

# Modeling Item-Level Dynamic Variability with Residual Diffusion for Bundle Recommendation

Dong Zhang<sup>1,2</sup>, Lin Li<sup>1\*</sup>, Ming Li<sup>1,3</sup>, Amran Bhuiyan<sup>3</sup>, Meng Sun<sup>1</sup>, Xiaohui Tao<sup>4</sup>, Jimmy Xiangji Huang<sup>3</sup>

<sup>1</sup>School of Computer Science and Artificial Intelligence, Wuhan University of Technology, China

<sup>2</sup>Engineering Research Center of Hubei Province for Clothing Information, Wuhan Textile University, China

<sup>3</sup>Information Retrieval and Knowledge Management Research Lab, School of Information Technology, York University, Canada

<sup>4</sup>School of Mathematics, Physics and Computing, University of Southern Queensland, Australia

{zhangdong23,cathylilin,liming7677,sunmeng}@whut.edu.cn, amran.apece@gmail.com, Xiaohui.Tao@unisq.edu.au, jhuang@yorku.ca

## Abstract

Existing solutions for bundle recommendation (BR) have achieved remarkable effectiveness for predicting the user’s preference for prebuilt bundles. However, bundle-item (B-I) affiliation will vary dynamically in real scenarios. For example, a bundle themed as ‘casual outfit’ may add ‘hat’ or remove ‘watch’ due to factors such as seasonal variations, changes in user preferences or inventory adjustments. Our empirical study demonstrates that the performance of mainstream BR models may fluctuate or decline under item-level variability. This paper makes the first attempt to address the above problem and proposes Residual Diffusion for Bundle Recommendation (RDiffBR) as a model-agnostic generative framework which can assist a BR model in adapting this scenario. During the initial training of the BR model, RDiffBR employs a residual diffusion model to process the item-level bundle embeddings which are generated by the BR model to represent bundle theme via a forward-reverse process. In the inference stage, RDiffBR reverses item-level bundle embeddings obtained by the well-trained bundle model under B-I variability scenarios to generate the effective item-level bundle embeddings. In particular, the residual connection in our residual approximator significantly enhances BR models’ ability to generate high-quality item-level bundle embeddings. Experiments on six BR models and four public datasets from different domains show that RDiffBR improves the performance of Recall and NDCG of backbone BR models by up to 23%, while only increases training time about 4%.

**Code** — <https://github.com/WUT-IDEA/RDiffBR>

**Extended version** — <https://arxiv.org/abs/2507.03280>

## Introduction

With the rapid development of online platforms, bundle recommendation (BR) has become an important tool for enhancing user experience and promoting sales. Specifically, users are provided with one-stop services by adding a group of related items to a bundle (such as booklists, playlists,




Figure 1: An example showing that bundle-item affiliation is changed while bundle theme remains relatively stable.

fashion sets, and meal sets in takeout services) while improve sales efficiency and revenue.

Bundle recommendation usually learns from the user-bundle (U-B) interactions, user-item (U-I) interactions and bundle-item (B-I) affiliations to recommend a list of bundles to users. However, in real scenarios B-I affiliations may undergo dynamic changes due to factors such as seasonal variations, shifts in user preferences, or inventory adjustments, while the theme of the bundle remains. As illustrated in Figure 1, Bundle 1 initially includes ‘T-shirt’, ‘pants’, ‘shoes’ and ‘watch’ with the theme of ‘casual outfit’. In practice, platforms can add the ‘hat’ to align with summer promotions or remove the ‘watch’ based on user feedback. An intuitive idea is to regard it as a new bundle. However, such item-level changes may not definitely cause bundle theme shifting although bundle theme is jointly determined by the items within the bundle. As shown in Figure 1, the theme of Bundle 2 and Bundle 3 remain ‘casual outfit’ by item variation. It is advisable to give priority to reusing the identification information of the original bundle because this dynamic variability maintains the theme of bundle.

In recent years, lots of works have focused on bundle recommendation and achieved good results. Including factorization models represented by LIRE (Liu, Xie, and Lakshmanan 2014), EFM (Cao et al. 2017), attention-based models represented by DAM (Chen et al. 2019a), Attlist (He et al. 2019) and GNN-based models represented by Cross-CBR (Ma et al. 2022), BundleGT (Wei et al. 2023) and

\*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

DCBR (Li 2025). However, our experiments have found that the performance of existing mainstream bundle recommendation models will fluctuate or even decline under B-I dynamic variability scenarios we have mentioned above. For instance, we found that on Youshu dataset, Recall@20 of BGCN will decrease by up to 19% as items in bundles increases or decreases (see Section 3.2 for details). This is because existing bundle recommendation models lack the ability to effectively model the item-level dynamic variability when generating bundle representations of item-level.

In this paper, we make the first attempt to address the issue of item-level dynamic variability in bundle recommendation and propose a novel Residual Diffusion for Bundle Recommendation (*RDiffBR*) model-agnostic generative framework to assist the model in adapting this scenario. During the initial training of the BR model, RDiffBR employs a residual diffusion model to process the item-level bundle embeddings which are generated by BR model to represent bundle theme. *In the forward process*, RDiffBR injects controlled Gaussian noises into item-level bundle embeddings. *In the reverse process*, RDiffBR iteratively removes noises by a designed residual approximator. RDiffBR jointly trains this residual approximator and BR model via above forward-reverse process. *In the inference stage*, to generate the effective item-level bundle embeddings, RDiffBR reverses item-level bundle embeddings obtained by the well-trained bundle model under B-I variability scenarios rather than using pure noise as the starting point and conduct bundle recommendation. Experiments on six bundle models and four public datasets from different domains show that RDiffBR significantly improves the performance of bundle models in dynamic B-I variability scenario.

The main contributions of this work are summarized as follows. First, we propose a generative residual diffusion framework for bundle recommendation to address item-level dynamic variability issue. Secondly, our residual approximator in diffusion can significantly enhance the item-level bundle embedding generation ability of BR models. Lastly, extensive experiments on four datasets from different domains show that RDiffBR enhances the performance of Recall and NDCG of the BR backbone models by up to 23%, while only increasing the training time by about 4%.

## Related Work

**Bundle Recommendation.** Bundle recommendation aims to recommend a bundle of items related to a specific theme. Early work employed neural network methods for bundle recommendation (Cao et al. 2017; Chen et al. 2019a; He et al. 2019; Liu, Xie, and Lakshmanan 2014). Recently, graph neural networks have become the mainstream technology for bundle recommendation. We can categorize these graph bundle recommendation models based on these three types of representation learning strategies (Sun et al. 2024).

- **Unified learning.** The unified representation learning strategy aims to learn representations of users and bundles within a unified user-bundle-item graph (Chen et al. 2019b; Yu et al. 2022).
- **Separate learning.** The main idea of this strategy is to first perform representation learning and preference prediction upon two views individually, and then

fuse two view-specific predictions to obtain the final prediction (Chang et al. 2020; Tan et al. 2021; Wei et al. 2023; Ke et al. 2023b; Li et al. 2023).

- **Cooperative learning.** This approach learns user and bundle representations from bundle view and item view and uses techniques such as contrastive learning (Zhao et al. 2022; Ma et al. 2022, 2024c; Zou et al. 2024), knowledge distillation (Ke et al. 2023a; Ren et al. 2023) or diffusion (Li 2025) to achieve cross views mutual enhancement. However, most of bundle recommendation lack the ability to adapt to item-level dynamic variability. It is necessary to study our task for real-world application scenarios.

