations of users and bundles. However, such representations fail to differentiate the behavioral similarity among users and content relatedness among bundles from these two views, thus obscuring their cooperation. Another line, such as BGCN [3], first performs representation learning and preference prediction upon the views individually, and then fuses these two view-specific predictions. While this loose modeling of the two views performs better than the first line, it only considers the cooperative signal at the level of predictions, rather than directly plugging such signal into the representations optimized for recommendation. Hence, no mutual enhancement of the two views is guaranteed to be captured. Considering the limitations of the two lines, we believe that it is crucial to properly model the cooperative association and encourage the mutual enhancement across the bundle and item views.

Towards this end, we propose a **Cross-view Contrastive Learning for Bundle Recommendation (CrossCBR)** which captures the cooperative association by cross-view contrastive learning and mutually enhances the view-aware representations. The basic idea is to treat the bundle and item views as two distinct but correlated viewpoints of user-bundle preferences, and apply contrastive learning on these viewpoints to encapsulate their agreements into representations of users and bundles. Specifically, upon the U-B graph, we build a LightGCN [14] as the backbone to obtain the bundle-view representations of users and bundles; analogously, upon the U-I graph, we employ another LightGCN to generate the item-view representations of users and items, and aggregate the representations of compositional items as the bundle representation based on the B-I graph. We jointly employ the BPR [25] and contrastive loss [13]

to optimize these representations. Benefiting from the cross-view contrastive learning, CrossCBR outperforms the state-of-the-art (SOTA) baselines by a large margin on three datasets.

To demystify the working mechanism behind CrossCBR, we further investigate the alignment-dispersion characteristics of the learned representations. Encouraging the cross-view alignment enables the view-aware representations to learn from each other and achieve mutual enhancement; meanwhile, enlarging the cross-view dispersion between different users/bundles is excel at enhancing the representations' discriminative power. Such a powerful representation learning comes with minimal space complexity and low time complexity. Our main contributions are as follows:

- To the best of our knowledge, we are among the first to formulate the cross-view cooperative association in bundle recommendation, providing a new research line worthy of further exploration.
- We propose a simple yet effective bundle recommender, CrossCBR, to model the cooperative association between two views via cross-view contrastive learning.
- Our model outperforms SOTA baselines by a large margin on three public datasets, while requires largely reduced training time. We also demonstrate how the idea of CrossCBR can be generalized to a broader scope of tasks.

2 METHODOLOGY

In this section, we first formulate the task of bundle recommendation and the present our CrossCBR, as shown in Figure 2. The in-depth discussion of the working mechanism and analysis of computational complexity for CrossCBR are followed.

2.1 Problem Formulation

Given a set of users $\mathcal{U} = \{u_1, u_2, \dots, u_M\}$, a set of bundles $\mathcal{B} = \{b_1, b_2, \dots, b_N\}$, and a set of items $\mathcal{I} = \{i_1, i_2, \dots, i_O\}$, where M , N , and O are the number of users, bundles, and items, respectively. The user-bundle interactions, user-item interactions, and bundle-item affiliations are denoted as $X_{M \times N} = \{x_{ub} | u \in \mathcal{U}, b \in \mathcal{B}\}$, $Y_{M \times O} = \{y_{ui} | u \in \mathcal{U}, i \in \mathcal{I}\}$, and $Z_{N \times O} = \{z_{bi} | b \in \mathcal{B}, i \in \mathcal{I}\}$, respectively. $x_{ub}, y_{ui}, z_{bi} \in \{0, 1\}$, where 1 represents an interaction between the user-bundle or user-item pair, or the item belongs to a certain bundle. Note that since we deduplicate the historical bundle and item interactions for each user, each element of X and Y is a binary value rather than an integer. In addition, X and Y are separately generated, where users are allowed to directly interact with both bundles and individual items. Therefore, X and Y contain different information, which heuristically enables the cooperative effect between the two different views. The goal of bundle recommendation task is to learn a model from the historical $\{X, Y, O\}$ and predict the unseen user-bundle interactions in X .

2.2 Learning of Two Views' Representations

For the first component of CrossCBR, we aim to learn the representations from the two views: bundle and item view. Despite the effectiveness of two views' representation learning module of BGCN [3], its partial designs of graph construction and graph learning are useless or even harmful [14], especially under the circumstance of utilizing the contrastive learning. Here we devise our simpler yet more effective representation learning approach.




Figure 2: The overall framework of CrossCBR consists of two parts: (1) representation learning for the two views of users and bundles and (2) the joint optimization of the BPR loss \mathcal{L}^{BPR} and contrastive loss \mathcal{L}^C .

2.2.1 Bundle-view Representation Learning. In order to learn the user and bundle representations from the bundle view, we first construct a user-bundle bipartite graph, *i.e.*, U-B graph, based on the user-bundle interaction matrix X . We then employ the prevailing GNN-based recommendation framework LightGCN [14] to learn the representations of both user and bundle. Specifically, we conduct information propagation over the U-B graph, and the k -th layer’s information propagation is denoted as:

$$\begin{cases} e_u^{B(k)} = \sum_{b \in \mathcal{N}_u} \frac{1}{\sqrt{|\mathcal{N}_u|} \sqrt{|\mathcal{N}_b|}} e_b^{B(k-1)}, \\ e_b^{B(k)} = \sum_{u \in \mathcal{N}_b} \frac{1}{\sqrt{|\mathcal{N}_b|} \sqrt{|\mathcal{N}_u|}} e_u^{B(k-1)}, \end{cases} \quad (1)$$

where $e_u^{B(k)}, e_b^{B(k)} \in \mathbb{R}^d$ are the k -th layer’s information propagated to user u and bundle b ; d is the embedding dimensionality; the superscript B indicates the bundle view; $e_u^{B(0)}$ and $e_b^{B(0)}$ are randomly initialized at the beginning of the training; \mathcal{N}_u and \mathcal{N}_b are the neighbors of the user u and bundle b in the U-B graph.

