

Document Version

Final published version

Citation (APA)

Nguyen, H. S., Bui, T. N., Nguyen, L. H., Hoang, H., Thi Nguyen, C. V., Le, H. Q., & Le, D. T. (2024). Bundle Recommendation with Item-Level Causation-Enhanced Multi-view Learning. In A. Bifet, P. Daniušis, J. Davis, T. Krilavičius, M. Kull, E. Ntoutsi, K. Puolamäki, & I. Žliobaitė (Eds.), *Machine Learning and Knowledge Discovery in Databases. Research Track and Demo Track - European Conference, ECML PKDD 2024, Proceedings* (pp. 324-341). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 14948 LNAI). Springer. https://doi.org/10.1007/978-3-031-70371-3_19

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Bundle Recommendation with Item-Level Causation-Enhanced Multi-view Learning

Huy-Son Nguyen, Tuan-Nghia Bui, Long-Hai Nguyen, Hung Hoang,
Cam-Van Thi Nguyen, Hoang-Quynh Le, and Duc-Trong Le^(✉)

University of Engineering and Technology, Hanoi Vietnam National University,
Hanoi, Vietnam
{huyson, 21020364, 21020624, 21020518, vanntc, lhquynh, trongld}@vnu.edu.vn

Abstract. Bundle recommendation aims to enhance business profitability and user convenience by suggesting a set of interconnected items. In real-world scenarios, leveraging the impact of asymmetric item affiliations is crucial for effective bundle modeling and understanding user preferences. To address this, we present **BunCa**, a novel bundle recommendation approach employing item-level causation-enhanced multi-view learning. **BunCa** provides comprehensive representations of users and bundles through two views: the Coherent View, leveraging the Multi-Perspective Causation Network for causation-sensitive relations among items, and the Cohesive View, employing LightGCN for information propagation among users and bundles. Modeling user preferences and bundle construction combined from both views ensures rigorous cohesion in direct user-bundle interactions through the Cohesive View and captures explicit intents through the Coherent View. Simultaneously, the integration of concrete and discrete contrastive learning optimizes the consistency and self-discrimination of multi-view representations. Extensive experiments with **BunCa** on three benchmark datasets demonstrate the effectiveness of this novel research and validate our hypothesis.

Keywords: Bundle Recommendation · Collaborative Filtering · Graph Neural Network · Contrastive Learning

1 Introduction

Recommendation systems are crucial in enhancing user experiences and shaping business strategies, particularly in e-commerce [19, 24]. While conventional recommendation systems traditionally prioritize individual item suggestions, bundle recommendations have been recognized as a superior strategic marketing tactic. Bundle recommendation, rooted in user behavior and item relevance, involves grouping pertinent items into bundles, such as assortments of items from the same category (e.g., *detective books, thrilling games*) [5, 19], complementary item bundles (e.g., *men's vest with cravat, phone with case*) [19], etc. Recommending bundles based on user preferences poses greater challenges compared to separate items due to data sparsity and the diverse composition of bundles.

H.-S. Nguyen and T.-N. Bui—Equal Contribution.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2024
A. Bifet et al. (Eds.): ECML PKDD 2024, LNAI 14948, pp. 324–341, 2024.
https://doi.org/10.1007/978-3-031-70371-3_19

Prior studies have mainly focused on selecting items for bundling [2] and identifying product bundles in commercial markets [20]. However, recent works [4, 12] have shifted towards recommending existing bundles, aiming to better align suggested bundles with users’ preferences. State-of-the-art (SOTA) bundle recommendation such as MIDGN [23], CrossCBR [12], and BundleGT [22] aggregate items to generate user and bundle representations. However, they often overlook the intricate relationships between individual items within bundles. There has been an initial exploration of the impact of item relationships, often grounded in symmetric correlations, on purchasing decisions for a bundle [1, 10]. In real-world scenarios, the influence between items within bundles is usually asymmetric, with anchor items playing a significant role. Symmetric correlation-based recommendation models, as illustrated in Fig. 1, fall short in accurately capturing user preferences. While *blazers* and *cravats* may be frequently purchased together, their influence on each other is not necessarily symmetrical. Users with an interest to *blazers* might consider to combine with *cravats* as accessories, while others primarily interested in *cravats* may already have suitable *blazers*, making the combination with *pocket squares* more logical. Understanding these dynamics requires examining item-item relations across diverse user preferences, which may help improve the bundle recommendation performance.

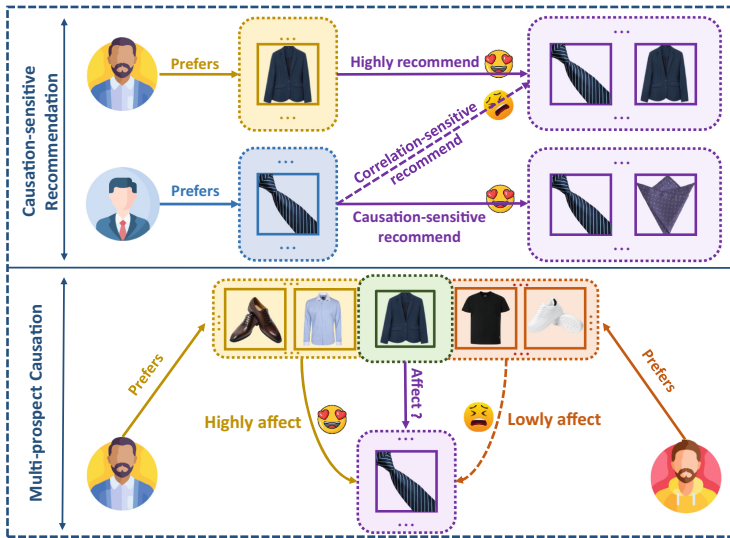


Fig. 1. Motivating examples of multi-prospect causation in bundle recommendation.

Approach and Contributions. Motivated by mentioned limitations of previous works, we introduce BunCa a novel neural architecture for bundle recommendation, whereby leverages asymmetric relationships between bundle items via the

item-level causation-enhanced multi-view learning. To the best of our knowledge, this paper represents the first comprehensive study to explore asymmetric item relationships, which is *our first contribution*. In Sect. 3, these information is transmitted in a multi-view learning framework which consists of two distinct views namely Coherent View and Cohesive View. The Cohesive View learning aims to encode rigorously high-order collaborative signals in user-bundle interactions. Besides, the Multi-Prospect Causation Networks for causation-sensitive relations between items are integrated to generate enhanced twin representations of objects in the Coherent View learning. Modeling user preferences and bundle construction combined from both views ensures rigorous cohesion in direct interactions between users-bundles via the Cohesive View, as well as explicit meaning of their intents via the Coherent View. *As our second contribution*, we employ concrete and discrete contrastive learning simultaneously, enhancing the consistency and self-discrimination of representations from both views. Section 4 presents *our final contribution* in conducting extensive experiments on three benchmark datasets. The experimental results demonstrate that BunCa has potential comparison with state-of-the-art methods for the bundle recommendation task, substantiating the rationalization of our proposed approach and new prominent issues in existing bundle-related tasks.

