

Decision-Aware Preference Modeling for Multi-Behavior Recommendation

Qingfeng Li¹, Wei Liu¹, Zaiqiao Meng¹ and Jian Yin^{1,2*}

¹School of Artificial Intelligence, Sun Yat-sen University, Zhuhai, China

²The Technology Innovation Center for Collaborative Applications of Natural Resources Data in GBA,
MNR

liqf37@mail2.sysu.edu.cn, liuw259@mail.sysu.edu.cn, zaiqiao.meng@gmail.com,
issjyin@mail.sysu.edu.cn

Abstract

In recommender systems, multi-behavior methods have demonstrated significant effectiveness in addressing issues such as data sparsity—challenges commonly encountered by traditional single-behavior recommendation methods. These methods typically infer user preferences from various auxiliary behaviors and apply them to recommendations for the target behavior. However, existing methods face challenges in uncovering the interaction patterns for different behaviors from multi-behavior implicit feedback, as users exhibit varying preference strengths for different items across behaviors. To address this issue, this paper introduces a novel approach, Decision-Aware Preference Modeling (DAPM), for multi-behavior recommendation. We first construct a behavior-agnostic graph to learn comprehensive representations that are not affected by behavior factors, complementing the behavior-specific representations. Subsequently, we introduce an innovative contrastive learning paradigm that emphasizes inter-behavior consistency and intra-behavior uniformity to alleviate the “false repulsion” problem in traditional contrastive learning. Furthermore, we propose a multi-behavior hinge loss with boundary constraints to explicitly model users’ decision boundaries across different behaviors, thereby enhancing the model’s ability to accurately capture users’ inconsistent preference intensities. Extensive experiments on three real-world datasets demonstrate the consistent improvements achieved by DAPM over thirteen state-of-the-art baselines. We release our code at <https://github.com/Breeze-del/DAPM>.

1 Introduction

Recommender systems are information filtering technologies designed to provide personalized services based on user preferences [He *et al.*, 2020]. They are widely used in areas like e-commerce [Wu *et al.*, 2021], social media [Qi *et al.*,

2021], and online video platforms [Wei *et al.*, 2020]. Traditional collaborative filtering methods [Cheng *et al.*, 2016; He *et al.*, 2020] typically focus on a single behavior like purchase and often encounter challenges with data sparsity, resulting in significant performance degradation.

In real-world scenarios, users can interact with items in multiple ways, including viewing, adding to the cart and purchasing. Different types of behaviors may characterize user preference from various intention dimensions and complement each other for better user preference learning [Tanjim *et al.*, 2020]. To overcome the challenges posed by data sparsity, researchers have introduced Multi-Behavior Recommendation (MBR) methods [Wei *et al.*, 2022; Chen *et al.*, 2020a; Meng *et al.*, 2023b; Li *et al.*, 2024], which draw on data from multiple user behaviors to provide rich insights into user preferences. MBR methods primarily focus on the target behavior (i.e., purchase) while treating other behaviors as auxiliary behaviors, thereby enhancing the robustness and accuracy of the recommendation system by utilizing a broader spectrum of user data.

Early studies [Singh and Gordon, 2008; Zhao *et al.*, 2015a] employed matrix factorization using shared embeddings for MBR. With the rise of deep learning, neural network-based methods [Gao *et al.*, 2019; Guo *et al.*, 2019; Xia *et al.*, 2020] have gained traction for their ability to model complex relationships between users and items, capturing nuanced user interests and item characteristics. Among these, graph neural networks [Chen *et al.*, 2020b; Schlichtkrull *et al.*, 2018; Meng *et al.*, 2023b; Zhu *et al.*, 2024] have been widely applied in MBR due to their effectiveness in utilizing high-order connectivity between users and items. For example, MBGCN [Jin *et al.*, 2020] and S-MBRec [Gu *et al.*, 2022] implement behavior-aware embedding propagation layers to learn behavior diversity by aggregating multi-behavior interactions from high-order neighbors. However, these methods do not consider using dependencies between behaviors to assist model learning. To bridge this gap, recent works have introduced a behavior hierarchical order to support recommendations. CRGCN [Yan *et al.*, 2023] and MB-CGCN [Cheng *et al.*, 2023] incorporate cascade dependencies between behaviors into graph convolutional networks, improving the representation learning of users and items. BCIPM [Yan *et al.*, 2024] further develops a behavior-contextualized item preference network to enhance item-aware preferences. Despite the tech-

* Corresponding author.




Figure 1: Examples of user interaction patterns influenced by inconsistent preference strengths.

nical difference, these methods share the common goal of integrating insights from auxiliary behaviors to guide recommendations for target behaviors.

Despite the effectiveness of existing methods, they face a significant challenge: accurately inferring users’ interaction patterns from complex multi-behavior signals. This challenge is exemplified in real-world scenarios. As depicted in Figure 1, Susan views four items—diamond, socks, phone and watch, and adds socks and phone to her cart but ultimately only purchases the socks. This diversity in her view list indicates an exploratory phase, where curiosity and information-seeking drive engagement with various items, including luxury and electronic items. During the cart list, practical considerations such as usability and price become more prominent, narrowing her focus to the phone and socks. By the time she makes a purchase, her decision is guided by immediate needs and budget constraints, leading her to choose the socks. Each behavior—view for information, cart for tentative choices, and purchase for final decisions—serves distinct objectives that influence the user’s focus and selection criteria, leading to varying preference intensities for different items across different behavioral contexts. Therefore, understanding users’ comprehensive preferences and decision-making boundaries is essential for capturing their interaction patterns, yet most existing MBR methods overlook these nuanced factors. Our approach aims to address this gap by emphasizing comprehensive preferences learning and explicitly modeling users’ personalized decision boundaries, thereby enhancing recommendation effectiveness.