We follow LightGCN to remove the self-connections from the U-B graph and the nonlinear transformation from the information propagation function. We will empirically demonstrate that such simplifications, which BGCN does not take into account, are truly helpful for better performance (*cf.* Section 3.3). More importantly, we do not incorporate the bundle-bundle connections, which are introduced by BGCN and calculated from the degree of overlapped items between the two bundles. The reason lies in the fact that bundle-bundle overlap information can be distilled from the item view through the cross-view contrastive learning (*cf.* Section 2.3.2). Meanwhile, the removal of extra bundle-bundle connections can further reduce the computational costs during the graph learning.

We concatenate all K layers’ embedding to combine the information received from neighbors of different depths. The final bundle-view representations e_u^{B*} and e_b^{B*} are denoted as:

$$e_u^{B*} = \sum_{k=0}^K e_u^{B(k)}, \quad e_b^{B*} = \sum_{k=0}^K e_b^{B(k)}. \quad (2)$$

2.2.2 Item-view Representation Learning. In order to learn the user and bundle representations from the item view, we first build two bipartite graphs, *i.e.*, U-I and B-I graph, according to the user-item interactions Y and bundle-item affiliations Z , respectively. Similar to the U-B graph learning, we learn user and item representations using LightGCN. The obtained user representations are the item-view user representations, and the item-view bundle representations are obtained by performing average pooling over the item-view item representations guided by the B-I graph. Specifically, the information propagation over the U-I graph is defined as:

$$\begin{cases} e_u^{I(k)} = \sum_{i \in \mathcal{N}_u} \frac{1}{\sqrt{|\mathcal{N}_u|} \sqrt{|\mathcal{N}_i|}} e_i^{I(k-1)}, \\ e_i^{I(k)} = \sum_{u \in \mathcal{N}_i} \frac{1}{\sqrt{|\mathcal{N}_i|} \sqrt{|\mathcal{N}_u|}} e_u^{I(k-1)}, \end{cases} \quad (3)$$

where $e_u^{I(k)}, e_i^{I(k)} \in \mathbb{R}^d$ are the k -th layer’s information propagated to user u and item i , respectively; the superscript I refers to the item view; $e_i^{I(0)}$ is randomly initialized; \mathcal{N}_u and \mathcal{N}_i are the neighbors of the user u and item i in the U-I graph. We follow BGCN and share the parameters of $e_u^{I(0)}$ with $e_u^{B(0)}$, which empirically does not affect the performance but largely reduces the number of parameters. Meanwhile, such initial layer’s parameters sharing between two views is too weak even impossible to capture the cross-view cooperative association (*cf.* CrossCBR-CL in Section 3.3). Similar to U-B graph, we also remove the self-connections from the U-I graph and nonlinear feature transformation from the information propagation function. And a layer aggregation operation is adopted after K layers of information propagation, formulated as follows:

$$e_u^{I*} = \sum_{k=0}^K e_u^{I(k)}, \quad e_i^{I*} = \sum_{k=0}^K e_i^{I(k)}, \quad (4)$$

where e_u^{I*} and e_i^{I*} are the item-view user and item representations, respectively. Based on the item-view item representation and the B-I graph, we can obtain the item-view bundle representations e_b^{I*}

through average pooling, denoted as:

$$e_b^{I*} = \frac{1}{|\mathcal{N}_b|} \sum_{i \in \mathcal{N}_b} e_i^{I*}, \quad (5)$$

where \mathcal{N}_b represents the set of items a certain bundle b contains.

In summary, we can learn the representations of all users and bundles from two views, denoted as $E_U^{B*}, E_U^{I*} \in \mathbb{R}^{M \times d}$ and $E_B^{B*}, E_B^{I*} \in \mathbb{R}^{N \times d}$, where the superscripts B and I stand for the bundle and item view, respectively; and the subscripts U and B indicate the whole user and bundle set, respectively ($E_I^{I*} \in \mathbb{R}^{O \times d}$ are the representations of all items in the item view). Thereafter, given a user u and a bundle b , we can obtain their bundle-view representations, i.e., e_u^{B*} and e_b^{B*} , and their item-view representations, i.e., e_u^{I*} and e_b^{I*} .

2.3 Cross-view Contrastive Learning

We devise the critical component to model the cross-view cooperative association via contrastive learning. We first present the data augmentation methods, followed by the contrastive loss.

2.3.1 Data Augmentation. The main idea of self-supervised contrastive learning is to encourage the representation affinity among various views of the same object, while at the same time enlarge the representation dispersion of different objects [27]. In practice, if multiple views naturally exist for each object, e.g., images taken from different angles, or the bundle and item view in bundle recommendation, the contrastive loss can be directly applied. On the other hand, in many scenarios, multiple views are not available, **and data augmentation is leveraged to generate multiple views from the original data [6, 11, 30]**. Proper data augmentation not only release the (multi-view) data constraint for applying contrastive learning, but also may improve the robustness to counter potential noise. Therefore, while keeping the original preservation (no augmentation) as the default setting, we also introduce two simple data augmentation methods: graph- and embedding-based augmentations.

Graph-based Augmentation. The main purpose of graph-based augmentation is to generate augmented data by revising the graph structure [30]. We adopt a simple random augmentation method of edge dropout (ED), which randomly removes a certain proportion (dropout ratio ρ) of edges from the original graph. The rationale behind edge dropout lies in that the core local structure of the graph is preserved. Therefore, the robustness of learned representations may be enhanced to counter certain noise.

Embedding-based Augmentation. Different from the graph-based augmentation, which can be applied only to graph data, embedding-based augmentations are more general and suitable for any deep representation learning based methods [11]. The major idea is to vary the learned representation embeddings regardless of how the embeddings are obtained. We employ message dropout (MD), which randomly masks some elements of the propagated embeddings with a certain dropout ratio ρ during the graph learning.

Original Preservation. We name the approach without any data augmentation as original preservation (OP), where no randomness is introduced and only the original representations are preserved. Since the two views in bundle recommendation are obtained from different sources of data, their representations are distinctive sufficiently to work well.