- **Diffusion Models for Recommendation.** The emerging diffusion model in generative models, compared with GANs and VAEs, possess stronger stability and representation. Diffusion model plays three main roles in recommendation systems (Lin et al. 2024a):

  - **Diffusion for data augmentation.** This approach generating realistic samples to expand the original training data, such as sequence data (Wu et al. 2023; Ma et al. 2024a) and side information (Jiang et al. 2024b,a).
  - **Diffusion for representation enhancement.** Diffusion models converts the input data into embeddings to address some key challenges of recommendation. Lots of studies have been proposed such as denoising the implicit user feedback (Zhao et al. 2024), sequential recommendation (Li, Sun, and Li 2024), out-of-distribution issue (Zhao et al. 2025) and multi-modal learning (Ma et al. 2024b).
  - **Diffusion as recommendation model.** This approach can be classified into three categories: collaborative recommendation (Wang et al. 2023; Hou, Park, and Shin 2024), context-aware recommendation (Li, Xia, and Huang 2024; Qin et al. 2024), and other applications (Lin et al. 2024b). However, a straightforward application of diffusion models on bundle recommendation is insufficient to address the issue of item-level dynamic variability. There are mainly two reasons: the loss targets of diffusion and bundle recommendation are inconsistent and the approximator needs to be designed based on item-level dynamic variability scenario.

## Methodology

While it is intuitive to leverage side information—such as item attributes—to enrich the feature space, doing so typically demands non-trivial architectural changes to the original model. In contrast, our RDiffBR operates without any external information and leaves the base architecture untouched. The technical details are presented below.

### Problem Definition

$U$  ( $|U| = M$ ) is denoted as the user set,  $B$  ( $|B| = N$ ) as the bundle set, and  $I$  ( $|I| = O$ ) as the item set. The user-bundle interactions, user-item interactions, and bundle-item affiliations are represented as  $X_{M \times N} = \{(u, b) \mid u \in U, b \in B\}$ ,  $Y_{M \times O} = \{(u, i) \mid u \in U, i \in I\}$  and  $Z_{N \times O} = \{(b, i) \mid b \in B, i \in I\}$  with a binary value at each entry respectively. The value 1 indicates an observed interaction of user-item or user-bundle, or affiliation of bundle-item. The item-level dynamic variability problem can be formulated as follows:

**Input:** The user-bundle interactions  $X_{M \times N}$ , user-item interactions  $Y_{M \times O}$ , and bundle-item affiliations  $Z_{N \times O}$ .

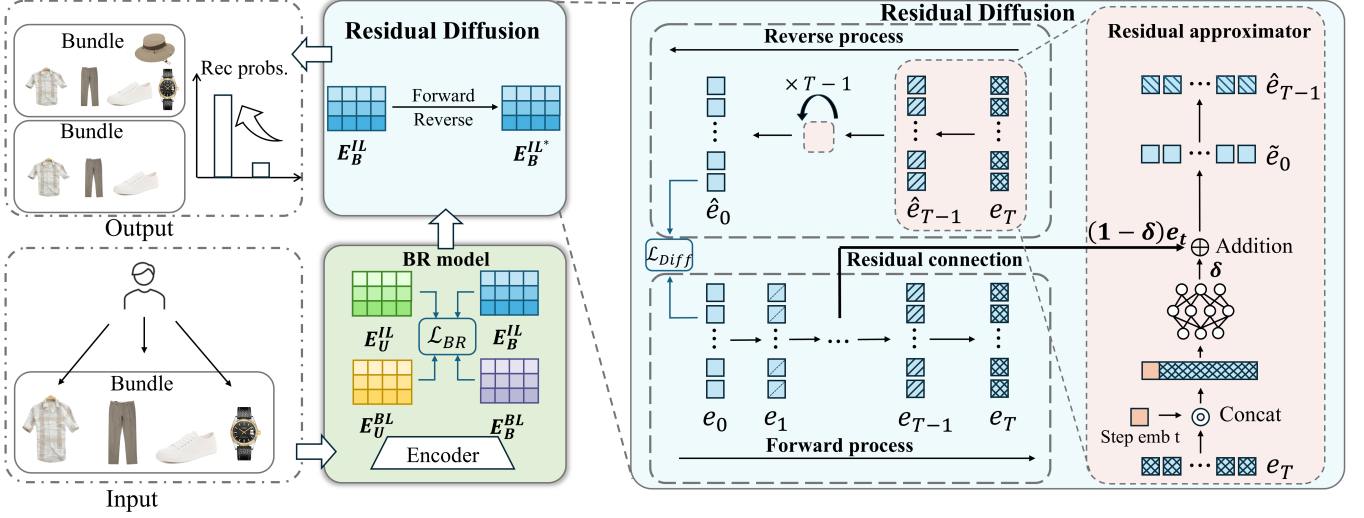


Figure 2: Overview of the proposed RDiffBR framework.

**Output:** A well-trained bundle recommendation estimates the probability  $\hat{y}_{ub} \in (0, 1)$  of a user  $u$  interacting with unseen bundles  $b$  based on  $Z^* = Z \cup \Delta^+ \setminus \Delta^-$  rather than  $Z$ . Here,  $\Delta^+ = \{(b, i) \mid (b, i) \notin Z\}$  and  $\Delta^- = \{(b, i) \mid (b, i) \in Z\}$  represent dynamic variability bundle-item affiliation. When an item is replaced by others, this can be treated as the joint operation of  $\Delta^+$  and  $\Delta^-$ .

### RDiffBR Framework

Considering the successful experience of diffusion as a generative model and the sparsity of input data, our RDiffBR framework introduces a residual diffusion model in the latent space of interactions. As shown in Figure 2, after BR model utilizes an encoder to get bundle-level user/bundle embeddings  $E_u^{BL}, E_b^{BL}$  and item-level embeddings user/bundle embeddings  $E_u^{IL}, E_b^{IL}$  from  $\{X, Y, Z\}$ . Since  $E_b^{IL}$  represents the bundle theme generated by the BR model, we iteratively process  $E_b^{IL}$  through a forward-backward process, jointly optimizing the BR model and the residual approximator module that we have designed based on the task. In the inference stage, RDiffBR applies this optimized forward-reverse process to  $E_b^{IL}$  from dynamic variability B-I affiliations  $Z^*$  to obtain new  $E_b^{IL}$ , enhancing BR model’s ability for this scenario.