In addressing the challenges identified within MBR, we propose a novel approach named Decision-Aware Preference Modeling (DAPM). In DAPM, we construct a behavior-agnostic user-item interaction graph that leverages multi-behavior interaction data to capture a broader range of collaborative information, complementing preference learning on type-specific behavior subgraphs and leading to a more comprehensive understanding of user preferences. In addition, to avoid the “false repulsion” problem in traditional contrastive learning—where users with highly similar interactions are overlooked and incorrectly treated as negative samples—we eliminate the global contrastive loss for negative samples across all users. Instead, we focus on measuring the cross-behavior embedding consistency and embedding uniformity within behavior for individual users. This enables our developed DAPM to effectively extract additional super-

vision signals from different types of user behaviors, thereby enhancing the model optimization process with sparse supervision labels. Inspired by Bayesian personalized ranking loss, we introduce a multi-behavior hinge loss that enlarges the gap between positive and negative interactions by modeling the user’s acceptance and rejection boundaries. Specifically, DAPM increases the likelihood of items that have interacted in a specific behavior, pushing it above the acceptance boundary, while decreasing the likelihood of the remaining uninteracted items, bringing them closer to the rejection boundary, thereby effectively exploring users’ decision boundaries. Furthermore, we introduce a boundary constraint to prevent the learnable acceptance boundary from continually approaching the rejection boundary, thereby enhancing the model’s ability to capture users’ personalized decision boundaries.

In summary, our major contributions are summarized as follows. (1) We develop a new multi-behavior learning paradigm DAPM for recommendation by emphasizing the importance of comprehensive preferences and personalized decision boundaries. (2) We propose a multi-behavior hinge loss that explicitly models users’ varying preference intensities by introducing acceptance and rejection boundaries, enabling the model to make accurate recommendations from the user’s perspective. (3) Comprehensive experiments on three real-world datasets demonstrate that our DAPM outperforms the state-of-the-art approaches in multi-behavior scenarios. Further experimental results verify the rationality and effectiveness of the designed sub-modules.

2 Related Works

Recently, MBR methods [Meng *et al.*, 2023b; Li *et al.*, 2024] have attracted considerable attention for their effectiveness in addressing data sparsity challenges. With advancements in technology, MBR methods can be broadly categorized into four main types: traditional machine learning methods, deep neural network (DNN) methods, graph convolutional network (GCN) methods, and self-supervised learning methods.

Traditional machine learning-based methods to MBR often tackle multi-behavior data through multiple matrix factorization techniques [Singh and Gordon, 2008; Tang *et al.*, 2016; Zhao *et al.*, 2015b] or innovative sampling strategies [Loni *et al.*, 2016; Ding *et al.*, 2018]. The former extends traditional matrix factorization by using multiple matrices with shared embeddings, such as CMF [Zhao *et al.*, 2015b]. The latter leverages various user behaviors as auxiliary data, designing sampling strategies that enhance the training process. For example, MF-BPR [Loni *et al.*, 2016] and VALS [Ding *et al.*, 2018], introduce and refine negative sampling strategies to improve recommendation performance.

As deep learning technologies advance, researchers are increasingly investigating MBR methods that leverage DNNs and GCNs. DNN-based methods [Liang *et al.*, 2023; Guo *et al.*, 2019; Xia *et al.*, 2020] typically design complex mechanisms to learn embeddings from different behaviors, integrating them into predictions for the target behavior. For instance, DIPN [Guo *et al.*, 2019] and MATN [Xia *et al.*, 2020] utilize various attention mechanisms to understand relationships between behaviors for effective embedding learning and ag-

gregation. In contrast, NMTR [Gao *et al.*, 2019] employs a multi-task learning strategy, treating all user behaviors as prediction objectives and using the prediction score of the previous behavior to inform the next.

GCN-based methods [Li *et al.*, 2023; Schlichtkrull *et al.*, 2018; Zhu *et al.*, 2024] focus on learning embeddings by constructing a unified user-item graph and performing graph convolution operations. For example, MBGCN [Jin *et al.*, 2020] emphasizes behavior semantics, leveraging an item-item propagation layer alongside user-item propagation to enhance score predictions. CRGCN [Yan *et al.*, 2023] and MB-CGCN [Cheng *et al.*, 2023] further refine these methods by incorporating hierarchical relationships between behaviors, utilizing cascading graph convolutional networks to effectively capture user preferences and achieve notable performance improvements. The recently proposed PKEF [Meng *et al.*, 2023a] builds on MB-CGCN by incorporating a parallel knowledge fusion module and a projected disentangled multi-expert network to address the issue of data distribution imbalance in multiple behaviors. In addition, BCIPM [Yan *et al.*, 2024] develops a behavior-contextualized item preference network to learn item-aware preferences, enhancing the high-order neighbor preferences of users captured by the GCNs.

Inspired by the success of self-supervised learning (SSL), some studies [Wei *et al.*, 2022; Liang *et al.*, 2023] have attempted to leverage contrastive learning to enhance representation learning. The fundamental principle of contrastive learning is to learn high-quality, discriminative representations by maximizing the similarity among positive samples while minimizing it among negative samples. For example, CML [Wei *et al.*, 2022] employs contrastive meta-learning to capture transferable user-item knowledge across various behaviors while maintaining users’ personalized multi-behavior patterns. MBSSL [Xu *et al.*, 2023] performs node self-identification at both inter-behavior and intra-behavior levels in contrastive learning to obtain refined node representations.

Despite the technical diversity among these MBR methods, they commonly depend on collaborative information from auxiliary behaviors to improve the understanding of user preferences. Traditional machine learning-based methods directly extract information from user-item interactions in auxiliary behaviors. In contrast, other methods model overall auxiliary behaviors. However, these methods often overlook the inconsistency in users’ preference intensities for different items across behaviors. Our proposed method stands in sharp contrast by focusing on finely modeling users’ diverse item preference strengths across behaviors through the introduction of acceptance and rejection boundaries.

3 Problem Formulation

We define \mathcal{U} ($u \in \mathcal{U}$) and \mathcal{V} ($v \in \mathcal{V}$) as the sets of users and items, respectively. In our multi-behavior recommendation scenario, we define the \mathcal{B} ($b \in \mathcal{B}$) as the set of behaviors, where the $|\mathcal{B}|$ -th behavior is the target behavior and other behaviors are auxiliary behaviors. The \mathbf{e}_b represents the b -th behavior type embedding. The multi-behavior interaction matrices can be represented as a set, i.e., $\mathcal{M} = \{\mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_{|\mathcal{B}|}\}$, where $\mathbf{M}_b = [m_{(b)uv}]_{|\mathcal{U}| \times |\mathcal{V}|} \in \{0, 1\}$

indicates whether user u interacted with item v under behavior b . The task we aim to address is formally described as follows: **Input**: observed user-item interactions with multiplex $|\mathcal{B}|$ types of behaviors \mathcal{M} among users \mathcal{U} and items \mathcal{V} . **Output**: a predictive function which estimates the likelihood of user u will interact with item v under the target type ($|\mathcal{B}|$) of behaviors.