To avoid the abuse of notations, after the data augmentation, we still use the same notations of $e_u^{B*}, e_b^{B*}, e_u^{I*}, e_b^{I*}$ to denote the embeddings for the bundle-view user, bundle-view bundle, item-view user, and item-view bundle, respectively.

2.3.2 Cross-view Contrastive Loss. We leverage the cross-view contrastive loss to optimize two-view representations. As the motivations illustrate in Figure 1, each view captures a distinctive aspect of user's preference, and the two views have to work cooperatively to maximize the overall modeling capacity. To model the cross-view cooperative association, we employ the cross-view contrastive loss (we leave other potential modeling solutions for future work). We adopt the popular InfoNCE [13] loss built upon the cross-view representations of users and bundles, respectively. More precisely, the contrastive loss is able to simultaneously encourage the alignment of the same user/bundle from different views and enforce the separation of different users/bundles. The equations are as follows:

$$\mathcal{L}_U^C = \frac{1}{|\mathcal{U}|} \sum_{u \in \mathcal{U}} -\log \frac{\exp(s(e_u^{B*}, e_u^{I*})/\tau)}{\sum_{v \in \mathcal{U}} \exp(s(e_u^{B*}, e_v^{I*})/\tau)}, \quad (6)$$

$$\mathcal{L}_B^C = \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} -\log \frac{\exp(s(e_b^{B*}, e_b^{I*})/\tau)}{\sum_{p \in \mathcal{B}} \exp(s(e_b^{B*}, e_p^{I*})/\tau)}, \quad (7)$$

where \mathcal{L}_U^C and \mathcal{L}_B^C denote the cross-view contrastive losses for users and bundles, respectively; $s(\cdot, \cdot)$ is the cosine similarity function; τ is a hyper-parameter known as the *temperature*. We follow SGL [30] to perform in-batch negative sampling to construct the negative pairs. By averaging the two cross-view contrastive losses, we obtain the final contrastive loss \mathcal{L}^C :

$$\mathcal{L}^C = \frac{1}{2} (\mathcal{L}_U^C + \mathcal{L}_B^C). \quad (8)$$

2.4 Prediction and Optimization

To obtain the final prediction for recommendation, we first utilize the inner-product to calculate the item-view and bundle-view predictions, and additively combine them for the final prediction.

$$y_{u,b}^* = e_u^{B* \top} e_b^{B*} + e_u^{I* \top} e_b^{I*}. \quad (9)$$

The conventional Bayesian Personalized Ranking (BPR) loss [25] is used as the main loss.

$$\mathcal{L}^{BPR} = \sum_{(u,b,b') \in Q} -\ln \sigma(y_{u,b}^* - y_{u,b'}^*). \quad (10)$$

where $Q = \{(u, b, b') | u \in \mathcal{U}, b, b' \in \mathcal{B}, x_{ub} = 1, x_{ub'} = 0\}$, $\sigma(\cdot)$ is the sigmoid function. We achieve the final loss \mathcal{L} by weighted combing the BPR loss \mathcal{L}^{BPR} , the contrastive loss \mathcal{L}^C , and the L2 regularization term $\|\Theta\|_2^2$:

$$\mathcal{L} = \mathcal{L}^{BPR} + \lambda_1 \mathcal{L}^C + \lambda_2 \|\Theta\|_2^2, \quad (11)$$

where λ_1 and λ_2 are the hyperparameters to balance the three terms, and $\Theta = \{E_U^{B(0)}, E_B^{B(0)}, E_I^{I(0)}\}$ are all of the model parameters.

2.5 Model Discussion

Integrating the cross-view contrastive loss into BPR loss can provide an additional regularization for representation learning [37]. Such




Figure 3: The illustration of the direct (a) and indirect (b) effects of the cross-view contrastive loss.

an effect encourages the bundle recommender to achieve the cross-view alignment of the same user/bundle and impose the dispersion of different user/bundle pairs in both ego-view and cross-view.

By enforcing the cross-view alignment of the same user/bundle, the distinctive information contained in each view will be distilled to the other view. Therefore, both views’ representations can be enhanced. Consequently, the combined prediction of the two views can be further boosted. See more details and results in Section 3.4.1.

In addition to the cross-view alignment, the effect of representation dispersion is also pivotal. Based on Equations 6 and 7, it seems only the item/bundle pairs across different views are pushed away. However, the cross-view alignment also acts as a bridge to make the distinct user/bundle pairs within the same view be widely separated. Let’s take a pair of users u_1 and u_2 as an example, shown in Figure 3. The direct effect of the contrastive loss is to pull close the pairs $(e_{u_1}^{B^*}, e_{u_1}^{I^*})$ and $(e_{u_2}^{B^*}, e_{u_2}^{I^*})$ while push apart the pairs $(e_{u_1}^{B^*}, e_{u_2}^{I^*})$ and $(e_{u_1}^{I^*}, e_{u_2}^{B^*})$. Consequently, as an indirect effect, the distance between the representations of u_1 and u_2 in the same view (*i.e.*, $(e_{u_1}^{B^*}, e_{u_2}^{B^*})$ and $(e_{u_1}^{I^*}, e_{u_2}^{I^*})$) are also enlarged. Therefore, our proposed cross-view contrastive loss can enhance the discriminative capability of representations in both ego-view and cross-view, resulting in better bundle recommendation quality. To be noted, solely enlarging the cross-view dispersion without encouraging the cross-view alignment cannot achieve the effect of ego-view dispersion, thus cannot enhance the self-discrimination of the representations. We will justify this effect by analyzing the alignment-dispersion characteristics of representations in Section 3.4.2.

2.6 Complexity Analysis

In terms of space complexity, the parameters of CrossCBR are minimal and only include three sets of embeddings: $E_U^{B(0)}$, $E_B^{B(0)}$, and $E_I^{I(0)}$. Therefore, the space complexity of CrossCBR is $O((M + N + O)d)$. Our model is more concise than BGCN due to the removal of the feature transformation matrices.