#### Residual Diffusion

• **Forward process.** Given item-level bundle embeddings  $E_b^{IL}$  from a BR model’s encoding we can obtain a bundle’s item-level bundle embedding  $e_b^{IL}$  and set  $e_0 = e_b^{IL}$ . As shown in the bottom of residual diffusion module in Figure 2, Gaussian noises are continuously incorporated into  $e_0$  with adjustable scales and steps:

$$q(e_t|e_{t-1}) = \mathcal{N}(e_t; \sqrt{1 - \beta_t}e_{t-1}, \beta_t \mathbf{I}). \quad (1)$$

Based on the reparameterization (Ho, Jain, and Abbeel 2020) and the additivity of independent Gaussian noises, we

can directly obtain  $e_t$  from  $e_0$ .

$$q(e_t|e_0) = \mathcal{N}(e_t, \sqrt{\bar{\alpha}_t}e_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (2)$$

where  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ ,  $\alpha_t = 1 - \beta_t$ ,  $\beta_t \in (0, 1)$  controls the Gaussian noise scales added at each step  $t$ . To regulate the added noises in  $e_{1:T}$ , we follow (Wang et al. 2023) using a linear variance noise schedule in the forward process:

$$1 - \bar{\alpha}_t = s \cdot \left[ \alpha_{\min} + \frac{t-1}{T-1}(\alpha_{\max} - \alpha_{\min}) \right], \quad (3)$$

where  $T$  is the total forward step and  $s \in (0, 1)$  controls the noise scale. The condition  $\alpha_{\min} < \alpha_{\max} \in (0, 1)$  indicates the upper and lower limits of the added noise.

• **Reverse process.** After obtaining the noisy embedding  $e_T$  by forward process, we iteratively denoise this embedding in reverse processing as shown in the top of residual diffusion module in Figure 2. In each reverse step, a residual approximator module to denoise them.

$$p_\theta(\hat{e}_{t-1}|\hat{e}_t) = \mathcal{N}(\hat{e}_{t-1}; \mu_\theta(\hat{e}_t, t), \Sigma_\theta(\hat{e}_t, t)), \quad (4)$$

where  $\hat{e}_t$  is denoised embedding in the reverse step  $t$ ,  $\theta$  is the learnable parameter of residual approximator. This module is executed iteratively in the reverse process until reconstructing final embedding  $\hat{e}_0$

• **Residual approximator.** We design a residual approximator module based on our task to denoise the noisy embedding in each reverse step as shown in the right of residual diffusion module. The diffusion training aims to approximate the distribution  $p_\theta(\hat{e}_{t-1}|\hat{e}_t)$  with the tractable posterior distribution  $q(e_{t-1}|e_t, e_0)$  in the reverse process. Through Bayes’ theorem, we can derive:

$$q(e_{t-1}|e_t, e_0) \propto \mathcal{N}(e_{t-1}; \tilde{\mu}(e_t, e_0, t), \sigma^2(t)\mathbf{I}), \quad \text{where} \quad (5)$$

$$\begin{cases} \tilde{\mu}(e_t, e_0, t) = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}e_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1-\alpha_t)}{1-\bar{\alpha}_t}e_0 \\ \sigma^2(t) = \frac{(1-\alpha_t)(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}, \end{cases}$$

$\tilde{\mu}(e_t, e_0, t)$  and  $\sigma^2(t)I$  are the mean and covariance of  $q(e_{t-1}|e_t, e_0)$ . Following (Zhao et al. 2024) we can similarly factorize  $p_\theta(e_{t-1}|e_t)$ :

$$p_\theta(\hat{e}_{t-1}|\hat{e}_t) = \mathcal{N}(\hat{e}_{t-1}; \mu_\theta(\hat{e}_t, t), \tilde{\beta}_t I), \quad \text{where} \quad (6)$$

$$\mu_\theta(\hat{e}_t, t) = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \hat{e}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \tilde{e}_0,$$

$\tilde{e}_0$  is the predicted  $e_0$  since the distribution of  $e_0$  is unknown in the reverse process. We employ the multi-layer perceptron (MLP) to reconstruct  $e_0$  in the residual approximator. Specifically, for embedding  $e_t$  in the reverse step t:

$$\tilde{e}_0 = \delta f_\theta(\hat{e}_t, t) + (1 - \delta)e_{t'}, \quad (7)$$

where  $\tilde{e}_0$  is the predicted value of  $e_0$  by residual approximator with parameter  $\theta$ .  $\delta$  is residual connection weight.  $e_{t'}$  is noisy embedding in the forward step  $t'$ . Step information  $t$  is encoded by sinusoidal positional encoding (Ho, Jain, and Abbeel 2020), and these two inputs are concatenated as input.

**Loss Function.** To optimize the forward-reverse process, it is necessary to minimize the variational lower bound of the predicted embedding. According to the KL divergence based on the multivariate Gaussian distribution, the reconstruction loss of within one training iteration can be expressed as:

$$\begin{aligned} \mathcal{L}_r &= \mathbb{E}_{q(e_t|e_0)} [D_{\text{KL}}(q(e_{t-1}|e_t, e_0) \| p_\theta(e_{t-1}|e_t))] \\ &= \mathbb{E}_{q(e_t|e_0)} \left[ \frac{1}{2\sigma^2(t)} \left[ \|\mu_\theta(e_t, t) - \mu(e_t, e_0, t)\|_2^2 \right] \right]. \end{aligned} \quad (8)$$

Furthermore, by substituting Eq.(5) and Eq.(6) into Eq.(8) and simplifying (Wang et al. 2023) we can drive:

$$\mathcal{L}_{\text{diff}} = \mathbb{E}_{t \sim \mathcal{U}(1, T)} \mathbb{E}_q [\|e_0 - \hat{e}_t(e_t, e_{t'}, t)\|_2^2], \quad (9)$$

where  $\mathcal{U}$  means uniformly sampling  $t$  from  $\{1, 2, \dots, T\}$ . The final loss function of RDiffBR consists of: a bundle model loss and a reconstruction loss for reverse process. We have designed a loss balancing factor  $\lambda$  to adjust the weight of reconstruction loss and jointly train the bundle model and residual approximator together:

$$\mathcal{L} = \mathcal{L}_{BR \text{ model}} + \lambda \mathcal{L}_{\text{diff}}. \quad (10)$$

**Inference.** In the inference stage, we take the item-level bundle embedding  $e_0^*$  in the B-I dynamic variability scenario as input. Subsequently, RDiffBR continuously introduces noise into  $e_0^*$  in the sequence  $e_0^* \rightarrow e_1^* \rightarrow \dots \rightarrow e_{T'}^*$  for  $T'$  steps in the forward process and then we set  $\hat{e}_{T'} = e_{T'}^*$  to execute reverse process  $\hat{e}_{T'} \rightarrow \hat{e}_{T'-1} \rightarrow \dots \rightarrow \hat{e}_0$  to generate new item-level bundle embedding  $\hat{e}_0$ . The reverse ignores the variance and uses  $\hat{e}_{t-1} = \mu_\theta(\hat{e}_t, e_0^*, t)$  for directly inference following (Liang et al. 2018). We replace the original  $e_0^*$  with  $\hat{e}_0$ , and then conduct bundle recommendation following the original way of the bundle model. Finally the ranking of bundles is enhanced.

## Experiments

Our experiments aim to answer the following questions:

- **RQ1:** How does RDiffBR perform with backbone bundle models under various degrees of B-I variability?
- **RQ2:** What is the contribution of residual connection in RDiffBR to its overall performance?
- **RQ3:** How does the efficiency of RDiffBR?

Dataset	Youshu	NetEase	iFashion	MealRec <sup>+</sup> <sub>H</sub>
# U	8,039	18,528	53,897	1,575
# I	32,770	123,628	42,563	7,280
# B	4,771	22,864	27,694	3,817
# U-I	138,515	1,128,065	2,290,645	151,148
# U-B	51,377	302,303	1,679,708	46,767
# B-I	176,667	1,778,838	106,916	11,451
# Avg. I/B	37.03	77.8	3.86	3
Dens.B-I	0.11%	0.06%	0.01%	0.77%

Table 1: Statistics of datasets.