4 Methodology

The motivation behind this work stems from the observation that users exhibit varying levels of preference for different items across different behaviors. Based on this observation, we propose a decision-aware modeling module that explicitly explores personalized decision boundaries, focusing on learning the varying preference intensities for each item within specific behaviors. To complement the behavior-specific preferences, we first construct a behavior-agnostic graph to capture more comprehensive user preferences. This graph is created from interaction data encompassing all behaviors, allowing us to utilize a broader and more nuanced spectrum of user-item interaction data. We then introduce a multi-behavior contrastive learning paradigm and a gated fusion mechanism to refine and enhance the behavior-specific preference using these comprehensive preferences. The architecture of DAPM is depicted in Figure 2, and the key components are explained in detail in the following subsections.

4.1 Behavior-specific Embedding Learning

To incorporate high-order connectivity into multiplex relation learning across users and items, we first establish a graph-based message passing framework that considers type-specific behavior context. Following the findings in the state-of-the-art model LightGCN [He *et al.*, 2020], our behavior-aware message passing scheme can be represented:

$$\mathbf{E}^{b,(l+1)} = (\mathbf{D}^{-\frac{1}{2}} \mathbf{A}^b \mathbf{D}^{-\frac{1}{2}}) \mathbf{E}^{b,(l)}, \quad (1)$$

where $\mathbf{E}^{b,(l)}$ represents the l -th layer node embedding under b -th behavior, initialized as $\mathbf{E}^{b,(0)} = \mathbf{E}^{(0)} + \mathbf{E}_{beh}$, $\mathbf{E}^{(0)}$ is the initialization of the user and item embeddings and \mathbf{E}_{beh} is composed of stacked behavior type embeddings \mathbf{e}_b . \mathbf{A}^b is adjacency matrix under b -th behavior, whose upper right block matrix is \mathbf{M}_b and lower left block matrix is \mathbf{M}_b^T . \mathbf{D} denotes the diagonal identity matrix of \mathbf{A}^b . The embeddings obtained from different layers emphasize the information passed from different hops. Thus, we further combine them to get the final embeddings: $\mathbf{E}^b = \frac{1}{L+1} \sum_{l=0}^L \mathbf{E}^{b,(l)}$. Finally, we obtain the embeddings under b -th behavior of a specific user u and item v from the \mathbf{E}^b as \mathbf{e}_u^b and \mathbf{e}_v^b .

4.2 Behavior-agnostic Embedding Learning

In real-world recommendation scenarios, users’ attention to items varies across different behaviors. For instance, as shown in Figure 1, compared to the more focused cart and purchase behaviors, Susan explores a wider range of items during viewing—ranging from luxury goods to electronics and daily necessities. This indicates that behavior-specific interactions capture only partial aspects of user preferences.

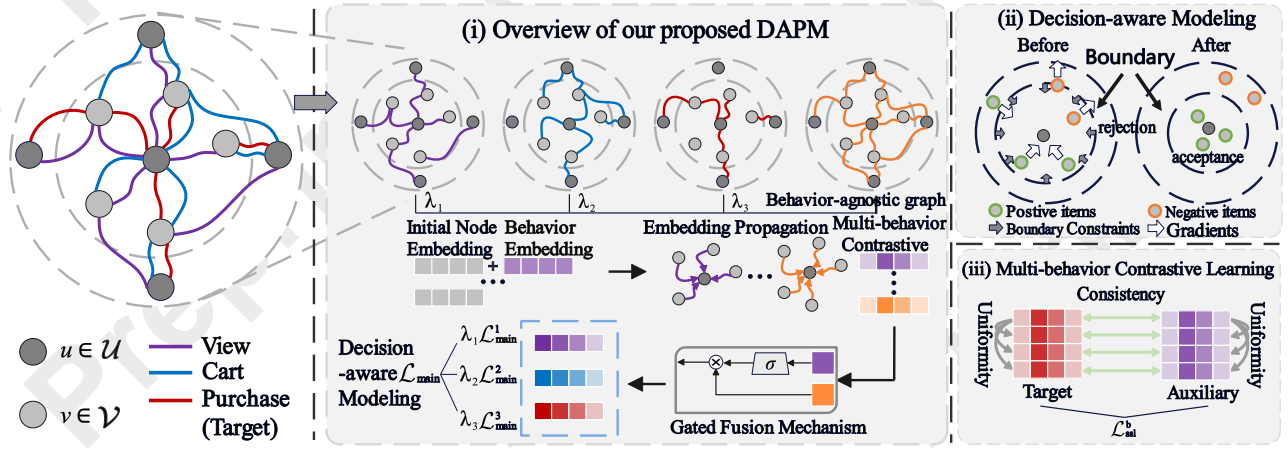


Figure 2: Illustration of the proposed DAPM framework.

Furthermore, the inconsistent sparsity across behaviors may compromise the quality of behavior-specific embeddings.

To address these limitations, we first construct a behavior-agnostic interaction graph by aggregating multi-behavior data, enabling the model to capture broader collaborative signals and enrich the learned representations. However, treating all behaviors equally may obscure users’ personalized behavioral tendencies, as repeated interactions with the same item across behaviors can imply stronger interest, and different behaviors inherently reflect varying preference intensities. To better model this heterogeneity, we introduce behavior-specific weighting coefficients λ to reweight the interaction matrices, thereby enhancing the expressiveness of user preferences across diverse behavioral signals. The behavior-agnostic interaction matrix $\mathbf{M}_{|\mathcal{B}|+1}$ is formalized as follows:

$$\mathbf{M}_{|\mathcal{B}|+1} = \lambda_1 \mathbf{M}_1 + \dots + \lambda_{|\mathcal{B}|} \mathbf{M}_{|\mathcal{B}|}. \quad (2)$$

Then, we apply the same message-passing mechanism as in Section 4.1 on the interaction matrix $\mathbf{M}_{|\mathcal{B}|+1}$ to learn embeddings $\mathbf{e}_u^{|\mathcal{B}|+1}, \mathbf{e}_v^{|\mathcal{B}|+1}$, treating it as an additional auxiliary behavior $|\mathcal{B}| + 1$ to provide more comprehensive preferences. Note that, in order to maintain behavior agnosticism during the message-passing procedure, we use the initial embeddings $\mathbf{E}^{(0)}$ directly, rather than $\mathbf{E}^{b,(0)}$.