In terms of time complexity, the main computational cost lies in the two views’ graph learning and the cross-view contrastive loss. Note we just focus on the main setting of original preservation. The time complexity of graph learning is $O(|E|Kds \frac{|E|}{T})$, where $|E|$ is the number of all edges in U-B and U-I graphs, K is the number of propagation layers, d is the embedding size, s is the number of epochs, T is the batch size. For comparison, the time complexity of BGCN is $O((|E|d + |V|d^2)Ks \frac{|E|}{T})$, where $|V|$ is the number of nodes in U-B and U-I graphs. And the time complexity of graph learning part in CrossCBR is smaller than that of BGCN due to the removal of the feature transformation layers and smaller $|E|$ (due to the removal of self-connections and bundle-bundle connections). The time complexity of calculating the contrastive loss during training

Table 1: Dataset Statistics.

Dataset	#U	#I	#B	#U-I	#U-B	#Avg.I/B
Youshu	8,039	32,770	4,771	138,515	51,377	37.03
NetEase	18,528	123,628	22,864	1,128,065	302,303	77.80
iFashion	53,897	27,694	42,563	2,290,645	1,679,708	3.86

is $O(2d(|E| + T^2)s \frac{|E|}{T})$. Even though the matrix multiplication of contrastive loss has cubic complexity ($O(dT^2)$), the time used in practice is very limited due to the acceleration of both hardware GPU and the optimized libraries, justified by the experiments in Section 3.5.2. To be noted, the inference of CrossCBR has the identical time complexity with that of BGCN.

3 EXPERIMENTS

To evaluate our proposed approach, we conduct experiments on three public datasets: Youshu, NetEase, and iFashion. In particular, we aim to answer the following research questions:

- **RQ1:** Can CrossCBR outperform the SOTA baseline models?
- **RQ2:** Are all the key components effective *w.r.t.* performance?
- **RQ3:** Whether the cross-view contrastive learning works as we expected, *i.e.*, achieving cross-view mutual enhancement and alignment and dispersion in the representation space?
- **RQ4:** What about the hyper-parameter sensitivity and training efficiency of the model?

3.1 Experimental Settings

We follow the previous works [3, 9] to adopt the two established bundle recommendation datasets: Youshu [5] for book list recommendation and NetEase [2] for music playlist recommendation. In addition, we introduce another online fashion outfit recommendation dataset iFashion [8], where the outfit consisted of individual fashion items is treated as bundle. We follow the outfit recommendation setting [20] to preprocess the iFashion dataset by the 20-core rule for users and 10-core rule for outfits. All the three datasets have all the required data, *i.e.*, user-bundle interactions, user-item interactions, and bundle-item affiliation. The statistics of the datasets are listed in Table 1. To be noted, the three datasets are diverse *w.r.t.* both application scenarios and the statistical characteristics (various scales of interactions and bundle sizes), ensuring the model’s robustness to such variance. The training/validation/testing sets are randomly split with the ratio of 70%/10%/20%. Recall@K and NDCG@K are used as the evaluation metrics, where $K \in \{20, 40\}$. And NDCG@20 is used to select the best model based on the validation set, and all the items are ranked during testing [28].

3.1.1 Compared Methods. In terms of baselines, we select both general user-item recommendation models and bundle-specific recommendation models to compare with our proposed method.

The User-item Recommendation Models treat the bundle as a special type of *item*, only using the user-bundle interactions without considering the affiliated items within the bundle. We select the following SOTA methods: (1) **MFBR** [25]: Matrix Factorization optimized with the Bayesian Personalized Ranking (BPR) loss; (2) **LightGCN** [14]: a GNN- and CF-based recommendation model, which utilizes a light-version graph learning kernel; and (3) **SGL** [30]: it enhances the LightGCN model with contrastive graph learning and achieves SOTA performance.

Table 2: The overall performance comparison, where Rec is short of Recall. Note that the improvement achieved by CrossCBR is significant (p -value $\ll 0.05$).

Model	Youshu				NetEase				iFashion			
	Rec@20	NDCG@20	Rec@40	NDCG@40	Rec@20	NDCG@20	Rec@40	NDCG@40	Rec@20	NDCG@20	Rec@40	NDCG@40
MFBPR	0.1959	0.1117	0.2735	0.1320	0.0355	0.0181	0.0600	0.0246	0.0752	0.0542	0.1162	0.0687
LightGCN	0.2286	0.1344	0.3190	0.1592	0.0496	0.0254	0.0795	0.0334	0.0837	0.0612	0.1284	0.0770
SGL	<u>0.2568</u>	<u>0.1527</u>	<u>0.3537</u>	<u>0.1790</u>	<u>0.0687</u>	<u>0.0368</u>	<u>0.1058</u>	<u>0.0467</u>	<u>0.0933</u>	<u>0.0690</u>	<u>0.1389</u>	<u>0.0851</u>
DAM	0.2082	0.1198	0.2890	0.1418	0.0411	0.0210	0.0690	0.0281	0.0629	0.0450	0.0995	0.0579
BundleNet	0.1895	0.1125	0.2675	0.1335	0.0391	0.0201	0.0661	0.0271	0.0626	0.0447	0.0986	0.0574
BGCN	0.2347	0.1345	0.3248	0.1593	0.0491	0.0258	0.0829	0.0346	0.0733	0.0531	0.1128	0.0671
CrossCBR	0.2813	0.1668	0.3785	0.1938	0.0842	0.0457	0.1264	0.0569	0.1173	0.0895	0.1699	0.1080
%Improv.	9.57	9.26	7.02	8.28	22.57	24.33	19.48	21.96	25.76	29.63	22.33	26.85

The Bundle-specific Recommendation Models are designed for bundle recommendation and utilize all the user-bundle interactions, user-item interactions, and bundle-item affiliation data. We consider the following models: (1) **DAM** [5]: it uses an attention mechanism to learn bundle representations over the affiliated items and employs multi-task learning to optimize both user-item and user-bundle interactions; (2) **BundleNet** [9]: it builds a user-bundle-item tripartite graph, leverages GCN to learn the representations, and applies multi-task learning; and (3) **BGCN** [3, 4]: it decomposes the user-bundle-item relations into two separate views, builds two graphs (*i.e.*, bundle-view graph and item-view graph), uses GCN to learn representations, makes prediction by summing the two views' predictions, and achieves SOTA performance. There are also some earlier works (*e.g.*, [2]) that have been turned to be inferior to the methods listed above, and we do not consider them.