2 Related Work

The field of bundle recommendation has witnessed diverse approaches aimed at accurately recommending pre-defined item sets to users. Early approaches [11], grounded in the BPR framework [15], deliberate users' past interactions with lists of items and individual items. DAM model [5] emerged as a pioneer by recognizing the importance of affiliated items within bundles and optimizing both user-item and user-bundle interactions using attention mechanisms and multi-task learning. Recently, BGCN [4] used Graph Convolutional Network operations to capture intricate relations between users, items, and bundles from two distinct views, establishing an effective multi-view learning approach for subsequent SOTA models.

In the realm of multi-view architectures, addressing challenges related to the inconsistency and integration of information concerning objects has been tackled through the adoption of contrastive learning (CL) techniques. CrossCBR [12], evolving from BGCN, stands out for its remarkable improvements, leveraging cross-view contrastive learning to highlight the significance of capturing cooperative information from two distinct views. MIDGN [23] introduces the concept of intent disentanglement, modeling multiple latent features for each view and employing contrastive loss between the two views. Our study introduces concrete contrastive learning to underscore the self-discrimination of augmented representations in multi-view learning. In addition, discrete contrastive learning is employed to address the problem of inconsistency between separate views.

The capture of intricate connections between items influencing bundling strategies and user preferences has gained attention. MIDGN [23], which tackles

hidden intents inside user preference and bundle construction, encounters performance and complexity challenges. BundleGT [22] achieves SOTA performance with a Hierarchical Graph Transformer network. Regarding the recommendation of item sets, Beacon [10] utilizes correlation-sensitive signals to enhance the prediction of user with next-basket. It is noteworthy that relationships between items in real-world scenarios are not symmetric, in contrast to the simulation in [10,22]. Integrating homogeneous graphs is also employed to leverage item-item relationships, solely utilizing the LightGCN operation [13,16]. Furthermore, some further approaches such as EBRec augmenting observed user-item interactions with pre-trained models [7], or DGMAE distillation framework [14] also introduce several new aspects. Our study focuses on exploiting asymmetric item-item relationships via the features of user preference and bundle construction, leading to enhanced user/bundle representations compared to previous studies.

3 Methodology

In this section, we initially formalize the bundle recommendation problem and define the detailed interaction graph construction. Subsequently, we demonstrate the four important modules of BunCa illustrated in Fig. 2 including: (1) Cohesive View representation learning; (2) Coherent View representation learning; (3) contrastive learning modules; and (4) prediction and joint optimization module.

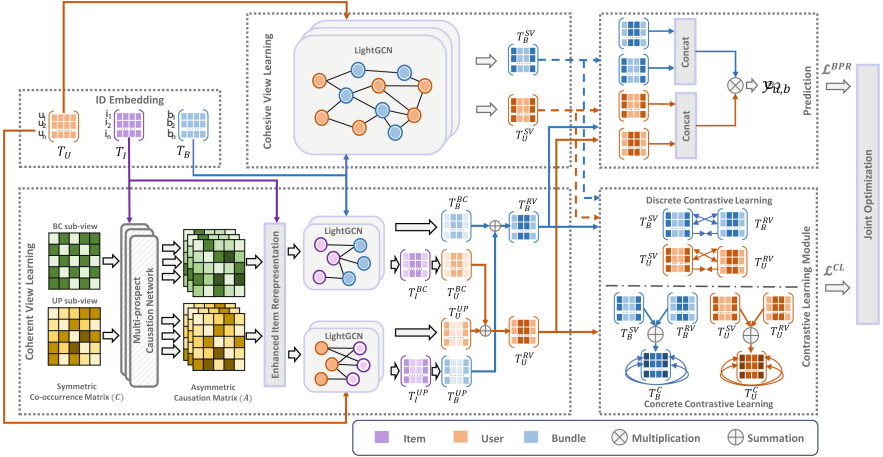


Fig. 2. The schematic illustration of our proposed model BunCa.

3.1 Preliminaries

Problem Formulation. Given the set of users $\mathcal{U} = \{u_1, u_2, \dots, u_{|\mathcal{U}|}\}$, the set of bundles $\mathcal{B} = \{b_1, b_2, \dots, b_{|\mathcal{B}|}\}$, and the set of items $\mathcal{I} = \{i_1, i_2, \dots, i_{|\mathcal{I}|}\}$.

The user-bundle interactions, user-item interactions, and bundle-item affiliations are respectively defined as three binary-valued matrices $X \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{B}|}$, $Y \in \{0, 1\}^{|\mathcal{U}| \times |\mathcal{I}|}$, and $Z \in \{0, 1\}^{|\mathcal{B}| \times |\mathcal{I}|}$ where the cell with value 1 denotes an observed relation between the user-bundle, user-item, or bundle-item pair, value 0 for otherwise. The objective of our work is to accurately predict unseen user-bundle interactions for recommendation system.

Graph Construction in the Cohesive View. Inspired by collaborative filtering concepts suggesting that users with mutual interactions share similar preferences [21], we assume that user preference patterns are hidden within the co-occurrence of purchased bundles. In addition, bundles often bought collectively by numerous individuals may indicate shared interests. To capture these hypotheses, the user co-occurrence matrix $C_{\mathcal{U}} = X \cdot X^{\top} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{U}|}$ and the bundle co-occurrence matrix $C_{\mathcal{B}} = X^{\top} \cdot X \in \mathbb{R}^{|\mathcal{B}| \times |\mathcal{B}|}$ are both derived from user-bundle interactions. Thereby, we explicitly establishes sets of connections denoted as: $\mathcal{E}_{\mathcal{UB}} = \{e_{u,b} | b \in \mathcal{B} \wedge u \in \mathcal{U} \wedge X_{(u,b)} = 1\}$, $\mathcal{E}_{\mathcal{U}} = \{e_{u,u'} | (u, u' \in \mathcal{U}) \wedge C_{\mathcal{U}(u,u')} \geq 1\}$, $\mathcal{E}_{\mathcal{B}} = \{e_{b,b'} | (b, b' \in \mathcal{B}) \wedge C_{\mathcal{B}(b,b')} \geq 1\}$. BunCa integrates user-bundle interactions with homogeneous correlations to thoroughly exploit the intrinsic relationships among users and bundles in graph $\bar{\mathcal{G}} = \{\bar{\mathcal{V}}, \bar{\mathcal{E}}\}$, where $\bar{\mathcal{V}} = \{\mathcal{U} \cup \mathcal{B}\}$ and $\bar{\mathcal{E}} = \{\mathcal{E}_{\mathcal{UB}} \cup \mathcal{E}_{\mathcal{U}} \cup \mathcal{E}_{\mathcal{B}}\}$ respectively describe the set of vertices and the set of edges.