4.3 Cross-behavior Aggregation

In the prediction layer of the DAPM framework, we propose using the more comprehensive embeddings learned from the behavior-agnostic graph to supplement and enhance the behavior-specific embeddings. To adaptively achieve this cross-behavior aggregation, we develop a gating network that learns the explicit importance of behavior-specific embeddings. Formally, the aggregation process is as follows:

$$\begin{aligned} \mathbf{g}_u^b &= \text{sigmoid}(\mathbf{W}_u^b \mathbf{e}_u^b + \mathbf{b}_u^b), \\ \mathbf{z}_u^b &= \mathbf{g}_u^b \odot \mathbf{e}_u^b + (1 - \mathbf{g}_u^b) \odot \mathbf{e}_u^{|\mathcal{B}|+1}, \end{aligned} \quad (3)$$

where \mathbf{g}_u^b denotes the importance of the behavior-specific embeddings \mathbf{e}_u^b during the aggregation procedure, $\mathbf{W}_u^b \in \mathbb{R}^{d \times d}$,

$\mathbf{b}_u^b \in \mathbb{R}^d$ are trainable parameters. \odot represents the element-wise product of vectors. Similar aggregation is applied for the item side. To sum up, by integrating user interactions across all behaviors, we can learn more comprehensive embeddings and then apply a gated fusion mechanism to adaptively enhance and refine the behavior-specific embeddings.

4.4 Decision-aware Modeling

Existing multi-behavior methods have made progress in modeling user preferences, but they often overlook the complexity of interaction patterns—specifically, the varying preference intensities users exhibit across different behavioral contexts. Inferring fine-grained preferences from implicit feedback remains challenging, as observed interactions may not accurately reflect true preference intensity. For example, $m_{(b)uv} = 1$ denotes that user u interacted with item v under the b -th behavior, but this does not necessarily indicate that u favors v . Moreover, for the same item, users may apply different decision criteria across behaviors such as viewing or purchasing, resulting in different levels of preference.

To model these nuanced signals, we introduce behavior-specific decision boundaries that represent each user’s interaction criteria under different behavioral contexts. Unlike prior works [Hsieh *et al.*, 2017; Luo *et al.*, 2023] that assume shared thresholds across behaviors, we argue that each behavior reflects a distinct aspect of preference and thus requires separate decision boundaries. Our DAPM enhances expressiveness by learning a personalized acceptance boundary t_{uv}^b for each user-item-behavior combination, providing an explicit interpretation of what it means to “accept” or “reject” an item. This enables us to move beyond binary interaction modeling and capture varying degrees of preference among observed interactions. Specifically, we estimate behavior-specific preference scores by projecting user and item behavior embeddings through a behavior-specific prediction layer:

$$\hat{y}_{uv}^b = \text{ReLu}(\alpha_b^T (\mathbf{z}_u^b \odot \mathbf{z}_v^b)), \quad (4)$$

where $\alpha_b \in \mathbb{R}^d$ denotes the predict layer under b -th behavior. To reduce computational cost, the acceptance boundary t_{uv}^b is parameterized via matrix factorization, and we fix the rejection boundary at 0. We then design a decision-aware hinge

loss with asymmetric constraints: (1) preference scores for positive interactions should exceed the acceptance boundary t_{uv}^b , indicating strong interest; and (2) those for negative interactions should remain close to zero, reflecting weak or no preference. As illustrated in Figure 2(ii), this design allows behavior- and instance-specific supervision. The main loss for b -th behavior can be formalized as follows:

$$\mathcal{L}_{main}^b = \sum_{u \in \mathcal{U}} \left(\sum_{v \in \mathcal{V}_u^{(b)+}} f^2(t_{uv}^b - \hat{y}_{uv}^b) + c \sum_{v' \in \mathcal{V}_u^{(b)-}} f^2(\hat{y}_{uv'}^b) \right), \quad (5)$$

where $\mathcal{V}_u^{(b)+}$, $\mathcal{V}_u^{(b)-}$ denotes the corresponding observed and unobserved interacted items of user u under b -th behavior; c denotes the weight of negative entry, $f(\cdot)$ denotes $\max(\cdot, 0)$. At inference time, the predicted score $\hat{y}_{uv}^{|\mathcal{B}|}$ for target behavior $|\mathcal{B}|$ is adjusted by its corresponding acceptance boundary $t_{uv}^{|\mathcal{B}|}$, yielding a final score $\hat{y}_{uv}^{|\mathcal{B}|}/t_{uv}^{|\mathcal{B}|}$ that reflects the user's preference relative to their personal behavioral threshold and guides recommendation more effectively.

However, this design may introduce an optimization issue where the acceptance boundary t_{uv}^b gradually decreases and approaches the rejection boundary. Since these boundaries play a crucial role in modeling the intensity and subtlety of user preferences, a vanishing gap between them undermines the model's capacity to distinguish fine-grained preference signals. To mitigate this, we introduce a boundary constraint loss to prevent collapse of the acceptance boundary: $\mathcal{L}_{cons}^b = \sum_{u \in \mathcal{U}} \sum_{v \in \mathcal{V}} -\log(t_{uv}^b)$. This regularization term penalizes excessively small acceptance boundaries, effectively maintaining a meaningful separation from the rejection threshold and preserving the expressiveness of user preference modeling.

4.5 Multi-behavior Contrastive Learning

Compared to other types of user-item interactions, interactions observed under the target behavior are often sparse. The sparse supervision signals of the target behavior may lead to severe bias of learned representations compared with those of auxiliary behaviors. In view of the fact that supervision signals in auxiliary behaviors are much richer than that in the target behavior, we perform selective contrastive learning between auxiliary behaviors and the target behavior to enable knowledge transfer, thereby alleviating the sparsity of target behavior data. Specifically, we treat each behavior type as a distinct view, considering different behaviors of the same user as positive pairs and those of different users as negative pairs. Given the target behavior embedding $\mathbf{e}_u^{|\mathcal{B}|}$, we construct positive pairs $\{\mathbf{e}_u^{|\mathcal{B}|}, \mathbf{e}_u^b | u \in \mathcal{U}\}$ and negative pairs $\{\mathbf{e}_u^{|\mathcal{B}|}, \mathbf{e}_{u'}^b | u, u' \in \mathcal{U}, u \neq u'\}$. This auxiliary supervision helps the model identify user-specific patterns across behaviors and capture the latent relations between target and auxiliary signals. Following [Wei *et al.*, 2022], we adopt the InfoNCE loss [Oord *et al.*, 2018] within this cross-behavior contrastive learning framework, computed as follows:

$$\mathcal{L}_{cl}^b = \sum_{u \in \mathcal{U}} -\log \frac{\exp(\phi(\mathbf{e}_u^{|\mathcal{B}|}, \mathbf{e}_u^b)/\tau)}{\sum_{u' \in \mathcal{U}} \exp(\phi(\mathbf{e}_u^{|\mathcal{B}|}, \mathbf{e}_{u'}^b)/\tau)}, \quad (6)$$