3.1.2 Hyper-parameter Settings. For all methods, the embedding size is set as 64, Xavier normal initialization [12] is adopted, the models are optimized using Adam optimizer [17] with the learning rate 0.001, and the batch size is set as 2048. For our method, we tune the hyper-parameters K , λ_1 , λ_2 , τ , and ρ with the ranges of $\{1, 2, 3\}$, $\{0.01, 0.04, 0.1, 0.5, 1\}$, $\{10^{-6}, 10^{-5}, 2 \times 10^{-5}, 4 \times 10^{-5}, 10^{-4}\}$, $\{0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.4, 0.5\}$, and $\{0, 0.1, 0.2, 0.5\}$. For graph augmentation, we follow SGL to drop edges by every epoch. For baseline methods, we adopt the results of MFBPR, DAM, and BGCN on Youshu and NetEase datasets based on those reported in [3], since their settings are the same with ours. We implement all the other baselines by ourselves. All the models are trained using Pytorch 1.9.0, NVIDIA Titan-XP and Titan-V GPUs.

3.2 Performance Comparison (RQ1)

We first compare the overall recommendation performance of CrossCBR with both user-item recommendation baselines and bundle-specific recommendation baselines on three datasets, as shown in Table 2. The best performing methods are bold, while the strongest baselines are underlined; %Improv. measures the relative improvements of CrossCBR over the strongest baselines. We observe that:

In terms of the general user-item recommendation models, LightGCN consistently outperforms MFBPR, indicating the GNN-based method especially the LightGCN graph learning module is effective in modeling the user-bundle CF signals. SGL further improves the performance of LightGCN, demonstrating the great power of contrastive loss on the user-bundle bipartite graph. Surprisingly, SGL is the strongest baseline, which is even better than the strongest

Table 3: Ablation study of the key components of CrossCBR.

Model	NetEase		iFashion	
	Rec@20	NDCG@20	Rec@20	NDCG@20
BGCN	0.0491	0.0258	0.0733	0.0531
CrossCBR-CL	0.0608	0.0320	0.0852	0.0626
CrossCBR_A	0.0602	0.0314	0.0664	0.0474
CrossCBR_D	0.0410	0.0208	0.0568	0.0402
CrossCBR_OP	0.0831	0.0452	0.1167	0.0891
CrossCBR_ED	0.0842	0.0457	0.1176	0.0891
CrossCBR_MD	0.0828	0.0457	0.1173	0.0895
CrossCBR+SC	0.0831	0.0455	0.1146	0.0873
CrossCBR+BB	0.0828	0.0444	0.1163	0.0887
CrossCBR	0.0842	0.0457	0.1173	0.0895
%Improv.	71.43	77.29	52.50	64.54

bundle-specific method (BGCN), implying the effectiveness of graph contrastive learning in recommendation. Our proposed CrossCBR performs better than SGL by a large margin, showing that the item view truly provides additional useful information and can enhance the discriminative power of the model.

When considering the bundle-specific models, BGCN performs best among bundle-specific baselines, *i.e.*, DAM and BundleNet. We attribute this success to decomposing the users' preference into two views. Unfortunately, BundleNet performs poorly since the user-bundle-item tripartite graph fails to differentiate behavioral similarity among users and content relatedness. CrossCBR achieves significant gains over all the bundle-specific baselines by a large margin, demonstrating the effectiveness of modeling the cross-view cooperative association. Our model performs consistently on all the three datasets that belong to varying application scenarios (including book, music, and fashion) and have different statistical characteristics (scales of interactions and bundle sizes). Therefore, our model turns out to be sufficiently robust.

3.3 Ablation Study (RQ2)

To further evaluate the key innovative components of CrossCBR, we conduct a list of ablation studies as shown in Table 3, where the %Improv. quantifies the relative improvement of CrossCBR over the SOTA bundle-specific model BGCN.

3.3.1 Effectiveness of Cross-view Contrastive Learning. To evaluate whether the cross-view contrastive learning contributes to the performance, we remove the contrastive loss \mathcal{L}^C in Equation (8) during training, named CrossCBR-CL. CrossCBR-CL inevitably shows a severe performance decline compared with CrossCBR, justifying

the crucial role of modelling cross-view information. Surprisingly, CrossCBR-CL still significantly outperforms BGCN. We ascribe its success to utilizing the LightGCN kernel, which has been proved to be more effective than the typical GCN kernel used by BGCN [14]. To further identify the characteristics of alignment and dispersion, we implement CrossCBR_A that only enlarges the cross-view alignment (using the negative cross-view cosine similarity to replace the contrastive loss), and CrossCBR_D that only encourages the cross-view dispersion (setting the cosine similarity in numerator of the contrastive loss as static 1). The results in Table 3 demonstrates that the alignment and dispersion collaboratively contribute to the performance, while only modeling either of them can degrade the performance. Especially when only enlarging the cross-view dispersion (CrossCBR_D), the model can hardly benefit from it or even collapse, justifying our discussion in Section 2.5.

3.3.2 Effectiveness of Data Augmentations. We try various settings of data augmentations during the representation learning of the two views. CrossCBR_OP corresponds to original preservation (*i.e.*, no augmentation), CrossCBR_ED represents Edge Dropout of the graph-based augmentation method, and CrossCBR_MD refers to Message Dropout of the embedding-based augmentation method. The results in Table 3 demonstrate that the differences among the three data augmentation settings for CrossCBR are negligible compared with the performance gain over baselines. This phenomenon indicates that the distinction within the original data of the two views provides sufficient variance for the cross-view contrastive learning, while the variance introduced by random data augmentation is insignificant. More advanced and effective data augmentation methods can be explored in the future.