Graph Construction in the Coherent View. The bundle-item bipartite graph $\mathcal{G}_{\mathcal{BI}}$, and user-item bipartite graph $\mathcal{G}_{\mathcal{UI}}$ are leveraged to effectively aggregate information on item-level representations. The graph $\mathcal{G}_{\mathcal{BI}} = \{\mathcal{V}_{\mathcal{BI}}, \mathcal{E}_{\mathcal{BI}}\}$ is constructed by the affiliations between bundles and items in Z , where $\mathcal{V}_{\mathcal{BI}} = \{\mathcal{B} \cup \mathcal{I}\}$ and $\mathcal{E}_{\mathcal{BI}} = \{e_{b,i} | b \in \mathcal{B}, i \in \mathcal{I}\}$ represent the set of vertices and edges. User preferences, derived from historical interactions with standalone items, also impact decision-making for bundle purchases. Thus, the graph $\mathcal{G}_{\mathcal{UI}} = \{\mathcal{V}_{\mathcal{UI}}, \mathcal{E}_{\mathcal{UI}}\}$, where $\mathcal{V}_{\mathcal{UI}} = \{\mathcal{U} \cup \mathcal{I}\}$ and $\mathcal{E}_{\mathcal{UI}} = \{e_{u,i} | u \in \mathcal{U}, i \in \mathcal{I}\}$ are established based on the interactions between users and items in Y . Furthermore, at the beginning of the training phase, the features of users $T_{\mathcal{U}} \in \mathbb{R}^{|\mathcal{U}| \times d}$, bundles $T_{\mathcal{B}} \in \mathbb{R}^{|\mathcal{B}| \times d}$, and items $T_{\mathcal{I}} \in \mathbb{R}^{|\mathcal{I}| \times d}$ are randomly initialized, where d is the embedding dimensionality.

3.2 Cohesive View Representation Learning

As mentioned, the heterogeneous graph $\bar{\mathcal{G}}$ may help enhance robust propagation of high-order collaborative signals by incorporating homogeneous correlations. BunCa adopts the LightGCN operation [9] to encode cohesive representations of users and bundles, excluding self-connections and non-linear transformations in the propagation function. The h -th layer's information propagation of $\bar{\mathcal{G}}$ is:

$$t_v^{SV(h)} = \sum_{v' \in \mathcal{N}_v} \frac{t_{v'}^{SV(h-1)}}{\sqrt{|\mathcal{N}_v|} \sqrt{|\mathcal{N}_{v'}|}} \quad (1)$$

where $t_v^{SV(h)} \in \mathbb{R}^d$, $t_{v'}^{SV(h)} \in \mathbb{R}^d$ represent the h -th layer’s information propagated to nodes $v, v' \in \bar{\mathcal{V}}$; \mathcal{N}_v and $\mathcal{N}_{v'}$ denote the first-hop neighbors of node v and v' in $\bar{\mathcal{G}}$ respectively. Thereby, the embedding of user t_u^{SV} and bundle t_b^{SV} learned via Cohesive View are inferred as follows:

$$t_u^{SV} = \sum_{h=0}^H t_u^{SV(h)} \quad , \quad t_b^{SV} = \sum_{h=0}^H t_b^{SV(h)} \quad (2)$$

where, the initial embeddings $t_u^{SV(0)}$ and $t_b^{SV(0)}$ are obtained from $T_{\mathcal{U}}$ and $T_{\mathcal{B}}$.

Overall, the representations of all users $T_{\mathcal{U}}^{SV} \in \mathbb{R}^{|\mathcal{U}| \times d}$ and bundles $T_{\mathcal{B}}^{SV} \in \mathbb{R}^{|\mathcal{B}| \times d}$ through the unification of all embeddings of different layers rigorously integrate information propagated from high-order neighbors.

3.3 Coherent View Representation Learning

For the Coherent View, **BunCa** aims to explicitly leverage the asymmetric relation between items based on the user preference (UP) sub-view and the bundle construction (BC) sub-view to enhance the modeling of decision-making factors. Learning these sub-views improves item embeddings, hence consolidating the representations for users and bundles. These representations are fused to highlight vital information that determines user preferences and bundle construction. Initially, we employ a coarse-grained mixture method to create item-item symmetric relations. The co-occurrence item matrices from different aspects are computed by the matrix multiplication as follows:

$$C_{\mathcal{I}}^{UP} = Y^{\top} \cdot Y \quad , \quad C_{\mathcal{I}}^{BC} = Z^{\top} \cdot Z \quad (3)$$

where $C_{\mathcal{I}}^{UP}, C_{\mathcal{I}}^{BC} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{I}|}$ represent item co-occurrence based on user-item interactions Y and bundle-item affiliations Z . The symmetric connections between items are normalized as follows:

$$\tilde{C}_{(i,j)} = \begin{cases} 1 & \text{if } C_{(i,j)} \geq 1 \wedge i \neq j, \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $i, j \in \mathcal{I}$ refer to the corresponding items. Filtering the threshold of $C_{(i,j)}$ can be considered empirically to achieve better performance on different data domains. It is notable that the symmetric matrix $\tilde{C}_{\mathcal{I}}$ is uniform symbol for $C_{\mathcal{I}}^{UP}, C_{\mathcal{I}}^{BC}$ due to the similar calculations of two sub-views in many scenarios.

Multi-prospect Causation Network. Assuming that causation-sensitive relationships exist among items frequently purchased together, **BunCa** employs Multi-Prospect Causation Network (MPCNet) to explicitly model asymmetric associations between items. MPCNet is constructed with L prospects, each represented by a learnable prospect vector $\mathbf{p}_l \in \mathbb{R}^d$. For the l -th prospect, the weight $r_{j \rightarrow i}$ signifies the influence from item j to item i based on various user preferences and bundling strategies, derived as follows:

$$r_{j \rightarrow i}^l = \mathbf{p}_l^\top \cdot \varphi(\Psi_{src}^l \cdot t_j \oplus \Psi_{dst}^l \cdot t_i \oplus \Phi) \quad (5)$$

where $\Psi_{src}^l, \Psi_{dst}^l \in \mathbb{R}^{d \times d}$ are learnable parameters for the source and destination object in the l -th prospect; $t_i, t_j \in \mathbb{R}^d$ denote the representations of item i and item j ; $\varphi(\cdot)$ is the non-linear activation function; \oplus denotes the element-wise summation; and Φ is the bias parameter.