Dataset	#User	#Item	#Interaction	#Target Interaction	#Interactive Behavior Type
Beibei	21,716	7,977	3.3×10^6	282,860	{View, Cart, purchase}
Taobao	48,749	39,493	2.0×10^6	146,247	{View, Cart, purchase}
Tmall	41,738	11,953	2.3×10^6	255,586	{View, Collect, Cart, purchase}

Table 1: Statistics of evaluation datasets.

where $\phi(\cdot)$ denotes the inner product between two embeddings, τ represents the temperature hyperparameter for the softmax function. However, forcing all user embeddings to be repelled may lead to the “false repulsion” phenomenon, where users with highly similar interactions are overlooked and incorrectly treated as negative samples, thereby interfering with the model's ability to accurately capture user preferences. To address this issue, we remove the global contrastive of negative samples across all users, instead focusing on measuring the similarity of embeddings for individual users across behaviors while enhancing uniformity to have more uniformly distributed user representations within their hyperspheres. Therefore, we conduct efficient self-supervised contrastive learning of user representation from the perspective of inter-behavior consistency and intra-behavior uniformity, formalized as follows:

$$\mathcal{L}_{ssl}^b = \sum_{u \in \mathcal{U}} \left(1 - \mathbf{e}_u^{|\mathcal{B}|} (\mathbf{e}_u^b)^T \right)^2 + \log \sum_{u \neq u' \in \mathcal{U}} \exp \left(-2 \|\mathbf{e}_u^b - \mathbf{e}_{u'}^b\|^2 \right). \quad (7)$$

These improvements allow the model to avoid forcibly increasing the distance between all negative user embeddings, thereby alleviating the “false repulsion” caused by behavioral similarities. Moreover, we constrain the Euclidean distance between different user embeddings to prevent excessive expansion and maintain the stability of the embedding space. These modifications simplify the computation, improve efficiency and stability, and enable the model to capture more label-irrelevant data structures and patterns, leading to more generalized consistency and uniformity, while also enhancing recommendation performance. Finally, we combine the above objective functions for joint optimization:

$$\mathcal{L}_{all} = \sum_{i=1}^{|\mathcal{B}|} \lambda_i (\mathcal{L}_{main}^{b_i} + \mathcal{L}_{cons}^{b_i}) + \sum_{i=1}^{|\mathcal{B}|+1} \beta \mathcal{L}_{ssl}^{b_i}, \quad (8)$$

where the normalized λ_i adjusts the effect of the i -th behavioral tasks, β controls the impact of contrastive learning.

5 Experiments

5.1 Experimental Setting

Dataset Description. To evaluate the effectiveness of the proposed DAPM, we conduct extensive experiments on three public multi-behavior datasets, including Beibei, Taobao and Tmall. For the three datasets, to eliminate duplicate data, we follow the previous works [Cheng *et al.*, 2023; Yan *et al.*, 2023; Meng *et al.*, 2023a] to retain only the earliest occurrence of each interaction. The specific statistical information of the three datasets is presented in Table 1.

Model	Beibei				Taobao				Tmall			
	H@10	N@10	H@20	N@20	H@10	N@10	H@20	N@20	H@10	N@10	H@20	N@20
MF-BPR	0.0191	0.0049	0.0237	0.0053	0.0178	0.0101	0.0221	0.0113	0.0230	0.0124	0.0278	0.0135
LightGCN	0.0351	0.0190	0.0473	0.0211	0.0254	0.0138	0.0328	0.0153	0.0393	0.0169	0.0499	0.0182
SGL	0.0422	0.0201	0.0561	0.0223	0.0426	0.0257	0.0537	0.0285	0.0456	0.0185	0.0575	0.0202
SimGCL	0.0466	0.0231	0.0615	0.0245	0.0415	0.0226	0.0531	0.0249	0.0444	0.0195	0.0546	0.0215
SGL_all	0.0599	0.0310	0.0767	0.0332	0.0466	0.0278	0.0592	0.0311	0.0577	0.0266	0.0698	0.0316
SimGCL_all	0.0613	0.0311	0.0736	0.0333	0.0458	0.0274	0.0563	0.0288	0.0592	0.0296	0.0711	0.0321
NMTR	0.0429	0.0198	0.0558	0.0224	0.0409	0.0212	0.0515	0.0242	0.0517	0.0250	0.0651	0.0280
MBGCN	0.0470	0.0259	0.0635	0.0282	0.0434	0.0259	0.0564	0.0282	0.0549	0.0285	0.0686	0.0314
GHCF	0.1722	0.0912	0.1964	0.0903	0.0807	0.0442	0.0993	0.0499	0.0433	0.0175	0.0528	0.0186
CML	0.0507	0.0292	0.0649	0.0310	0.0298	0.0150	0.0378	0.0171	0.0388	0.0127	0.0493	0.0138
MBSSL	0.2114	0.1271	0.2706	0.1436	0.1031	0.0583	0.1299	0.0659	0.0858	0.0436	0.1082	0.0477
CRGCN	0.0459	0.0324	0.0560	0.0356	0.1152	0.0629	0.1463	0.0717	0.0840	0.0442	0.1092	0.0482
MB-CGCN	0.0579	0.0381	0.0724	0.0400	0.0989	0.0470	0.1197	0.0512	0.1073	0.0416	0.1352	0.0474
PKEF	0.1130	0.0582	0.1503	0.0663	0.1097	0.0627	0.1382	0.0683	0.1118	0.0630	0.1487	0.0693
BCIPM	0.1822	0.0985	0.2024	0.1105	0.1292	0.0716	0.1594	0.0843	0.1414	0.0741	0.1833	0.0855
DAPM	0.2603*	0.1534*	0.3140*	0.1670*	0.1666*	0.1106*	0.1832*	0.1149*	0.1558*	0.0820*	0.2072*	0.0951*
#Improve	23.13%	20.69%	16.04%	16.28%	28.95%	54.47%	14.93%	36.30%	10.18%	10.66%	13.04%	11.23%

Table 2: The performance comparison on three datasets. Note that baselines with the “all” suffix use data from all the behaviors to build the single-behavior model. The best results are illustrated in bold and the number underlined is the runner-up. Superscript * indicates the significant improvement between our DAPM and the best performing baseline with p -value ≤ 0.05 .