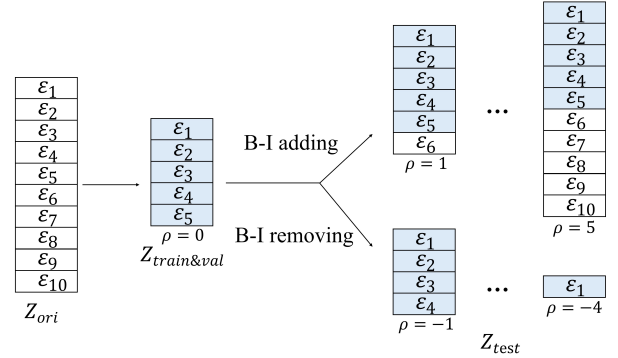


Figure 3: The processing of the B-I affiliation to match the scenario of B-I dynamic variability.

## Experimental Settings

• **Datasets.** Four publicly bundle datasets from different domains are summarized in Table 1. Youshu (Chen et al. 2019a) for book list recommendation; NetEase (Cao et al. 2017) for music playlist recommendation; iFashion (Chen et al. 2019b) for fashion outfit recommendation; In addition, we introduce another meal recommendation dataset MealRec+ (Li et al. 2024), where the meal consisted of individual course items is treated as bundle.

• **BR Backbones.** We implement RDiffBR on six representative BR models with official codes released by their authors. **BGCN** (Chang et al. 2020) utilizes graph convolutional networks to separately learn the representations of bundle view and item view. **HyperMBR** (Ke et al. 2023a) models two views in hyperbolic space to learn their accurate representations. **MIDGN** (Zhao et al. 2022) refines the user-item/bundle-item graph of different intents from multiple views. **CrossCBR** (Ma et al. 2022) improves two separate views by contrastive learning. **MultiCBR** (Ma et al. 2024c) first fuses the multi-view representations into a unified one, and then adopts a contrastive loss on the unified representations. **BundleGT** (Wei et al. 2023) learns the strategy-aware user and bundle representation in bundle recommendation.

• **Experimental Preprocess and Metrics.** According to previous works (Ma et al. 2022) (Li et al. 2024), the training, validation and testing sets of user-bundle interactions  $X$  are randomly splited with specific ratio. It is not feasible to simply add new bundles in bundle-item affiliations  $Z$  under the B-I variability assumptions. Therefore, we adjust existing bundle datasets to match this scenario: Let

Dataset	Youshu								NetEase											
	-4		-2		0		2		4		-4		-2		0		2		4	
$\rho$	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20
Metrics	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20
BGCN	.1897	.1047	.2254	.1278	.2367	.1380	.2214	.1247	.2138	.1204	.0581	.0309	.0634	.0333	.0643	.0339	.0646	.0341	.0622	.0332
+RDiffBR	.2178	.1278	.2391	.1415	.2453	.1466	.2470	.1423	.2407	.1396	.0614	.0321	.0657	.0340	.0662	.0343	.0668	.0345	.0649	.0340
%improv.	<b>14.83</b>	<b>22.15</b>	<b>6.06</b>	<b>10.74</b>	<b>3.61</b>	<b>6.25</b>	<b>11.57</b>	<b>14.11</b>	<b>12.60</b>	<b>16.03</b>	<b>5.56</b>	<b>3.98</b>	<b>3.63</b>	<b>1.99</b>	<b>2.89</b>	<b>1.14</b>	<b>3.47</b>	<b>1.07</b>	<b>4.42</b>	<b>2.54</b>
MIDGN	.2388	.1373	.2557	.1472	.2598	.1496	.2578	.1495	.2566	.1489	.0540	.0278	.0625	.0330	.0636	.0337	.0639	.0337	.0640	.0337
+RDiffBR	.2498	.1441	.2621	.1498	.2645	.1515	.2652	.1517	.2643	.1516	.0568	.0293	.0649	.0340	.0661	.0347	.0667	.0348	.0672	.0350
%improv.	<b>4.62</b>	<b>4.98</b>	<b>2.53</b>	<b>1.72</b>	<b>1.79</b>	<b>1.32</b>	<b>2.87</b>	<b>1.49</b>	<b>3.00</b>	<b>1.82</b>	<b>5.21</b>	<b>5.27</b>	<b>3.89</b>	<b>2.91</b>	<b>3.96</b>	<b>2.92</b>	<b>4.45</b>	<b>3.25</b>	<b>4.95</b>	<b>4.14</b>
HyperMBR	.1464	.0953	.2126	.1290	.2649	.1581	.2167	.1285	.1897	.1162	.0588	.0313	.0675	.0354	.0717	.0377	.0689	.0363	.0661	.0345
+RDiffBR	.1789	.1109	.2229	.1327	.2669	.1594	.2280	.1347	.2128	.1275	.0627	.0332	.0694	.0369	.0739	.0386	.0704	.0369	.0667	.0350
%improv.	<b>22.20</b>	<b>16.42</b>	<b>4.85</b>	<b>2.87</b>	<b>0.75</b>	<b>0.83</b>	<b>5.20</b>	<b>4.87</b>	<b>12.14</b>	<b>9.75</b>	<b>6.67</b>	<b>5.97</b>	<b>2.88</b>	<b>4.41</b>	<b>3.07</b>	<b>2.51</b>	<b>2.25</b>	<b>1.79</b>	<b>0.85</b>	<b>1.50</b>
CrossCBR	.2677	.1584	.2763	.1626	.2776	.1640	.2788	.1649	.2800	.1656	.0715	.0386	.0776	.0419	.0794	.0431	.0802	.0436	.0798	.0434
+RDiffBR	.2744	.1640	.2814	.1659	.2827	.1663	.2816	.1664	.2816	.1662	.0772	.0419	.0818	.0436	.0826	.0443	.0829	.0444	.0826	.0445
%improv.	<b>2.52</b>	<b>3.57</b>	<b>1.84</b>	<b>2.02</b>	<b>1.85</b>	<b>1.42</b>	<b>1.01</b>	<b>0.92</b>	<b>0.56</b>	<b>0.40</b>	<b>7.99</b>	<b>8.50</b>	<b>5.36</b>	<b>4.04</b>	<b>4.09</b>	<b>2.67</b>	<b>3.43</b>	<b>1.70</b>	<b>3.50</b>	<b>2.51</b>
MultiCBR	.2770	.1649	.2769	.1649	.2791	.1652	.2767	.1641	.2770	.1648	.0843	.0463	.0867	.0472	.0876	.0475	.0880	.0479	.0879	.0479
+RDiffBR	.2831	.1674	.2824	.1677	.2824	.1673	.2825	.1669	.2842	.1675	.0859	.0469	.0879	.0478	.0881	.0480	.0883	.0482	.0893	.0487
%improv.	<b>2.20</b>	<b>1.55</b>	<b>2.56</b>	<b>1.71</b>	<b>1.19</b>	<b>1.25</b>	<b>2.10</b>	<b>1.71</b>	<b>2.60</b>	<b>1.61</b>	<b>1.91</b>	<b>1.32</b>	<b>1.37</b>	<b>1.27</b>	<b>0.57</b>	<b>1.18</b>	<b>0.40</b>	<b>0.75</b>	<b>1.60</b>	<b>1.52</b>
BundleGT	.2687	.1590	.2872	.1683	.2877	.1694	.2873	.1697	.2837	.1679	.0729	.0367	.0869	.0465	.0874	.0468	.0874	.0468	.0871	.0465
+RDiffBR	.2829	.1653	.2927	.1708	.2922	.1707	.2939	.1712	.2903	.1702	.0759	.0384	.0900	.0479	.0896	.0479	.0897	.0477	.0895	.0476
%improv.	<b>5.27</b>	<b>3.96</b>	<b>1.90</b>	<b>1.49</b>	<b>1.57</b>	<b>0.79</b>	<b>2.30</b>	<b>0.93</b>	<b>2.31</b>	<b>1.40</b>	<b>4.06</b>	<b>4.75</b>	<b>3.49</b>	<b>3.16</b>	<b>2.52</b>	<b>2.42</b>	<b>2.67</b>	<b>1.86</b>	<b>2.76</b>	<b>2.37</b>