In the l -th prospect, the asymmetric causation matrix $A^l \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{I}|}$, representing the causation-sensitive relationships among items at fine-grained level, is computed by the attention mechanism concept of GATv2 [3]. The weight $A_{(i,j)}^l$ describes how much item i is influenced by item j , defined as follows:

$$A_{(i,j)}^l = \frac{\exp(r_{j \rightarrow i}^l) \odot \tilde{C}_{(i,j)}}{\max(\sum_{j' \in \mathcal{I}} \exp(r_{j' \rightarrow i}^l) \odot \tilde{C}_{(i,j')}, \epsilon)} \quad (6)$$

where $\epsilon \in R$ is a fixed constant; \odot display the multiplication between two scalars.

Enhancing Item Representation. The asymmetric relationships obtained from MPC-Net are utilized to encode the latent representation of item $t_i^l \in \mathbb{R}^d$ in the l -th prospect, formulated as follows:

$$t_i^l = \sum_{j \in \mathcal{I}} A_{(i,j)}^l \Psi_{src}^l \cdot t_j \quad (7)$$

where $\Psi_{src}^l \in \mathbb{R}^{d \times d}$ denotes the dense layer to learn item features; and $t_j \in \mathbb{R}^d$ denotes other item representation. Subsequently, the multi-prospect item representation is devised using the residual connection method as:

$$t_i^f = \alpha \frac{1}{L} \sum_{k=1}^L t_i^k \oplus (1 - \alpha) t_i \quad (8)$$

where $t_i^f \in \mathbb{R}^d$ is the enhanced item representation, combining individual features and aggregated information from connected items with a controlled influence $\alpha \in [0, 1]$. It is notable that t_i^f serves as a uniform placeholder, representing both $\hat{t}_i^{UP} \in \hat{T}_{\mathcal{I}}^{UP}$ and $\hat{t}_i^{BC} \in \hat{T}_{\mathcal{I}}^{BC}$, denoting the enhanced item representation of the UP sub-view and BC sub-view.

For UP sub-view, the item representation \hat{t}_i^{UP} is employed alongside the user's preference $t_u \in T_{\mathcal{U}}$ as the node's initial feature within the user-item graph $\mathcal{G}_{\mathcal{UI}}$. The graph-based propagation process utilizes the LightGCN operation similar to Eq. (1). The ultimate representation of users and items is denoted as $T_{\mathcal{U}}^{UP} \in \mathbb{R}^{|\mathcal{U}| \times d}$ and $T_{\mathcal{I}}^{UP} \in \mathbb{R}^{|\mathcal{I}| \times d}$. The item features is aggregated to derive the bundle features $T_{\mathcal{B}}^{UP} \in \mathbb{R}^{|\mathcal{B}| \times d}$ utilizing the mean pooling operation as:

$$t_b^{UP} = \sum_{i \in \mathcal{N}_b^{BI}} \frac{t_i^{UP}}{|\mathcal{N}_b^{BI}|} \quad (9)$$

where $t_b^{UP}, t_i^{UP} \in \mathbb{R}^d$ describe the aggregated representation of bundle b and item i in the UP sub-view; \mathcal{N}_b^{BI} is the neighbor set of bundle b in graph \mathcal{G}_{BI} .

Similar operations are employed in the BC sub-view learning procedure. The graph \mathcal{G}_{BI} takes $t_b \in T_B$ and t_i^{BC} as input representations for nodes in the first layer. The information propagation is performed using the LightGCN operation, which results in the final representations of bundles and items respectively denoted as $T_B^{BC} \in \mathbb{R}^{|\mathcal{B}| \times d}$ and $T_I^{BC} \in \mathbb{R}^{|\mathcal{I}| \times d}$. Subsequently, the final user representations $T_U^{BC} \in \mathbb{R}^{|\mathcal{U}| \times d}$ in the BC sub-view is derived as follows:

$$t_u^{BC} = \sum_{i \in \mathcal{N}_u^{UI}} \frac{t_i^{BC}}{|\mathcal{N}_u^{UI}|} \quad (10)$$

where $t_u^{BC} \in T_U^{BC}, t_i^{BC} \in T_I^{BC}$ describe the aggregated representation of user u and item i in BC sub-view; \mathcal{N}_u^{UI} denotes the neighbors of user u in graph \mathcal{G}_{UI} .

The comprehensive information is integrated by the attentive summation as:

$$\begin{aligned} t_u^{RV} &= \beta t_u^{BC} \oplus (1 - \beta) t_u^{UP}, \\ t_b^{RV} &= \beta t_b^{BC} \oplus (1 - \beta) t_b^{UP} \end{aligned} \quad (11)$$

where the hyper-parameter $\beta \in [0, 1]$ is to maintain a balance between the propagated features from the two sub-view learning phases. $t_u^{RV}, t_b^{RV} \in \mathbb{R}^d$ denotes the final representation of user u and bundle b through Coherent View learning.

3.4 Contrastive Learning Module

The discrete contrastive learning is directed at minimizing the inconsistency between two different views, whereas the concrete contrastive learning strives to enhance the distinctiveness of augmented representations.

Discrete Contrastive Learning. Inspired by CrossCBR [12], BunCa addresses variations in user preferences for items and bundles by minimizing inconsistencies in representations across different views. The in-batch negative sampling is adopted to construct the negative pairs [17]. Relied on InfoNCE [8], the discrete contrastive loss function for user representations is computed from contrastive pairs (t_u^{SV}, t_u^{RV}) as follows:

$$\mathcal{L}_u^{DC} = -\frac{1}{|\mathcal{U}|} \sum_{u' \in \mathcal{U}} \log \frac{\exp((\cos(t_u^{SV}, t_u^{RV})/\tau))}{\sum_{u' \in \mathcal{U}} \exp((\cos(t_u^{SV}, t_{u'}^{RV})/\tau)} \quad (12)$$

where u' represents a user other than u ; $\cos(\cdot, \cdot)$ denotes the cosine similarity function; τ is the temperature hyper-parameter that help to distinguish the positive and negative samples. The formula for calculating \mathcal{L}_b^{DC} is derived similarly.