Evaluation Protocols. In all our experiments, we assess the performance of our proposed DAPM model and baseline methods based on the top- k recommended items, using two evaluation metrics: Hit Ratio ($H@k$) and Normalized Discounted Cumulative Gain ($N@k$). For more experimental analysis, please refer to the supplementary materials.

Baseline Methods. To demonstrate the effectiveness of DAPM, we compare it with several state-of-the-art methods, which can be divided into three categories: (1) **Single-behavior methods:** MF-BPR [Rendle *et al.*, 2012] and LightGCN [He *et al.*, 2020], (2) **Self-supervised learning methods:** SGL [Wu *et al.*, 2021] and SimGCL [Yu *et al.*, 2022], (3) **Multi-behavior methods:** NMTR [Gao *et al.*, 2019], MBGCN [Jin *et al.*, 2020], GHCF [Chen *et al.*, 2021], CML [Wei *et al.*, 2022], MBSSL [Xu *et al.*, 2023], CRGCN [Yan *et al.*, 2023], MB-CGCN [Cheng *et al.*, 2023], PKEF [Meng *et al.*, 2023a] and BCIPM [Yan *et al.*, 2024].

Parameter Settings. For all methods, we uniformly set the batch size to 1024 and the embedding size to 64 during the training phase. The parameters are optimized by Adam, while the learning rate is set to $1e^{-3}$. We adjust the behavior coefficients for each behavior in $[0, 1/6, 2/6, 3/6, 4/6, 5/6, 1]$. To determine the optimal values for the hyperparameters, including c and β , we perform a grid search on the set $[1e^{-2}, 1e^{-1}, 3e^{-1}, 5e^{-1}, 7e^{-1}, 1, 10, 100]$. To ensure fairness, we also set parameters for the baselines according to the descriptions in their papers and perform a grid search to find the optimal values.

5.2 Performance Evaluation

Table 2 showcases the comparative experimental results of our proposed model alongside thirteen benchmark baseline models, evaluated across three distinct datasets.

LightGCN consistently outperforms MF-BPR across all three datasets, leveraging the strong representation learning capabilities of Graph Convolutional Networks. Moreover, SSL-based methods surpass single-behavior methods,

demonstrating that SSL enhances representation learning and improves the generalization ability of recommender models.

Most multi-behavior recommendation methods outperform single-behavior methods, highlighting the benefits of incorporating multi-behavioral information into user preference modeling. Among various multi-behavior recommendation models, BCIPM is the best baseline in most cases. This suggests that capturing users’ comprehensive preferences better guides model optimization.

Our DAPM outperforms all baselines across all datasets. Specifically, DAPM achieves average performance improvements of 19.16%, 33.66%, and 11.28% over the best baseline on the Beibei, Taobao, and Tmall datasets, respectively. This success is attributed to DAPM’s ability to effectively capture the varying preference intensities of users across different behavioral contexts at a fine-grained level, while the behavior-agnostic graph aids in understanding users’ overall preferences. Additionally, we introduce a novel contrastive learning paradigm that improves embedding quality through inter-behavior consistency and intra-behavior uniformity. The significant boost in recommendation accuracy underscores the effectiveness of our approach.

5.3 Ablation Studies

To evaluate the effectiveness of the individual design components in the DAPM framework, we consider five model variants: (1) removing the behavior-agnostic graph and learning embeddings only from type-specific behavior graph (w/o. BAG.), (2) removing the multi-behavior contrastive learning module (w/o. MBC.), (3) adopting the traditional InfoNCE-based contrastive loss (w/o. INC.), (4) replacing our hinge loss with a BPR loss (w/o. MHL.), and (5) removing the boundary constraint component (w/o. BDL.).

The experimental results are presented in Table 3. From the results, we observe that the w/o. BAG. variant shows a performance drop across all datasets, highlighting the importance of learning more comprehensive preferences to complement

Model Variants	Beibei		Taobao		Tmall	
	H@10	N@10	H@10	N@10	H@10	N@10
DAPM	0.2603	0.1534	0.1666	0.1106	0.1558	0.0820
w/o. BAG.	0.2185	0.1308	0.1132	0.0672	0.1311	0.0705
w/o. MBC.	0.2215	0.1331	0.1459	0.0969	0.1371	0.0720
w/o. INC.	0.2385	0.1428	0.1518	0.1018	0.1415	0.0731
w/o. MHL.	0.2082	0.1242	0.1361	0.0913	0.1268	0.0687
w/o. BDL.	0.2504	0.1469	0.1511	0.0993	0.1504	0.0774

Table 3: Performances of different DAPM variants.

Model	Beibei		Taobao		Tmall	
	E	T	E	T	E	T
DAPM	16 (± 1)s	86m	20 (± 1)s	111m	21 (± 1)s	126m
NMTR	165 (± 5)s	550m	180 (± 5)s	600m	196 (± 5)s	670m
GHCF	13 (± 1)s	45m	34 (± 1)s	115m	44 (± 1)s	140m
BCIPM	402 (± 10)s	890m	903 (± 10)s	2211m	1597 (± 10)s	1822m

Table 4: Computational Time Cost Investigation (Second/Minute [s/m]). Here “E” and “T” represent the training time for each epoch and the total training time, respectively.

behavior-specific preferences. A similar trend is observed with the w/o. MCL. variant, indicating the effectiveness of inter-behavior consistency and intra-behavior uniformity in improving embedding quality. Furthermore, the performance degradation of the w/o. INC. variant validates the “false repulsion” mentioned before, where potential incorrect negative samples interfere with the learning of representations. The performance of the w/o. MHL. variant significantly drops, demonstrating that our multi-behavior hinge loss captures varying inherent preference intensities, which notably improves performance. Finally, DAPM outperforms the w/o. BDL. variant, confirming that continuously pushing the acceptance boundary toward the rejection boundary harms the model’s ability to accurately model preference strength.