3.3.3 The Impact of Simplification of Graph Structure. CrossCBR simplifies the graph structure of BGCN by removing self-connections (SC) in the U-B and U-I graphs and the bundle-bundle (BB) connections in the U-B graph. To justify that these removals do not affect the performance, we specifically add SC and BB connections to our graph construction and obtain CrossCBR+SC and CrossCBR+BB, respectively. The results of both CrossCBR+SC and CrossCBR+BB indicate that both SC and BB contribute little or even none to CrossCBR. The reasons are two-fold. First, in terms of SC, the LightGCN kernel has no feature transformation and activation layers, resulting in the SC a simple summation of the node itself (layer 0), which adds no additional information. Second, for the BB connections, they are obtained by calculating the overlap degree of the two bundles according to the B-I graph. However, the BB overlap information can be distilled from the item view to the bundle view representations through the alignment effect of the cross-view contrastive loss.

3.4 Model Study (RQ3)

In this section, we conduct experiments to study: (1) whether the two views are mutually enhanced by the cross-view contrastive loss? and (2) whether the cross-view alignment and dispersion happen as we expected in the representation space?

3.4.1 Mutual Enhancement Effect. To directly justify whether the cross-view contrastive learning can achieve cross-view mutual enhancement, we present the performance which is calculated solely based on ego-view predictions, *i.e.*, the bundle-view prediction uses




Figure 4: The performance of CrossCBR and CrossCBR-CL w.r.t. predictions based on different views.

Table 4: The cross-view alignment and dispersion analysis of the representations. A denotes Alignment; D denotes Dispersion; superscripts (C, B, I) denote the cross, bundle, and item view; subscripts (U, B) stand for users and bundles.

Metrics	NetEase		iFashion	
	CrossCBR-CL	CrossCBR	CrossCBR-CL	CrossCBR
A_U^C	0.638	0.932	0.878	0.950
D_U^B	0.313	0.049	0.331	0.014
D_U^I	0.044	0.004	0.193	0.004
A_B^C	0.351	0.632	0.635	0.910
D_B^B	0.040	0.042	0.052	0.033
D_B^I	0.075	0.016	0.011	0.030

$y_{u,b}^* = e_u^{B* \top} e_b^{B*}$, the item-view prediction uses $y_{u,b}^* = e_u^{I* \top} e_b^{I*}$, and the both-view prediction is identical with Equation 9. The results in Figure 4 indicate that using contrastive loss significantly improves the recommendation performance on all the three types of predictions. Interestingly, the bundle view prediction performs much better than that of the item view (even slightly better than the both-view prediction in NetEase), demonstrating that the bundle view plays a more significant role in bundle recommendation. This also helps explain why SGL, which is solely based on the user-bundle interactions, can outperform a list of bundle-specific methods.

3.4.2 Cross-View Alignment and Dispersion Analysis. We analyze the cross-view alignment and dispersion characteristics of the representations regularized by the cross-view contrastive learning. Inspired by the alignment-uniformity analysis [11, 27], we adopt a simplified version to portray the cross-view alignment and dispersion of the representations. In particular, we just calculate the average cross-view *cosine similarity* of the users (bundles) as the indication of alignment. Similarly, the average *cosine similarity* between different users (bundles) of the same view (either item or bundle view) is the indication of the dispersion. Several interesting findings can be derived from the results shown in Table 4. First, the cross-view alignment metrics of both user and bundle representations (A_U^C and A_B^C) significantly increase after applying the contrastive loss, justifying that the cross-view contrastive loss can effectively pull the two views close to each other. Thereafter, each view can distill cooperative information from the other and they are mutually enhanced. Second, the dispersion of the user representations of both ego-views (D_U^B and D_U^I) significantly reduces after applying the contrastive loss, verifying that the cross-view contrastive loss can improve the discrimination of the users in the ego-view (*cf.* Section 2.5). Third, the dispersion of the bundle representations (D_B^B and D_B^I) does not consistently reduce like that of




Figure 5: The performance (NDCG@20) variance of Cross-CBR w.r.t. the temperature τ and the batch size on both datasets of NetEase and iFashion.

Table 5: The statistics of one-epoch training time (seconds) for CrossCBR and baselines on different devices, where the "Cr" is short of "CrossCBR".

Device	Dataset	BGCN	Cr-CL	Cr_OP	Cr+SC	Cr+BB
Titan XP	NetEase	32.05	6.80	7.04	7.27	28.02
	iFashion	63.47	46.05	46.74	47.61	56.42
Titan V	NetEase	20.76	4.67	5.09	5.48	18.59
	iFashion	38.76	29.48	30.02	30.31	35.01

user's. It may be because that the B-I graph directly determines the dispersion of item-view bundle representations via the pooling (cf. Equation 5), which is distilled to the bundle view.

3.5 Hyper-parameter and Computational Efficiency Analysis (RQ4)

3.5.1 Hyper-parameter Analysis. As shown in Figure 5, CrossCBR is sensitive to the temperature τ , and deviating from the best setting degrades the performance remarkably. To test how the batch size affects the performance, we gradually increase the batch size from 128 to 8192, and the performance first grows quickly and later reaches a plateau as shown in Figure 5. We keep our default batch size as 2048, since it is widely adopted by the baselines' original implementation and performs sufficiently well in CrossCBR.

3.5.2 Computational Efficiency. To evaluate the computational efficiency of our model, we compare the one-epoch training time among the variants of our model and BGCN. We record ten consecutive training epochs and average them to obtain the one-epoch training time, as shown in Table 5¹. First of all, Cr_OP is significantly more efficient than BGCN, demonstrating the efficiency of CrossCBR. Second, we compare Cr_OP with several variants to further explicitly attribute the efficiency boost. In particular, Cr-CL approximates Cr_OP, showing that the contrastive loss brings negligible computational costs. Cr+SC and Cr+BB cost more training time than Cr_OP, demonstrating that both SC and BB connections introduce extra costs during training. Especially on NetEase, the costs brought by the BB connections are about three times of Cr_OP.