Concrete Contrastive Learning. The concrete contrastive learning is devised to distinguish the unique characteristics of each individual user and bundle. We create the combination of user preferences and bundle features as follows:

$$t_u^C = t_u^{RV} \oplus t_u^{SV} \quad , \quad t_b^C = t_b^{RV} \oplus t_b^{SV} \tag{13}$$

where t_u^C, t_b^C denote the fused multi-view representation of user and bundle, respectively. The formulation of the concrete contrastive loss function for user representations is derived as follows:

$$\mathcal{L}_u^{CC} = -\frac{1}{|\mathcal{U}|} \sum_{u' \in \mathcal{U}} \log \frac{\exp(1/\tau)}{\sum_{u' \in \mathcal{U}} \exp((\cos(t_u^C, t_{u'}^C)/\tau))} \tag{14}$$

where u' denotes a different user of u ; $\cos(.,.)$ denotes the cosine similarity function; and τ is the temperature hyper-parameter that aids in distinguishing between positive and negative samples. The same operation is repeated for the concrete contrastive loss of bundle \mathcal{L}_b^{CC} .

The final contrastive loss is derived from the dyadic contrastive learning techniques described above as follows:

$$\mathcal{L}^{CL} = \gamma \left(\frac{\mathcal{L}_u^{DC} + \mathcal{L}_b^{DC}}{2} \right) + (1 - \gamma) \left(\frac{\mathcal{L}_u^{CC} + \mathcal{L}_b^{CC}}{2} \right) \tag{15}$$

where $\gamma \in [0, 1]$ controls which contrastive learning is more important.

3.5 Prediction and Joint Optimization

For the comprehensive representations, we effectively capture augmented information from both views, derived as follows:

$$\tilde{t}_u = [\mu t_u^{SV} \parallel t_u^{RV}] \quad , \quad \tilde{t}_b = [\mu t_b^{SV} \parallel t_b^{RV}] \tag{16}$$

where the hyper-parameter $\mu \in [0, 1]$ controls the impact of views on the final user’s preference and bundle’s feature; \parallel denotes the concatenation; and $\tilde{t}_u, \tilde{t}_b \in \mathbb{R}^{2d}$ respectively represent the fusion of user’s preferences and bundle’s features, combining Cohesive View and Coherent View learning.

Finally, the interaction probability $\hat{y}_{u,b} \in \mathbb{R}$ between user u and bundle b is calculated by the inner-product as follows:

$$\hat{y}_{u,b} = (\tilde{t}_u)^\top \cdot \tilde{t}_b \tag{17}$$

The Bayesian Personalized Ranking loss [15] is utilized to enhance probabilities for interacted user-bundle pairs while reducing for non-interacted pairs:

$$\mathcal{L}^{BPR} = \sum_{(u,b,b') \in \mathcal{S}} -\ln \sigma(\hat{y}_{u,b} - \hat{y}_{u,b'}) \tag{18}$$

Table 1. Statistics of Youshu, NetEase and iFashion datasets.

Dataset	$ \mathcal{U} $	$ \mathcal{I} $	$ \mathcal{B} $	$ \mathcal{E}_{\mathcal{UI}} $	$ \mathcal{E}_{\mathcal{UB}} $	#Avg.I/B	$c-score$
Youshu	8,039	32,770	4,771	138,515	51,377	37.03	0.0812
NetEase	18,528	123,628	22,864	1,128,065	302,303	77.80	0.0599
iFashion	53,897	42,563	27,694	2,290,645	1,679,708	3.86	0.2917

where $\mathcal{S} = \{(u, b, b') | u \in \mathcal{U}; b, b' \in \mathcal{B}; (u, b) \in X; (u, b') \notin X\}$ and $\sigma(\cdot)$ is the *sigmoid* activation function. The final loss \mathcal{L} is combined by weighted summation of \mathcal{L}^{BPR} , \mathcal{L}^{CL} and the $L2$ regularization $\|\Theta\|_2^2$ as:

$$\mathcal{L} = \mathcal{L}^{BPR} + \lambda_1 \mathcal{L}^{CL} + \lambda_2 \|\Theta\|_2^2 \quad (19)$$

where the hyper-parameter λ_1 controls the contrastive loss; λ_2 , is the regularization weight and Θ denotes the model’s trainable parameters.

4 Experiments

In this section, we conduct extensive experiments on benchmark datasets to validate the effectiveness of BunCa and investigate the importance of its components.

4.1 Experimental Setup

Datasets. According to popular studies [4, 12, 14], we adopt three benchmark datasets for the bundle recommendation evaluation, including: The *Youshu*¹ dataset consists of bundles formed by lists of books that is purchased in each user session; the *NetEase*² dataset captures user-generated song lists as bundles; the *iFashion* [6] dataset establishes outfits composed of clothes and accessories as bundles. The statistics of these datasets are shown in Table 1. Specially, $c-score$ metric, according to [14], present the high consistency between user-bundle and user-item collaborative relations on iFashion but limited on Youshu and NetEase.

Evaluation Metrics. Recall ($R@K$) and Normalized Discounted Cumulative Gain ($N@K$) are two commonly employed metrics for the performance of method in bundle recommendation task. $R@K$ measures the proportion of test bundles within the top- K ranking list. $N@K$ manifests normalized discounted cumulative gain scores aimed at obtaining relevant items at higher positions on the ranking list. Both metrics are employed with $K \in \{10, 20\}$ for performance validation. The smaller the top- K , the more clearly it demonstrates the model’s practical recommendation performance. Average performances in 5 runs with various random initialization are reported. Comparisons are evaluated by two-tailed paired-sample Student’s t-test with p -value of 0.05.

¹ <https://www.yousuu.com/>.

² <https://music.163.com/>.

Baselines. We compare BunCa to the three groups of state-of-the-art models:

- **Traditional Bundle Recommendation (T):** BPR [15], LightGCN [9], and DAM [5].
- **Multi-view Learning Bundle Recommendation (M):** CrossCBR [12], BGCN [4], MIDGN [23], BundleGT [22], and EBRec [7].
- **Distillation Bundle Recommendation (D):** DGMAE [14].

For BPR and LightGCN, we simply consider bundles at the item level similar to previous reputable studies [5, 7, 19]. For MIDGN, embedding each user or bundle into k chunks within different feature spaces presents challenges for optimizing performance on large datasets like iFashion, particularly when computational resources are limited. For DGMAE, obtaining complete materials from the authors to ensure absolute reproducibility still faces some difficulties. To facilitate a fair comparison, the available results have been compiled from relevant reputable papers.