5.4 Efficiency Analysis of DAPM

In this section, we evaluate the computational time of our proposed DAPM and compare it with three state-of-the-art multi-behavior methods: NMTR, GHCF, and BCIPM. All methods are trained on a single NVIDIA GeForce GTX 3090 GPU with the same hidden state dimensionality setting to ensure a fair comparison of efficiency. Table 4 presents the runtime for each method during each epoch and the total training time. DAPM outperforms NMTR by effectively integrating users’ diverse preferences into decision boundary modeling, thereby enhancing preference learning efficiency. GHCF improves performance by integrating edge relationships into the graph message-passing mechanism to learn better representations of users and items, with the cost of low efficiency. BCIPM achieves near-optimal performance across the three datasets, benefiting from its accurate learning of item preferences and high-order neighbor preferences. However, this two-pronged design and the additional embedding pre-training module considerably increase training time and reduce model efficiency. In contrast, we improve efficiency by removing the computation of negative samples in multi-behavior contrastive learning and simplifying the decision boundary modeling. In a nutshell, our DAPM shows a competitive model scalability because of the comparable compu-

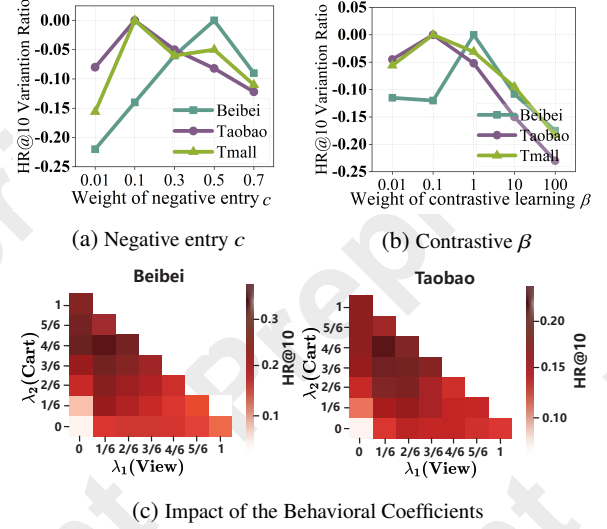


Figure 3: Hyperparameter study of DAPM.

tational complexity and great performance.

5.5 Parameter Analysis

We conduct extensive experiments to examine the impact of several key hyperparameters, including the weight for negative entry c , the weight for contrastive learning β , and behavior coefficients λ . Figure 3 illustrates the percentage decrease in performance relative to the best performance achieved. From Figures 3(a)(b), we observe a similar performance trend across different settings, i.e., performance initially improves significantly as the weight increases, but then quickly declines. For the behavior coefficient, both the Beibei and Taobao datasets consist of three behavior types—view, cart, and purchase—which correspond to the three loss coefficients λ_1 , λ_2 , and λ_3 . The HR@10 results are shown in Figure 3(c), in which darker blocks mean better performance. It can be found that a relatively large coefficient for the cart behavior performs best on Beibei and Taobao. The potential reason is that purchase interaction is too sparse to offer much information and view interaction is relatively far from the target behavior. For both datasets, DAPM achieves optimal performance when the coefficients are set to [1/6, 4/6, 1/6], which not only effectively utilizes all three behavior types but also places more emphasis on cart behavior.

6 Conclusion

In this paper, we revisit the use of multi-behavior data in recommendation and propose a novel method, DAPM. Our approach constructs a behavior-agnostic graph for comprehensive representation learning, incorporates multi-behavior contrastive learning and a gating mechanism to adaptively refine behavior-specific embeddings, and explicitly models users’ varying preference strengths via personalized decision boundaries. Extensive experiments on three real-world datasets demonstrate the effectiveness of DAPM.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (U22B2060, 62276279), Key-Area Research and Development Program of Guangdong Province (2024B0101050005), Research Foundation of Science and Technology Plan Project of Guangzhou City(2023B01J0001, 2024B01W0004), Guangdong Basic and Applied Basic Research Foundation (2024A151501024, 2024B1515020032), Technology Innovation Center for Collaborative Applications of Natural Resources Data in GBA, MNR(No.2025NRDZ02), Zhuhai City Industry-University-Research Project (2320004002797), 2024 Guangdong Province Science and Technology Innovation Strategic Special Fund (College Student Science and Technology Innovation Cultivation, pdjh2024b015).

References

- [Chen *et al.*, 2020a] Chong Chen, Min Zhang, Yongfeng Zhang, Weizhi Ma, Yiqun Liu, and Shaoping Ma. Efficient heterogeneous collaborative filtering without negative sampling for recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 19–26, 2020.
- [Chen *et al.*, 2020b] Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. Revisiting graph based collaborative filtering: A linear residual graph convolutional network approach. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 27–34, 2020.
- [Chen *et al.*, 2021] Chong Chen, Weizhi Ma, Min Zhang, Zhaowei Wang, Xiuqiang He, Chenyang Wang, Yiqun Liu, and Shaoping Ma. Graph heterogeneous multi-relational recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 3958–3966, 2021.
- [Cheng *et al.*, 2016] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, et al. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*, pages 7–10, 2016.
- [Cheng *et al.*, 2023] Zhiyong Cheng, Sai Han, Fan Liu, Lei Zhu, Zan Gao, and Yuxin Peng. Multi-behavior recommendation with cascading graph convolution networks. In *Proceedings of the ACM Web Conference 2023*, pages 1181–1189, 2023.
- [Ding *et al.*, 2018] Jingtao Ding, Guanghui Yu, Xiangnan He, Yuhuan Quan, Yong Li, Tat-Seng Chua, Depeng Jin, and Jiajie Yu. Improving implicit recommender systems with view data. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 3343–3349, 2018.
- [Gao *et al.*, 2019] Chen Gao, Xiangnan He, Dahua Gan, Xiangning Chen, Fuli Feng, Yong Li, Tat-Seng Chua, and Depeng Jin. Neural multi-task recommendation from multi-behavior data. In *2019 IEEE 35th international conference on data engineering (ICDE)*, pages 1554–1557, 2019.
- [Gu *et al.*, 2022] Shuyun Gu, Xiao Wang, Chuan Shi, and Ding Xiao. Self-supervised graph neural networks for multi-behavior recommendation. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, pages 2052–2058, 2022.
- [Guo *et al.*, 2019] Long Guo, Lifeng Hua, Rongfei Jia, Bin-qiang Zhao, Xiaobo Wang, and Bin Cui. Buying or browsing?: Predicting real-time purchasing intent using attention-based deep network with multiple behavior. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1984–1992, 2019.
- [He *et al.*, 2020] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 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*, pages 639–648, 2020.
- [Hsieh *et al.*, 2017] Cheng-Kang Hsieh, Longqi Yang, Yin Cui, Tsung-Yi Lin, Serge Belongie, and Deborah Estrin. Collaborative metric learning. In *Proceedings of the 26th international conference on world wide web*, pages 193–201, 2017.
- [Jin *et al.*, 2020] Bowen Jin, Chen Gao, Xiangnan He, Depeng Jin, and Yong Li. Multi-behavior recommendation with graph convolutional networks. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 659–668, 2020.
- [Li *et al.*, 2023] Qingfeng Li, Huifang Ma, Ruoyi Zhang, Wangyu Jin, and Zhixin Li. Intra-and inter-behavior contrastive learning for multi-behavior recommendation. In *Proceedings of the 28th International Conference on Database Systems for Advanced Applications*, pages 147–162, 2023.
- [Li *et al.*, 2024] Qingfeng Li, Huifang Ma, Wangyu Jin, Yungang Ji, and Zhixin Li. Multi-interest network with simple diffusion for multi-behavior sequential recommendation. In *Proceedings of the SIAM International Conference on Data Mining*, pages 734–742, 2024.
- [Liang *et al.*, 2023] Ke Liang, Yue Liu, Sihang Zhou, Wenxuan Tu, Yi Wen, Xihong Yang, Xiangjun Dong, and Xinwang Liu. Knowledge graph contrastive learning based on relation-symmetrical structure. *IEEE Transactions on Knowledge and Data Engineering*, 36(1):226–238, 2023.
- [Loni *et al.*, 2016] Babak Loni, Roberto Pagano, Martha Larson, and Alan Hanjalic. Bayesian personalized ranking with multi-channel user feedback. In *Proceedings of the 10th ACM conference on recommender systems*, pages 361–364, 2016.
- [Luo *et al.*, 2023] Xiao Luo, Daqing Wu, Yiyang Gu, Chong Chen, Luchen Liu, Jinwen Ma, Ming Zhang, Minghua Deng, Jianqiang Huang, and Xian-Sheng Hua. Criterion-based heterogeneous collaborative filtering for multi-behavior implicit recommendation. *ACM Transactions on Knowledge Discovery from Data*, 18(1):1–26, 2023.