4 RELATED WORK

In this section, we briefly review the related works in two areas: (1) graph-based and bundle recommendation and (2) contrastive learning in recommendation.

¹The CPU is Intel(R) Xeon(R) CPU E5-2620 v3 with 12 cores, and the memory is 256GB. We use multiprocess dataloader to assure that the computation on CPU is not the bottleneck. In addition, the GPU utility is almost 100% during the training of all the settings, therefore, the capacity of GPU dominates the overall training time.

Graph-based and Bundle Recommendation. Graph-based model has dominated the CF-based recommendation methods due to its superior capability in modeling the higher order interactions between users and items, especially the recent graph neural network-based methods [7, 14, 28, 35]. Wang *et al.* propose NGCF [28] to build a bipartite graph based on the user-item interaction matrix and employ graph convolutional network as the graph learning module. Following NGCF, He *et al.* propose to remove some redundant modules (*e.g.*, nonlinear feature transformation and activation function layers) from the NGCF model and significantly improve the performance, resulting in the LightGCN [14] model. LightGCN has achieved great performance in various recommendation tasks [10], and our model is also base on this backbone.

Bundle recommendation aims to solve a special scenario of recommendation, *i.e.*, the recommended object is a bundle of items that related with a certain theme. Initial works just ignore the affiliated items of the bundle and just use an id to represent a bundle [26]. Following works recognize the importance of affiliated items and develop various models to capture the additional user-item interaction and bundle-item affiliation relations, such as EFM [2] and DAM [5]. With the proliferation of GNN-based recommendation models, Deng *et al.* propose BundleNet [9] and Chang *et al.* propose BGCN [3, 4]. However, BundleNet mixup the three types of relations among user, bundle, and item, while BGCN decompose the users' preference into item view and bundle view. The two-view representations effectively capture the two types of preferences, resulting in better performance. Our work is based on this two-view modeling framework, and we further emphasize the significance of the cooperative association modeling between the two views. Some related topics, such as set, basket, or package recommendation [16, 18, 24] and bundle generation [1, 4], are different with our scenario in either the recommended object (a loosely/arbitrary co-occurred set/basket/package vs a pre-defined bundle of items related with a theme) or the task (generation of bundle from items vs recommending pre-defined bundles).

Contrastive Learning in Recommendation. Recently, contrastive learning regains popularity and achieves great success in computer vision [6, 15, 23], natural language processing [11, 22], and graph learning [19, 31]. The community of recommender systems also seizes this trend and adapts contrastive learning into various recommendation tasks, such as general CF-based recommendation [30, 38, 40], sequential and session recommendation [21, 32–34, 39], multimedia and social recommendation [29, 36], *etc.* The key of introducing contrastive learning into recommender systems lies in proper construction of contrastive pairs. One branch of current approaches are based on various data augmentations to create more views from the original data. For example, SGL [3] adopts various graph augmentation methods (*e.g.*, dege dropout or random walk), and CL4SRec [34] and CoSeRec [21] apply different sequence augmentation methods (*e.g.*, insertion, deletion, and reordering *etc.*). Another branch of methods aim at mining multiple views that exist in the data. For example, COTREC [32] builds two views (*i.e.*, an item view and a session view) to learn the session representations from two sources of data (*i.e.*, item transition graph of a session and session-session similarity graph), and apply contrastive learning based on the two views. CLCRec [29] treats different modalities and the user/item as different views to build contrastive pairs. In

this work, we unify both types of methods: build two views from different data sources and apply data augmentations.

5 CONCLUSION AND FUTURE WORK

In this work, we applied the cross-view contrastive learning to model the cross-view cooperative association in bundle recommendation. We introduced the cross-view contrastive learning to regularize the cross-view representations and proposed a simple, efficient, yet effective method CrossCBR, which significantly enhanced the SOTA performance of bundle recommendation on three public datasets. Various ablation and model studies demystified the working mechanism behind such huge performance leap.

Even though CrossCBR has achieved great performance, the study of contrastive learning on bundle or even general recommendation is still in its infancy, and several directions are promising in the future. First, model-based data augmentations, which can introduce both hard negative and diverse positive samples, should be helpful for further performance improvements. Second, more potential approaches are to be explored for modeling the cross-view cooperative association. Third, the cross-view contrastive learning paradigm is easy to be generalized to other similar tasks, as long as two distinctive views exist.

ACKNOWLEDGEMENT

This research/project is supported by the Sea-NExT Joint Lab, and CCCD Key Lab of Ministry of Culture and Tourism.