Dataset	iFashion								MealRec <sub>H</sub> <sup>+</sup>											
	-4		-2		0		2		4		-4		-2		0		2		4	
$\rho$	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20
Metrics	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20	R20	N20
BGCN	.0473	.0331	.0560	.0395	.0602	.0426	.0607	.0428	.0604	.0427	.1068	.1888	.1213	.2094	.1325	.2247	.1326	.2250	.1375	.2323
+RDiffBR	.0488	.0341	.0568	.0398	.0606	.0428	.0608	.0429	.0607	.0429	.1180	.2078	.1293	.2180	.1408	.2401	.1387	.2371	.1388	.2384
%improv.	<b>3.09</b>	<b>2.91</b>	<b>1.55</b>	<b>0.58</b>	<b>0.61</b>	<b>0.45</b>	<b>0.16</b>	<b>0.23</b>	<b>0.57</b>	<b>0.47</b>	<b>10.07</b>	<b>10.55</b>	<b>4.08</b>	<b>6.53</b>	<b>6.82</b>	<b>6.26</b>	<b>5.38</b>	<b>4.57</b>	<b>2.63</b>	<b>0.92</b>
MIDGN	.0372	.0257	.0549	.0380	.0672	.0482	.0428	.0279	.0441	.0294	.0820	.1564	.1017	.1816	.1121	.1908	.1129	.1919	.1115	.1880
+RDiffBR	.0426	.0290	.0628	.0437	.0763	.0553	.0534	.0348	.0545	.0363	.0826	.1604	.1075	.1930	.1179	.2050	.1145	.2012	.1146	.2027
%improv.	<b>14.55</b>	<b>12.83</b>	<b>14.36</b>	<b>14.87</b>	<b>13.53</b>	<b>14.65</b>	<b>24.66</b>	<b>24.83</b>	<b>23.48</b>	<b>23.77</b>	<b>2.60</b>	<b>0.82</b>	<b>6.29</b>	<b>5.69</b>	<b>7.44</b>	<b>5.13</b>	<b>4.87</b>	<b>1.48</b>	<b>7.82</b>	<b>2.79</b>
HyperMBR	.0448	.0336	.0687	.0494	.0815	.0580	.0773	.0545	.0744	.0521	.0859	.1498	.0939	.1595	.1069	.1778	.1057	.1781	.1106	.1887
+RDiffBR	.0462	.0349	.0692	.0504	.0830	.0595	.0797	.0567	.0773	.0550	.0874	.1585	.0995	.1760	.1122	.1944	.1205	.1988	.1247	.2144
%improv.	<b>3.21</b>	<b>4.09</b>	<b>0.81</b>	<b>2.01</b>	<b>1.75</b>	<b>2.69</b>	<b>3.01</b>	<b>4.01</b>	<b>3.92</b>	<b>5.42</b>	<b>5.82</b>	<b>1.70</b>	<b>10.33</b>	<b>5.90</b>	<b>9.34</b>	<b>4.94</b>	<b>11.61</b>	<b>14.00</b>	<b>13.62</b>	<b>12.82</b>
CrossCBR	.0624	.0491	.0877	.0683	.0999	.0772	.1046	.0799	.1081	.0831	.0612	.1048	.0819	.1417	.0938	.1688	.1084	.2041	.1504	.2599
+RDiffBR	.0720	.0554	.0928	.0716	.1027	.0794	.1060	.0812	.1088	.0834	.0681	.1203	.0926	.1531	.1026	.1788	.1241	.2243	.1611	.2669
%improv.	<b>15.36</b>	<b>12.87</b>	<b>5.75</b>	<b>4.78</b>	<b>2.83</b>	<b>2.82</b>	<b>1.31</b>	<b>1.64</b>	<b>0.60</b>	<b>0.32</b>	<b>14.80</b>	<b>11.40</b>	<b>8.01</b>	<b>13.15</b>	<b>5.95</b>	<b>9.37</b>	<b>9.89</b>	<b>14.48</b>	<b>2.69</b>	<b>7.12</b>
MultiCBR	.0672	.0588	.1080	.0905	.1297	.1064	.1343	.1101	.1381	.1133	.0989	.1782	.1147	.1841	.1305	.2092	.1476	.2376	.1625	.2592
+RDiffBR	.0694	.0608	.1094	.0916	.1313	.1080	.1350	.1111	.1385	.1139	.1066	.1889	.1217	.1945	.1375	.2164	.1524	.2436	.1633	.2598
%improv.	<b>3.29</b>	<b>3.31</b>	<b>1.23</b>	<b>1.16</b>	<b>1.20</b>	<b>1.47</b>	<b>0.53</b>	<b>0.97</b>	<b>0.27</b>	<b>0.58</b>	<b>6.04</b>	<b>7.83</b>	<b>5.63</b>	<b>6.10</b>	<b>3.44</b>	<b>5.40</b>	<b>2.51</b>	<b>3.19</b>	<b>0.25</b>	<b>0.50</b>
BundleGT	.0444	.0352	.0690	.0517	.0954	.0724	.0899	.0653	.0960	.0707	.0732	.1281	.0923	.1556	.1093	.1761	.1217	.2134	.1579	.2585
+RDiffBR	.0456	.0361	.0701	.0530	.0970	.0740	.0909	.0660	.0970	.0712	.0863	.1512	.1082	.1770	.1315	.2077	.1414	.2346	.1707	.2733
%improv.	<b>2.68</b>	<b>2.58</b>	<b>1.64</b>	<b>2.34</b>	<b>1.61</b>	<b>2.10</b>	<b>1.02</b>	<b>1.04</b>	<b>1.10</b>	<b>0.71</b>	<b>18.03</b>	<b>17.88</b>	<b>13.75</b>	<b>17.20</b>	<b>17.97</b>	<b>20.35</b>	<b>9.91</b>	<b>16.21</b>	<b>5.73</b>	<b>8.05</b>

Table 2: Recall@20 and NDCG@20 performance comparison between backbone models and our RDiffBR on four datasets ( $\rho=-4, -2, 0, 2, 4$ ), while  $\rho < 0$  indicates  $Z_{\text{test}} = Z_{\text{train}\&\text{val}} \setminus \Delta^-$ ,  $\rho > 0$  indicates adding set  $Z_{\text{test}} = Z_{\text{train}\&\text{val}} \cup \Delta^+$ , and  $\rho = 0$  indicates  $Z_{\text{train}\&\text{val}}$  invariability. Improv. indicates the relative improvement of a backbone model+RDiffBR over the itself. The performance of  $K = \{10, 40, 80\}$  is consistent with  $K = 20$ .