- [Meng *et al.*, 2023a] Chang Meng, Chenhao Zhai, Yu Yang, Hengyu Zhang, and Xiu Li. Parallel knowledge enhancement based framework for multi-behavior recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 1797–1806, 2023.
- [Meng *et al.*, 2023b] Chang Meng, Hengyu Zhang, Wei Guo, Huifeng Guo, Haotian Liu, Yingxue Zhang, Hongkun Zheng, Ruiming Tang, Xiu Li, and Rui Zhang. Hierarchical projection enhanced multi-behavior recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4649–4660, 2023.
- [Oord *et al.*, 2018] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [Qi *et al.*, 2021] Tao Qi, Fangzhao Wu, Chuhan Wu, and Yongfeng Huang. Personalized news recommendation with knowledge-aware interactive matching. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 61–70, 2021.
- [Rendle *et al.*, 2012] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th International Conference on Uncertainty in Artificial Intelligence*, page 452–461, 2012.
- [Schlichtkrull *et al.*, 2018] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *The semantic web: 15th international conference*, pages 593–607, 2018.
- [Singh and Gordon, 2008] Ajit P Singh and Geoffrey J Gordon. Relational learning via collective matrix factorization. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 650–658, 2008.
- [Tang *et al.*, 2016] Liang Tang, Bo Long, Bee-Chung Chen, and Deepak Agarwal. An empirical study on recommendation with multiple types of feedback. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 283–292, 2016.
- [Tanjim *et al.*, 2020] Md Mehrab Tanjim, Congzhe Su, Ethan Benjamin, Diane Hu, Liangjie Hong, and Julian McAuley. Attentive sequential models of latent intent for next item recommendation. In *Proceedings of The Web Conference 2020*, pages 2528–2534, 2020.
- [Wei *et al.*, 2020] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, and Tat-Seng Chua. Graph-refined convolutional network for multimedia recommendation with implicit feedback. In *Proceedings of the 28th ACM international conference on multimedia*, pages 3541–3549, 2020.
- [Wei *et al.*, 2022] Wei Wei, Chao Huang, Lianghao Xia, Yong Xu, Jiashu Zhao, and Dawei Yin. Contrastive meta learning with behavior multiplicity for recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, pages 1120–1128, 2022.
- [Wu *et al.*, 2021] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. Self-supervised graph learning for recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pages 726–735, 2021.
- [Xia *et al.*, 2020] Lianghao Xia, Chao Huang, Yong Xu, Peng Dai, Bo Zhang, and Liefeng Bo. Multiplex behavioral relation learning for recommendation via memory augmented transformer network. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 2397–2406, 2020.
- [Xu *et al.*, 2023] Jingcao Xu, Chaokun Wang, Cheng Wu, Yang Song, Kai Zheng, Xiaowei Wang, Changping Wang, Guorui Zhou, and Kun Gai. Multi-behavior self-supervised learning for recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 496–505, 2023.
- [Yan *et al.*, 2023] Mingshi Yan, Zhiyong Cheng, Chen Gao, Jing Sun, Fan Liu, Fuming Sun, and Haojie Li. Cascading residual graph convolutional network for multi-behavior recommendation. *ACM Transactions on Information Systems*, 42(1):1–26, 2023.
- [Yan *et al.*, 2024] Mingshi Yan, Fan Liu, Jing Sun, Fuming Sun, Zhiyong Cheng, and Yahong Han. Behavior-contextualized item preference modeling for multi-behavior recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 946–955, 2024.
- [Yu *et al.*, 2022] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 1294–1303, 2022.
- [Zhao *et al.*, 2015a] Zhe Zhao, Zhiyuan Cheng, Lichan Hong, and Ed H Chi. Improving user topic interest profiles by behavior factorization. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1406–1416, 2015.
- [Zhao *et al.*, 2015b] Zhe Zhao, Zhiyuan Cheng, Lichan Hong, and Ed H Chi. Improving user topic interest profiles by behavior factorization. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1406–1416, 2015.
- [Zhu *et al.*, 2024] Xi Zhu, Fake Lin, Ziwei Zhao, Tong Xu, Xiangyu Zhao, Zikai Yin, Xueying Li, and Enhong Chen. Multi-behavior recommendation with personalized directed acyclic behavior graphs. *ACM Transactions on Information Systems*, 2024.