REFERENCES

- [1] Jinze Bai, Chang Zhou, Junshuai Song, Xiaoru Qu, Weiting An, Zhao Li, and Jun Gao. 2019. Personalized Bundle List Recommendation. In *WWW*. ACM, 60–71.
- [2] Da Cao, Liqiang Nie, Xiangnan He, Xiaochi Wei, Shunzhi Zhu, and Tat-Seng Chua. 2017. Embedding Factorization Models for Jointly Recommending Items and User Generated Lists. In *SIGIR*. ACM, 585–594.
- [3] Jianxin Chang, Chen Gao, Xiangnan He, Depeng Jin, and Yong Li. 2020. Bundle Recommendation with Graph Convolutional Networks. In *SIGIR*. ACM, 1673–1676.
- [4] Jianxin Chang, Chen Gao, Xiangnan He, Depeng Jin, and Yong Li. 2021. Bundle Recommendation and Generation with Graph Neural Networks. *IEEE Transactions on Knowledge and Data Engineering* (2021).
- [5] Liang Chen, Yang Liu, Xiangnan He, Lianli Gao, and Zibin Zheng. 2019. Matching User with Item Set: Collaborative Bundle Recommendation with Deep Attention Network. In *IJCAI*. 2095–2101.
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. In *ICML (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 1597–1607.
- [7] Weijian Chen, Yulong Gu, Zhaochun Ren, Xiangnan He, Hongtao Xie, Tong Guo, Dawei Yin, and Yongdong Zhang. 2019. Semi-supervised User Profiling with Heterogeneous Graph Attention Networks. In *IJCAI*, Vol. 19. 2116–2122.
- [8] Wen Chen, Pipei Huang, Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, and Binqiang Zhao. 2019. POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion. In *KDD*. ACM, 2662–2670.
- [9] Qilin Deng, Kai Wang, Minghao Zhao, Zhene Zou, Runze Wu, Jianrong Tao, Changjie Fan, and Liang Chen. 2020. Personalized Bundle Recommendation in Online Games. In *CIKM*. ACM, 2381–2388.
- [10] Yujuan Ding, Yunshan Ma, Wai Keung Wong, and Tat-Seng Chua. 2021. Leveraging Two Types of Global Graph for Sequential Fashion Recommendation. In *ICMR*. ACM, 73–81.
- [11] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *EMNLP (1)*. Association for Computational Linguistics, 6894–6910.
- [12] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS (JMLR Proceedings, Vol. 9)*. JMLR.org, 249–256.
- [13] Michael Gutmann and Aapo Hyvärinen. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS (JMLR Proceedings, Vol. 9)*. JMLR.org, 297–304.
- [14] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *SIGIR*. ACM, 639–648.
- [15] R. Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. Learning deep representations by mutual information estimation and maximization. In *ICLR*. OpenReview.net.
- [16] Haoji Hu, Xiangnan He, Jinyang Gao, and Zhi-Li Zhang. 2020. Modeling Personalized Item Frequency Information for Next-basket Recommendation. In *SIGIR*. ACM, 1071–1080.
- [17] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [18] Chen Li, Yuanfu Lu, Wei Wang, Chuan Shi, Ruobing Xie, Haili Yang, Cheng Yang, Xu Zhang, and Leyu Lin. 2021. Package Recommendation with Intra- and Inter-Package Attention Networks. In *SIGIR*. ACM, 595–604.
- [19] Sihang Li, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. 2022. Let Invariant Rationale Discovery Inspire Graph Contrastive Learning. In *ICML*.
- [20] Xingchen Li, Xiang Wang, Xiangnan He, Long Chen, Jun Xiao, and Tat-Seng Chua. 2020. Hierarchical Fashion Graph Network for Personalized Outfit Recommendation. In *SIGIR*. ACM, 159–168.
- [21] Zhiwei Liu, Yongjun Chen, Jia Li, Philip S Yu, Julian McAuley, and Caiming Xiong. 2021. Contrastive self-supervised sequential recommendation with robust augmentation. *arXiv preprint arXiv:2108.06479* (2021).
- [22] Lajanugen Logeswaran and Honglak Lee. 2018. An efficient framework for learning sentence representations. In *ICLR (Poster)*. OpenReview.net.
- [23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [24] Yuqi Qin, Pengfei Wang, and Chenliang Li. 2021. The World is Binary: Contrastive Learning for Denoising Next Basket Recommendation. In *SIGIR*. ACM, 859–868.
- [25] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [26] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized Markov chains for next-basket recommendation. In *WWW*. ACM, 811–820.
- [27] Tongzhou Wang and Phillip Isola. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In *ICML (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 9929–9939.
- [28] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *SIGIR*. ACM, 165–174.
- [29] Yinwei Wei, Xiang Wang, Qi Li, Liqiang Nie, Yan Li, Xuanping Li, and Tat-Seng Chua. 2021. Contrastive Learning for Cold-Start Recommendation. In *ACM Multimedia*. ACM, 5382–5390.
- [30] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. 2021. Self-supervised Graph Learning for Recommendation. In *SIGIR*. ACM, 726–735.
- [31] Yingxin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. 2022. Discovering Invariant Rationales for Graph Neural Networks. In *ICLR*.
- [32] Xin Xia, Hongzhi Yin, Junliang Yu, Yingxia Shao, and Lizhen Cui. 2021. Self-Supervised Graph Co-Training for Session-based Recommendation. In *CIKM*. ACM, 2180–2190.
- [33] Xin Xia, Hongzhi Yin, Junliang Yu, Qinyong Wang, Lizhen Cui, and Xiangliang Zhang. 2021. Self-Supervised Hypergraph Convolutional Networks for Session-based Recommendation. In *AAAI*. AAAI Press, 4503–4511.
- [34] Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Bolin Ding, and Bin Cui. 2020. Contrastive Learning for Sequential Recommendation. *arXiv preprint arXiv:2010.14395* (2020).
- [35] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. 2018. Graph Convolutional Neural Networks for Web-Scale Recommender Systems. In *KDD*. ACM, 974–983.
- [36] Junliang Yu, Hongzhi Yin, Min Gao, Xin Xia, Xiangliang Zhang, and Nguyen Quoc Viet Hung. 2021. Socially-Aware Self-Supervised Tri-Training for Recommendation. In *KDD*. ACM, 2084–2092.
- [37] Yifan Zhang, Bryan Hooi, Dapeng Hu, Jian Liang, and Jiashi Feng. 2021. Unleashing the Power of Contrastive Self-Supervised Visual Models via Contrast-Regularized Fine-Tuning. In *NeurIPS*. 29848–29860.
- [38] Chang Zhou, Jianxin Ma, Jianwei Zhang, Jingren Zhou, and Hongxia Yang. 2021. Contrastive Learning for Debaised Candidate Generation in Large-Scale Recommender Systems. In *KDD*. ACM, 3985–3995.
- [39] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-Rec: Self-Supervised Learning for Sequential Recommendation with Mutual Information Maximization. In *CIKM*. ACM, 1893–1902.
- [40] Xin Zhou, Aixun Sun, Yong Liu, Jie Zhang, and Chunyan Miao. 2021. SelfCF: A Simple Framework for Self-supervised Collaborative Filtering. *arXiv preprint arXiv:2107.03019* (2021).