$\{\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_{10}\}$  denote a uniform partition of  $Z$  into ten disjoint subsets. Bundle models are trained and validated on  $Z_{\text{train}\&\text{val}} = \cup_{i=1}^5 \mathcal{E}_i$ . In testing phase, we design  $Z_{\text{test}}$  by adjusting  $Z_{\text{train}\&\text{val}}$  with adding or removing subsets:

$$Z_{\text{test}} = \begin{cases} Z_{\text{train}\&\text{val}} \cup \left( \cup_{i=6}^{5+\rho} \mathcal{E}_i \right) & \text{if } \rho > 0 \\ Z_{\text{train}\&\text{val}} & \text{if } \rho = 0 \\ Z_{\text{train}\&\text{val}} \setminus \left( \cup_{i=6-|\rho|}^5 \mathcal{E}_i \right) & \text{if } \rho < 0 \end{cases} \quad (11)$$

Where  $\cup$  indicates set union and  $\setminus$  indicates set difference.  $\rho \in \{-4, -3, -2, -1, 0, 1, 2, 3, 4, 5\}$  reflects level of B-I variation.  $\rho > 0$  indicates adding set  $\Delta^+$  to  $Z_{\text{test}}$ ,  $\rho < 0$  indicates removing set  $\Delta^-$  from  $Z_{\text{test}}$  and  $\rho = 0$  indicates  $Z_{\text{test}}$  invariability. The well-trained bundle models predict user-bundle interactions  $X$  on  $Z_{\text{test}}$ . The process-

ing of the B-I affiliation is shown in Figure 3. In this way, we can control the degree of B-I affiliations variation and simulate the scenario of B-I dynamic variability. Recall@K and NDCG@K are used as the evaluation metrics, where  $K \in \{10, 20, 40, 80\}$ .

• **Hyperparameter Settings** The hyperparameters of six backbone models are set by following their publications and RDiffBR is optimized with the Adam optimizer with a learning rate of 0.001. The time steps  $T$  and sampling steps  $T'$  are set to values ranging from [1,200]. The loss balance factor  $\lambda$  is chosen from (0,5), while the residual connection weight  $\delta$  is selected from (0,1). The number of MLP layers is selected within 1, 2, 4, 8. Its hidden sizes are from [8,4096]. Details of other hyperparameters can be found in our released code. All the models are trained using two GTX 4090.

## Overall Performance (RQ1)

The following observations are from Table 2.

- The performance of BR backbones when meeting item-level dynamic variability.** When  $\rho < 0$  ( $\mathbf{Z}_{\text{test}}$  removes  $\Delta^-$ ), the performance of bundle models is generally lower than that under  $\rho = 0$  in most cases. The greater the reduction in B-I, the more significant the performance degradation. When  $\rho > 0$  ( $\mathbf{Z}_{\text{test}}$  adds  $\Delta^+$ ), some models exhibit performance decline, others maintain comparable performance, while a few show improvement. This indicates that BR models possess capability to handle B-I variations to some extent, yet they tend to fail under drastic reductions of B-I. Notably, HyperMBR performs the worst, where both B-I adding and removing frequently lead to performance deterioration. This phenomenon may stem from hyperbolic space embeddings being more sensitive to distance errors compared to Euclidean space. In contrast, models like MIDGN, CrossCBR, and MultiCBR exhibit relatively stable performance, potentially attributed to their adoption of contrastive learning that leverages multi-view information through cooperative learning. Additionally, BundleGT demonstrates enhanced adaptability for dynamic B-I variations by modeling B-I affiliation via attention mechanisms.

- The performance of RDiffBR dealing with item-level dynamic variability.** Our experiments reveal that RDiffBR consistently enhances the performance of six backbone BR models across four datasets from different domains under different B-I variation levels. When  $\rho < 0$  ( $\mathbf{Z}_{\text{test}}$  removes  $\Delta^-$ ), The performance of BR models+RDiffBR outperform BR models with the improvement up to 22.15% at most. When  $\rho > 0$  ( $\mathbf{Z}_{\text{test}}$  adds  $\Delta^+$ ), RDiffBR also enhances the performance of BR models, reaching up to 24.83% at most. In some cases, RDiffBR can even enable BR models to outperform its performance under  $\rho = 0$ , which highlights the powerful generative capability of RDiffBR. Furthermore, performance improvements is observed even under  $\rho = 0$ , which demonstrates RDiffBR has certain effect even in the B-I remain scenario. Overall, RDiffBR can enhance the performance of BR models with different architectures and metric spaces under different variation scenarios. This remarkable performance improvement can be attributed to RDiffBR’s generation capability, which reconstructs item-level bundle embeddings in B-I change scenarios to better align with bundle themes rather than overfitting to the B-I affiliations observed during training.

## Ablation Study (RQ2)

To assess whether residual connections contribute to improving model performance, we set a variant **RDiffBR w/o R that omits the residual connection** in Eq.(7) and conduct a comprehensive ablation study on three BR models across three datasets, with the results shown in Table 3.

From this table, we can draw the following conclusions: In most cases, the performance of using RDiffBR is superior to that of not using it or using RDiffBR w/o R. Using RDiffBR w/o R may lead to worse performance than RDiffBR, and in some cases it may even be worse than only using backbone BR model. Figure 4 shows the visualiza-

$\rho = 3$	Youshu		NetEase		MealRec <sup>+</sup> <sub>H</sub>	
	R@20	N@20	R@20	N@20	R@20	N@20
BGCN	.2129	.1197	.0632	.0336	.2298	.1350
+RDiffBR	<b>.2394</b>	<b>.1374</b>	<b>.0660</b>	<b>.0343</b>	<b>.2390</b>	<b>.1404</b>
+RDiffBR w/o R	<u>.2337</u>	<u>.1229</u>	.0627	.0321	.2089	.1221
CrossCBR	.2795	.1653	.0801	.0435	.2395	.1295
+RDiffBR	<b>.2816</b>	<b>.1664</b>	<b>.0827</b>	<b>.0445</b>	<b>.2474</b>	<b>.1440</b>
+RDiffBR w/o R	.2763	.1639	.0798	.0426	.1843	.1052
BundleGT	.2858	.1691	.0871	.0466	.2394	.1416
+RDiffBR	<b>.2925</b>	<b>.1711</b>	<b>.0895</b>	<b>.0477</b>	<b>.2589</b>	<b>.1593</b>
+RDiffBR w/o R	.2519	.1454	.0520	.0267	.2312	.1417

Table 3: Ablation study where bold indicates the best method, and underlined indicates the second best method ( $\rho = 3$ ).