Implementation Details. BunCa is implemented using PyTorch on NVIDIA P100 and T4 GPUs, with evaluated datasets in Table 1. Baseline methods are evaluated with the same below settings based on source code and results in reputable works [4, 12, 22]. The initial embedding size is configured as 64. The learning rate is set as $1e-3$. We adopt Xavier initialization and Adam optimizer for trainable parameters. We tune the regularization weight $\lambda_2 \in \{1e-6, 4e-6, 1e-5, 4e-5, 1e-4, 4e-4\}$. The $\alpha, \gamma, \mu, \beta, \lambda_1 \in [0; 1]$ are tuned by grid search.

4.2 Performance Comparison to Baselines

Table 2 illustrates performance comparison between BunCa and related baselines in term of $R@K, N@K$ metrics with $K \in \{10, 20\}$. Generally, the recommendation methods with multi-view learning consistently outperforms the traditional techniques for all metrics. This observation implies the efficiency of the decomposition and modeling of multiple views for the bundle recommendation task.

Among the multi-view learning models, BunCa achieves the best performance in the benchmark datasets. Compared to applied contrastive learning methods (e.g., MIDGN, BGCN, CrossCBR), BunCa represents persistent enhancements, which verify the effectiveness of assembling discrete and concrete contrastive learning in providing discriminative representation across different views. Compared to the second-best results (BundleGT, EBRec), BunCa produces statistically significant improvements in terms of $R@10$ and $N@10$. Especially, for iFashion, the notable enhancements are 8.25% to 9.57% for both $K \in \{10, 20\}$. Compared to the DGMAE distillation framework, the performance of BunCa is 13% and 17% superior on $R@20$ and $N@20$ for iFashion, but weak for Youshu and NetEase. They affirm the effectiveness of exploiting item causation-sensitive relationships in small-sized bundles, in which bundle items are more coherent.

For the case of large-sized bundle datasets such as NetEase and Youshu, the superiority of BunCa could be negatively affected due to noisy connections.

Table 2. Performance comparison to baselines. \star denotes values in the previous research paper, while others are in our experiments. The best performances is in **bold** and the second best values is underlined. *Imp* shows the relative improvements of BunCa over the second-best model. \dagger indicates statistically significant improvements. “—” presents missing values due to the unavailability or run-time problems of the source codes. (*T*), (*M*), (*D*) respectively represent the best results of each baseline group.

Dataset	Youshu				NetEase				iFashion			
Metric	<i>R@10</i>	<i>R@20</i>	<i>N@10</i>	<i>N@20</i>	<i>R@10</i>	<i>R@20</i>	<i>N@10</i>	<i>N@20</i>	<i>R@10</i>	<i>R@20</i>	<i>N@10</i>	<i>N@20</i>
BPR	15.03	21.49	10.11	11.96	0.12	0.22	0.08	0.11	0.46	0.81	0.41	0.55
LightGCN	16.93	24.34	12.22	14.35	3.24	5.42	2.16	2.84	3.25	5.33	2.92	3.78
DAM	11.03	20.82	8.73	11.98	2.65	4.11	1.43	2.10	3.17	6.29	2.66	4.50
BGCN	15.55	23.47 \star	10.99	13.45 \star	2.88	4.91 \star	1.88	2.58 \star	4.83	7.33 \star	4.34	5.31 \star
BundleGT	<u>20.01</u>	29.05 \star	14.16	17.37 \star	5.35	<u>9.03\star</u>	3.51	4.78 \star	8.16	12.63 \star	7.65	9.69 \star
MIDGN	18.87	26.82 \star	12.91	15.27 \star	3.74	6.78 \star	2.46	3.43 \star	—	6.22 \star	—	4.41 \star
EBRec	19.62	28.50	<u>14.27</u>	16.80	<u>5.49</u>	8.83	<u>3.69</u>	4.76	<u>9.01</u>	<u>13.22</u>	<u>8.57</u>	<u>10.28</u>
CrossCBR	19.61	28.13 \star	14.21	16.68 \star	5.11	8.42 \star	3.44	4.57 \star	7.72	11.73 \star	7.31	8.95 \star
DGMAE	—	32.45\star	—	19.18\star	—	9.45\star	—	5.12\star	—	12.62 \star	—	9.59 \star
BunCa	20.85\dagger	<u>29.78</u>	14.65\dagger	<u>17.62</u>	5.55	8.99	3.73	<u>4.79</u>	9.80\dagger	14.31\dagger	9.39\dagger	11.23\dagger
% \uparrow (<i>T</i>)	23.15	22.35	19.88	22.78	71.29	65.86	72.68	68.66	301	227	321	250
% \uparrow (<i>M</i>)	4.20	2.51	2.66	1.44	1.72	0.44	1.09	0.21	8.77	8.25	9.57	9.24
% \uparrow (<i>D</i>)	—	−8.22	—	−8.13	—	−4.87	—	−6.44	—	13.39	—	17.10

In real-world scenarios, users tend to buy bundles less frequently compared to individual items due to the sensitive prices. To effectively promote user consumption and business marketing strategy, the constructed bundles need to be oriented towards specific and explicit intentions [19, 22, 24]. Therefore, advanced approaches that exploit the interaction between items to model the bundle’s targets (e.g., BundleGT, EBRec, BunCa) are clearly more effective on iFashion compared to the other architectures. On the other hand, Youshu and Netease data are solely built from large pre-defined item sets via user session and the limited consistency between the U-B and U-I collaborative relations (*c-score*).

Moreover, we investigate the distribution of high-level influence items within bundles, determined by their popularity surpassing the average popularity of the items in the same bundle. According to Fig. 3, the distribution on iFashion indicates that each bundle’s characteristics are mainly shaped by one or two key components. Conversely, many items within the same bundle demonstrate high-level influence in NetEase dataset featuring large-sized bundles (similar to Youshu). This phenomenon causes noise and hinders the identification of anchor items necessary for transparent bundle modeling. DGMAE hardly models bundle-item affiliations in depth but only synthesizes bundle representations based on the average pooling of item representations [14]. Therefore, DGMAE’s performance can be superior on Youshu and NetEase but are significantly limited on iFashion. The above evidence points to a contrast with the practical intuition of bun-




Fig. 3. Statistical distribution of High-level Influence Items within Bundles.

Table 3. The performance of **BunCa** once omitting various sub-view separately. The \downarrow denotes the percentage decrease when each specific view is ablated from **BunCa**.