Figure 4: Blue is distribution  $E_B^{IL}$  under  $\rho = 0$ . Red is distribution  $E_B^{IL}$  under  $\rho = 3$ .

tion results of CrossCBR, CrossCBR+RDiffBR and CrossCBR+RDiffBR w/o R under B-I variation levels  $\rho = 3$  for item-level bundle embeddings  $E_B^{IL}$  in Youshu dataset. It can be observed that  $E_B^{IL}$  generated by CrossCBR after B-I change is more dispersed than original  $E_B^{IL}$  and more compact  $E_B^{IL}$  distribution which retains the basic shape can be obtained after using CrossCBR+RDiffBR. This indicates that  $E_B^{IL}$  of the BR model cannot effectively represent the bundle theme distribution due to the B-I change and RDiffBR enables the model to obtain more bundle information at the item-level rather than overfitting to the existing B-I affiliations. While using CrossCBR+RDiffBR w/o R,  $E_B^{IL}$  distribution changes little so the effect is limited, which is consistent with the experimental results obtained in Table 3. Overall, RDiffBR outperforms backbone models, highlighting the effectiveness of RDiffBR. Residual connection plays an important role in RDiffBR, and merely using a diffusion model cannot generate better item-level bundle embeddings (The performance of  $\rho = -3$  is consistent with  $\rho = 3$ , see details in extended version).

Although different B-I affiliations are used during model testing, the bundle theme remains, which means items within a bundle remain to follow the similar distribution before and after the change. This can also be verified by Figure 4 that the shapes of two  $E_B^{IL}$  distributions under different B-I variation  $\rho$  are similar. Therefore, with using RDiffBR, our task can be approached without considering the Out-Of-Distribution (OOD) where the distribution of the training data and the test data is inconsistent. In extreme cases where the distribution of  $E_B^{IL}$  changes completely, the problem becomes equivalent to cold-start bundle recommenda-




Figure 5: The impact of loss balance factor  $\lambda$ .




Figure 6: The impact of residual weight  $\delta$ .

7249		25380		B-I change		25380	
6909		19793		→		19793	
28999 - bundle 10		23154				23154 - bundle 10	
29520		29520				29520	
item	$\cos(E_{B=10}^{IL}, E_I^{IL})$	item	$\cos(E_{B=10}^{IL}, E_I^{IL})$	item	$\cos(E_{B=10}^{IL}, E_I^{IL})$	item	$\cos(E_{B=10}^{IL}, E_I^{IL})$
19793	0.983	2105	0.945	19793	0.970	26391	0.966
23154	0.981	19793	0.943	26391	0.966	21811	0.949
28999	0.978	26391	0.934	21811	0.949	23154	0.942
13972	0.975	3331	0.925	23154	0.942	32222	0.941
20999	0.974	14200	0.918	32222	0.941		
CrossCBR		CrossCBR + RDiffBR w/o R		CrossCBR + RDiffBR			

Figure 7: A case study for Youshu under  $\rho = -3$ . The highlighted items are those before the bundle change.

tion (Jeon et al. 2024; Bui et al. 2025) that the original bundle id cannot be directly associated.

### Impact of Hyperparameters

In order to achieve a more detailed understanding, two key hyperparameters are studied, i.e., the loss balance factor  $\lambda$  and the residual weight  $\delta$ . The results of our hyperparameter study are presented in Figure 5-6.

- **Loss balance factor  $\lambda$ .** The loss balance factor  $r$  affects the performance of RDiffBR, which adjusts the emphasis between the recommendation task and the reconstruction task. A high value of  $r$  may neglect the core recommendation and thereby compromise the overall performance.
- **Residual weight  $\delta$ .** The appropriate selection of residual weights is crucial for embedding generation, this is because too high or too low values will lead to a performance decline.

### Case Study

In this section, we analyze the real data instances as show in Figure 7. We randomly select a bundle from Youshu and calculate cosine similarity between item-level bundle embeddings  $E_B^{IL}$  and all items under B-I variation levels  $\rho = -3$ . We listed the top 5 items with the highest similarity scores and their corresponding similarity values. It can be seen

Dataset	Youshu	NetEase	iFashion	MealRec <sup>+</sup> <sub>H</sub>
BGCN	1.46	23.5	68.3	1.07
+RDiffBR	1.51	24.5	71.4	1.09
Extra cost	3.42%	4.25%	4.53%	1.86%
CrossCBR	1.06	8.70	45.4	1.10
+RDiffBR	1.11	9.10	46.2	1.14
Extra cost	4.71%	4.59%	1.76%	3.63%
BundleGT	22.7	426.2	118.4	1.23
+RDiffBR	23.3	434.8	121.1	1.29
Extra cost	2.64%	2.01%	2.28%	4.87%

Table 4: The statistics of one-epoch training time (seconds) for three backbone models and using RDiffBR on them.

that the similarity of items (marked in red) that appeared within the bundle will be higher when only using CrossCBR, resulting in insufficient generalization ability. When using RDiffBR w/o R, the model cannot well capture the main items of the bundle. When using RDiffBR, the model can not only express the main items of the bundle but also introduce other items based on the theme to enhance the generalization ability.

### Training Efficiency Study (RQ3)

In order to evaluate the additional computing time of RDiffBR, we compare the one-epoch training time (seconds) of using BR model and BR model+RDiffBR on four datasets (The training rounds are the same for both modes). We record five consecutive training epochs and averaged them to obtain the one-epoch training time, as shown in Table 4. It shows that RDiffBR only increased training time about 4%.

From the perspective of time complexity, taking lightweight backbone CrossCBR as an example, the time complexity of graph learning is  $\mathcal{O}(|E|Kd)$ , 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. However, the diffusion process in RDiffBR costs  $\mathcal{O}(|B|d) + \mathcal{O}(|B|d)d'n$ , where  $|B|$  is the number of bundles,  $d'$  is the embedding size of MLP in approximator,  $n$  is MLP layer number. It shows that  $|E|$  is much greater than  $|B|$  (e.g., in iFashion  $|E| = 1679708 \gg |B| = 27694$ ). Therefore, our RDiffBR only increases training time about 4% on two GTX 4090.

### Conclusion and Future Work

In this work, to address the issue of item-level dynamic variability, we propose RDiffBR, a model-agnostic generative framework that can help BR models adapt to this scenario. A number of experiments show RDiffBR enables backbone BR models to achieve significant performance improvement under dynamic changes of B-I. Future research will consider extending RDiffBR as a plug-and-play module for well-performing pre-trained bundle recommenders and evaluating its applicability to cold-start bundle recommendation.

## Acknowledgments

This work is partially supported by NSFC, China (No.62276196), the research grant (RGPIN-2020-07157) from the Natural Science and Engineering Research Council (NSERC) of Canada and York Research Chairs (YRC) program. The authors also appreciate the anonymous reviewers for their valuable comments and constructive suggestions that greatly helped to improve the quality of the paper.