Dataset	Metric	w/o RV	w/o SV	w/o UP	w/o BC	BunCa
Youshu	$R@20$	26.54(\downarrow 10.84)	26.06(\downarrow 12.46)	29.38(\downarrow 1.32)	27.88(\downarrow 6.36)	29.77
	$N@20$	15.38(\downarrow 12.69)	15.03(\downarrow 14.70)	16.85(\downarrow 4.34)	16.65(\downarrow 5.47)	17.62
NetEase	$R@20$	5.65(\downarrow 37.17)	6.00(\downarrow 33.28)	8.08(\downarrow 10.21)	8.56(\downarrow 4.81)	8.99
	$N@20$	2.92(\downarrow 39.00)	3.11(\downarrow 35.03)	4.39(\downarrow 8.21)	4.65(\downarrow 2.82)	4.79
iFashion	$R@20$	5.36(\downarrow 62.52)	12.79(\downarrow 10.61)	14.20(\downarrow 0.74)	11.48(\downarrow 19.79)	14.31
	$N@20$	3.77(\downarrow 66.37)	10.53(\downarrow 6.07)	11.11(\downarrow 10.91)	9.15(\downarrow 18.39)	11.21

dle recommendation tasks. This statement demystifies the notable consideration of the reasonableness and quality of popular benchmark datasets for bundle recommendation, which is also mentioned recent studies [18, 19].

4.3 Ablation Study

To investigate the impact of key components in **BunCa**, we conduct various ablation studies described below.

Effects of Learning from Different Views. To assess the impact of integrating latent representations from different views, learning Cohesive View (SV), Coherent View (RV), UP sub-view, and BC sub-view are selectively omitted and evaluated performance on the three datasets. According to Table 3, the results indicate that learning representations from each view encodes valuable information to enhance prediction performance. Excluding SV learning significantly diminishes **BunCa**'s performance on all datasets, demonstrating its ability to effectively leverage high-order collaborative signals between users and bundles. Besides, the performance drops remarkably once RV-view learning is omitted. It demystifies the effectiveness of learning causation-sensitive relationships between items to enhance latent representations. UP sub-view representations exhibit a

Table 4. The impact of different contrastive learning modules on the BunCa performance. The \downarrow denotes the ratio decrease when ablating each corresponding module.

Dataset	Metric	w/o CL	w/o DC	w/o CC	BunCa
Youshu	$R@20$	27.49(\downarrow 7.64)	25.50(\downarrow 14.33)	29.16(\downarrow 2.06)	29.77
	$N@20$	15.97(\downarrow 9.34)	15.03(\downarrow 14.70)	17.03(\downarrow 3.34)	17.62
NetEase	$R@20$	5.78(\downarrow 35.71)	6.55(\downarrow 27.16)	8.72(\downarrow 3.00)	8.99
	$N@20$	3.05(\downarrow 36.30)	3.61(\downarrow 24.52)	4.67(\downarrow 2.54)	4.79
iFashion	$R@20$	6.39(\downarrow 55.35)	7.25(\downarrow 49.35)	13.46(\downarrow 5.91)	14.31
	$N@20$	4.55(\downarrow 59.42)	5.33(\downarrow 52.51)	10.28(\downarrow 8.32)	11.21

marginal decrease of approximately 1% when omitted on iFashion, contributing the least valuable information. Conversely, the performance of UP sub-view significantly affected on NetEase with about 10%. The results on iFashion without the BC sub-view highlight the explicit bundle construction more than the results on Youshu and NetEase. This demonstration also contributes to clarifying the observation in Sect. 4.2. In summary, integrating latent representations through multi-view learning ensures both cohesion and coherence in bundle construction and user preferences, effectively improving performance.

Effect of Contrastive Learning Module. As depicted in Table 4, we examine the effectiveness of contrastive learning module through experiments involving the removal of the entire module (*w/o CL*), discrete contrastive learning (*w/o DC*), and concrete contrastive learning (*w/o CC*). The results across the three datasets simultaneously emphasize the pivotal role of modeling multi-view information integrated with contrastive learning. Notably, the substantial impact of discrete contrastive learning demonstrates its effectiveness in aligning the representations of separate views. This implies that each view can extract cooperative information from the other, leading to mutual enhancement. The addition of concrete contrastive learning further confirms that the contrastive learning module can enhance the discrimination of each user/bundle in the fused representations.

Importance of Asymmetric Causation Matrix. The causation matrix A learned by the MPCNet module enhances item-level representation t_i^f for bundle recommendation. In order to assess its significance, we compare it against alternatives: i) the original symmetric co-occurrence matrix from UP sub-view (*w UP-sym*) or BC sub-view (*w BC-sym*); ii) both original symmetric co-occurrence matrices from UP and BC sub-views (*w UP-BC-sym*); iii) normalized Laplacian symmetric correlation matrix as in [10] (*w SnLm*). Figure 4 shows a notable performance decrease of BunCa without asymmetric matrix-enhanced item representations, instead simply employing symmetric item-item matrices. This demonstration validates the significance of asymmetric relationships in real-world scenarios, affirming the hypothesis discussed in Sect. 1.




Fig. 4. Importance of asymmetric causation matrix A on the performance of BunCa.




Fig. 5. The impact of key hyper-parameters on the performance of BunCa for iFashion. Figures (a) and (b) show the model performance calculated by $R@20$ and $N@20$ over the respective value ranges of β and L .

Effects of Key Hyper-parameters. In Fig. 5, the impact of key hyperparameters, β and L , on BunCa’s performance for iFashion is depicted. The parameter β regulates the contribution of enhanced item representations from the BC sub-view and UP sub-view. As β increases, the performance follows an ascending trend until reaching its peak at $\beta = 0.8$, underscoring the crucial role of the BC sub-view in refining item representations. Regarding L , the performance exhibits minor fluctuations within the range of 1 to 5 prospects before a decline is observed with $L = 6$. This decline is attributed to noise problems that increase simultaneously with latent features of high-order neighbors. After tuning, the optimal settings for iFashion are determined to be $\beta = 0.8$ and $L = 5$. Similar tuning processes for Youshu and NetEase yield comparable observations.