## References

- Bui, T.-N.; Nguyen, H.-S.; Nguyen, C.-V. T.; Le, H.-Q.; and Le, D.-T. 2025. Personalized Diffusion Model Reshapes Cold-Start Bundle Recommendation. In *Companion Proceedings of the ACM on Web Conference*, 3088–3091.
- Cao, D.; Nie, L.; He, X.; Wei, X.; Zhu, S.; and Chua, T.-S. 2017. Embedding Factorization Models for Jointly Recommending Items and User Generated Lists. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 585–594.
- Chang, J.; Gao, C.; He, X.; Jin, D.; and Li, Y. 2020. Bundle Recommendation with Graph Convolutional Networks. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1673–1676.
- Chen, L.; Liu, Y.; He, X.; Gao, L.; and Zheng, Z. 2019a. Matching User with Item Set: Collaborative Bundle Recommendation with Deep Attention Network. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2095–2101.
- Chen, W.; Huang, P.; Xu, J.; Guo, X.; Guo, C.; Sun, F.; Li, C.; Pfadler, A.; Zhao, H.; and Zhao, B. 2019b. POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2662–2670.
- He, Y.; Wang, J.; Niu, W.; and Caverlee, J. 2019. A Hierarchical Self-attentive Model for Recommending User-generated Item Lists. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 1481–1490.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, 6840–6851.
- Hou, Y.; Park, J.; and Shin, W. 2024. Collaborative Filtering Based on Diffusion Models: Unveiling the Potential of High-Order Connectivity. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1360–1369.
- Jeon, H.; Lee, J.; Yun, J.; and Kang, U. 2024. Cold-start Bundle Recommendation via Popularity-based Coalescence and Curriculum Heating. In *Companion Proceedings of the ACM on Web Conference*, 3277–3286.
- Jiang, Y.; Xia, L.; Wei, W.; Luo, D.; Lin, K.; and Huang, C. 2024a. DiffMM: Multi-Modal Diffusion Model for Recommendation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 7591–7599.
- Jiang, Y.; Yang, Y.; Xia, L.; and Huang, C. 2024b. DiffKG: Knowledge Graph Diffusion Model for Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 313–321.
- Ke, H.; Li, L.; Wang, P.; Yuan, J.; and Tao, X. 2023a. Hyperbolic Mutual Learning for Bundle Recommendation. In *International Conference on Database Systems for Advanced Applications*, 417–433.
- Ke, H.; Li, L.; Wang, P.; Yuan, J.; and Tao, X. 2023b. Tree-Like Interaction Learning for Bundle Recommendation. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 1–5.
- Li, J. 2025. Disentangled Contrastive Bundle Recommendation with Conditional Diffusion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 12067–12075.
- Li, M.; Li, L.; Tao, X.; and Huang, J. X. 2024. MealRec+: A Meal Recommendation Dataset with Meal-Course Affiliation for Personalization and Healthiness. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 564–574.
- Li, M.; Li, L.; Tao, X.; Xie, Q.; and Yuan, J. 2023. Category-Wise Meal Recommendation. In *International Conference on Neural Information Processing*, 282–294. Springer.
- Li, Z.; Sun, A.; and Li, C. 2024. DiffuRec: A Diffusion Model for Sequential Recommendation. *ACM Trans. Inf. Syst.*, 42(3): 66:1–66:28.
- Li, Z.; Xia, L.; and Huang, C. 2024. RecDiff: Diffusion Model for Social Recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 1346–1355.
- Liang, D.; Krishnan, R. G.; Hoffman, M. D.; and Jebara, T. 2018. Variational Autoencoders for Collaborative Filtering. In *Proceedings of the 27th World Wide Web Conference*, 689–698.
- Lin, J.; Liu, J.; Zhu, J.; Xi, Y.; Liu, C.; Zhang, Y.; Yu, Y.; and Zhang, W. 2024a. A Survey on Diffusion Models for Recommender Systems. *arXiv preprint arXiv:2409.05033*.
- Lin, X.; Chen, X.; Wang, C.; Shu, H.; Song, L.; Li, B.; and Jiang, P. 2024b. Discrete Conditional Diffusion for Reranking in Recommendation. In *Companion Proceedings of the ACM Web Conference*, 161–169.
- Liu, Y.; Xie, M.; and Lakshmanan, L. V. 2014. Recommending User Generated Item Lists. In *Proceedings of the 8th ACM Conference on Recommender systems*, 185–192.
- Ma, H.; Xie, R.; Meng, L.; Chen, X.; Zhang, X.; Lin, L.; and Kang, Z. 2024a. Plug-In Diffusion Model for Sequential Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 8886–8894.
- Ma, H.; Yang, Y.; Meng, L.; Xie, R.; and Meng, X. 2024b. Multimodal Conditioned Diffusion Model for Recommendation. In *Companion Proceedings of the ACM on Web Conference*, 1733–1740.
- Ma, Y.; He, Y.; Wang, X.; Wei, Y.; Du, X.; Fu, Y.; and Chua, T. 2024c. MultiCBR: Multi-view Contrastive Learning for Bundle Recommendation. *ACM Trans. Inf. Syst.*, 42(4): 100:1–100:23.

Ma, Y.; He, Y.; Zhang, A.; Wang, X.; and Chua, T.-S. 2022. CrossCBR: Cross-view Contrastive Learning for Bundle Recommendation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1233–1241.

Qin, Y.; Wu, H.; Ju, W.; Luo, X.; and Zhang, M. 2024. A Diffusion Model for POI Recommendation. *ACM Trans. Inf. Syst.*, 42(2): 54:1–54:27.

Ren, Y.; Haonan, Z.; Fu, L.; Wang, X.; and Zhou, C. 2023. Distillation-Enhanced Graph Masked Autoencoders for Bundle Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1660–1669.

Sun, M.; Li, L.; Li, M.; Tao, X.; Zhang, D.; Wang, P.; and Huang, J. X. 2024. A Survey on Bundle Recommendation: Methods, Applications, and Challenges. *arXiv preprint arXiv:2411.00341*.

Tan, M.; Chen, W.; Wang, W.; Liu, A.; and Zhao, L. 2021. Intention-oriented Hierarchical Bundle Recommendation with Preference Transfer. In *2021 IEEE International Conference on Web Services*, 107–116. IEEE.

Wang, W.; Xu, Y.; Feng, F.; Lin, X.; He, X.; and Chua, T. 2023. Diffusion Recommender Model. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 832–841.

Wei, Y.; Liu, X.; Ma, Y.; Wang, X.; Nie, L.; and Chua, T.-S. 2023. Strategy-aware Bundle Recommender system. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1198–1207.

Wu, Z.; Wang, X.; Chen, H.; Li, K.; Han, Y.; Sun, L.; and Zhu, W. 2023. Diff4Rec: Sequential Recommendation with Curriculum-scheduled Diffusion Augmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, 9329–9335.

Yu, Z.; Li, J.; Chen, L.; and Zheng, Z. 2022. Unifying Multi-associations Through Hypergraph for Bundle Recommendation. *Knowledge-Based Systems*, 255: 109755.

Zhao, C.; Yang, E.; Liang, Y.; Lan, P.; Liu, Y.; Zhao, J.; Guo, G.; and Wang, X. 2025. Graph Representation Learning via Causal Diffusion for Out-of-Distribution Recommendation. *arXiv preprint arXiv:2408.00490*.

Zhao, J.; Wang, W.; Xu, Y.; Sun, T.; Feng, F.; and Chua, T. 2024. Denoising Diffusion Recommender Model. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1370–1379.

Zhao, S.; Wei, W.; Zou, D.; and Mao, X. 2022. Multi-view Intent Disentangle Graph Networks for Bundle Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 4379–4387.

Zou, D.; Zhao, S.; Wei, W.; Mao, X.; Li, R.; Chen, D.; Fang, R.; and Fu, Y. 2024. Towards Hierarchical Intent Disentanglement for Bundle Recommendation. *IEEE Trans. Knowl. Data Eng.*, 36(7): 3556–3567.