4.4 Qualitative Showcase

Some qualitative examples in the iFashion dataset is illustrated in Fig.6 due to the coherence and specificity of interconnected items, whereas Youshu and NetEase only provide the IDs of objects. The asymmetric weights indicate the difference effect among items, resulting in determining the explicit target of bundle features. Thereby, within each bundle, certain anchor items can be identified, and their influence on complementary items varies depending on causation signals. For instance, considering the bundle 15707 containing a *shirt* and *earrings*. The weights reveal that the *shirt* serves as the anchor item, as evidenced by its

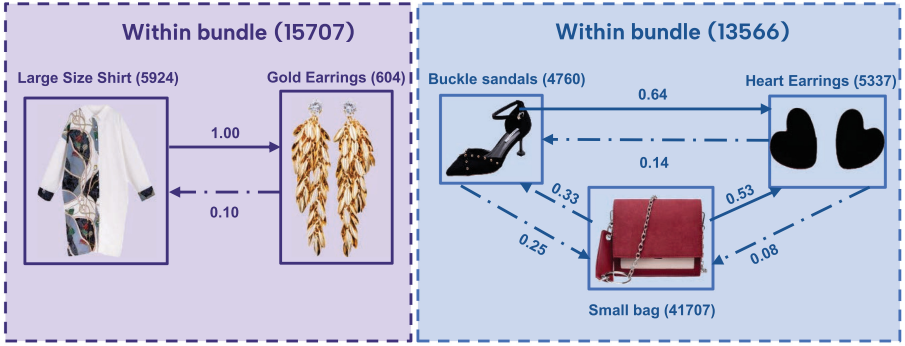


Fig. 6. The illustration of real cases in iFashion. Solid arrows indicate highly affected causation with asymmetric weight ≥ 0.5 whilst dash-dotted ones denote low effects.

influence on the *earring* with a weight of 1.0, while the impact of the *earring* on the *shirt* is only 0.1 in asymmetric weight. Likewise, the bundle 13566 shows reasonable causation weights from *buckle sandals* and a *small bag* to *heart earrings*. These evidence can be readily understood through the decision-making psychology of female customers. This emphasizes the ability to model key components within bundles via asymmetric relationships that adapt to user engagements. Therefore, it validates our hypothesis in leveraging item-level causation effects for bundle recommendation.

5 Conclusion

This study design a novel approach to address bundle recommendation by emphasizing the significance of asymmetric relationships between items. Our proposed model **BunCa** leverages item-level causation-enhanced multi-view learning, showcasing significant improvements and notable points, as evidenced by extensive analytics conducted on three benchmark datasets. The architecture of **BunCa** lies in learning two distinct views: the Coherent View, employing the Multi-Prospect Causation Network for causation-enhanced representation learning; and the Cohesive View, utilizing high-order collaborative signals of user-bundle interaction to aggregate comprehensive representations. In addition, **BunCa** simultaneously integrates concrete and discrete contrastive learning, enhancing the consistency and self-discrimination of representations from both views. Our work not only validates the efficacy of **BunCa** but also emphasizes the importance of considering asymmetric item relationships for bundle-related tasks.

Acknowledgements. This research was funded by the Vingroup Innovation Foundation under the project code VINIF.2022.DA00087 and the Master & PhD scholarship program code VINIF.2023.ThS.114.

References

1. Ariannezhad, M., Li, M., Jullien, S., de Rijke, M.: Complex item set recommendation. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 3444–3447 (2023)
2. Bai, J., et al.: Personalized bundle list recommendation. In: The World Wide Web Conference, pp. 60–71 (2019)
3. Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? In: International Conference on Learning Representations (2021)
4. Chang, J., Gao, C., He, X., Jin, D., Li, Y.: Bundle recommendation with graph convolutional networks. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1673–1676 (2020)
5. Chen, L., Liu, Y., He, X., Gao, L., Zheng, Z.: Matching user with item set: collaborative bundle recommendation with deep attention network. In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, pp. 2095–2101 (2019)
6. Chen, W., et al.: POG: personalized outfit generation for fashion recommendation at Alibaba iFashion. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, pp. 2662–2670 (2019)
7. Du, X., Qian, K., Ma, Y., Xiang, X.: Enhancing item-level bundle representation for bundle recommendation. *ACM Trans. Recommender Syst.* (2023)
8. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: a new estimation principle for unnormalized statistical models. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. *JMLR Workshop and Conference Proceedings*, pp. 297–304 (2010)
9. He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., Wang, M.: LightGCN: simplifying and powering graph convolution network for recommendation. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 639–648 (2020)
10. Le, D.T., Lauw, H.W., Fang, Y.: Correlation-sensitive next-basket recommendation. In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, pp. 2808–2814 (2019)
11. Liu, Y., Xie, M., Lakshmanan, L.V.: Recommending user generated item lists. In: Proceedings of the 8th ACM Conference on Recommender Systems, pp. 185–192 (2014)
12. Ma, Y., He, Y., Zhang, A., Wang, X., Chua, T.S.: CrossCBR: cross-view contrastive learning for bundle recommendation. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pp. 1233–1241 (2022)
13. Zhou, H., Zhou, X., Zhang, L., Shen, Z.: Enhancing dyadic relations with homogeneous graphs for multimodal recommendation. In: *ECAI 2023*, pp. 3123–3130. IOS Press (2023)
14. Ren, Y., Haonan, Z., Fu, L., Wang, X., Zhou, C.: Distillation-enhanced graph masked autoencoders for bundle recommendation. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1660–1669 (2023)
15. Rendle, S., Freudenthaler, C., Gantner, Z., Schmidt-Thieme, L.: BPR: Bayesian personalized ranking from implicit feedback. In: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, pp. 452–461 (2009)

16. Nguyen, H.S., et al.: HHMC: a heterogeneous x homogeneous graph-based network for multimodal cross-selling recommendation. In: 2023 15th International Conference on Knowledge and Systems Engineering (KSE), pp. 1–6. IEEE (2023)
17. Wu, J., et al.: Self-supervised graph learning for recommendation. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 726–735 (2021)
18. Sun, Z., et al.: Revisiting bundle recommendation for intent-aware product bundling. *ACM Trans. Recommender Syst.* **2**(3), 1–34 (2024)
19. Sun, Z., Yang, J., Feng, K., Fang, H., Qu, X., Ong, Y.S.: Revisiting bundle recommendation: datasets, tasks, challenges and opportunities for intent-aware product bundling. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 2900–2911 (2022)
20. Tzaban, H., Guy, I., Greenstein-Messica, A., Dagan, A., Rokach, L., Shapira, B.: Product bundle identification using semi-supervised learning. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 791–800 (2020)
21. Wang, X., He, X., Wang, M., Feng, F., Chua, T.S.: Neural graph collaborative filtering. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 165–174 (2019)
22. Wei, Y., Liu, X., Ma, Y., Wang, X., Nie, L., Chua, T.S.: Strategy-aware bundle recommender system. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1198–1207 (2023)
23. Zhao, S., Wei, W., Zou, D., Mao, X.: Multi-view intent disentangle graph networks for bundle recommendation. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, pp. 4379–4387 (2022)
24. Zhu, T., Harrington, P., Li, J., Tang, L.: Bundle recommendation in ecommerce. In: Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval, pp. 657–666 (2014)