

Pre-train, Align, and Disentangle: Empowering Sequential Recommendation with Large Language Models

Yuhao Wang
City University of Hong Kong
Hong Kong, China
yhwang25-c@my.cityu.edu.hk

Junwei Pan
Tencent Inc.
Shenzhen, China
jonaspan@tencent.com

Pengyue Jia
City University of Hong Kong
Hong Kong, China
jia.pengyue@my.cityu.edu.hk

Wanyu Wang
City University of Hong Kong
Hong Kong, China
wanyuwang4-c@my.cityu.edu.hk

Maolin Wang
City University of Hong Kong
Hong Kong, China
MorinWang@foxmail.com

Zhixiang Feng
Tencent Inc.
Shenzhen, China
lionelfeng@tencent.com

Xiaotian Li
Tencent Inc.
Shenzhen, China
dylanxtli@tencent.com

Jie Jiang
Tencent Inc.
Shenzhen, China
zeus@tencent.com

Xiangyu Zhao ✉
City University of Hong Kong
Hong Kong, China
xianzhao@cityu.edu.hk

Abstract

Sequential Recommendation (SR) aims to leverage the sequential patterns in users' historical interactions to accurately track their preferences. However, the primary reliance of existing SR methods on collaborative data results in challenges such as the cold-start problem and sub-optimal performance. Concurrently, despite the proven effectiveness of large language models (LLMs), their integration into commercial recommender systems is impeded by issues such as high inference latency, incomplete capture of all distribution statistics, and catastrophic forgetting. To address these issues, we introduce a novel Pre-train, Align, and Disentangle (PAD) framework to enhance SR models with LLMs. In particular, we initially pre-train both the SR and LLM models to obtain collaborative and textual embeddings. Subsequently, we propose a characteristic recommendation-anchored alignment loss using multi-kernel maximum mean discrepancy with Gaussian kernels. Lastly, a triple-experts architecture, comprising aligned and modality-specific experts with disentangled embeddings, is fine-tuned in a frequency-aware manner. Experimental results on three public datasets validate the efficacy of PAD, indicating substantial enhancements and compatibility with various SR backbone models, particularly for cold items. The code and datasets are accessible for reproduction¹.

CCS Concepts

• Information systems → Recommender systems.

¹<https://github.com/Applied-Machine-Learning-Lab/PAD>

✉Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '25, July 13–18, 2025, Padua, Italy

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1592-1/2025/07
<https://doi.org/10.1145/3726302.3730059>

Keywords

Sequential Recommendation, Recommender System, Large Language Model, Reproducing Kernel Hilbert Space

ACM Reference Format:

Yuhao Wang, Junwei Pan, Pengyue Jia, Wanyu Wang, Maolin Wang, Zhixiang Feng, Xiaotian Li, Jie Jiang, and Xiangyu Zhao ✉. 2025. Pre-train, Align, and Disentangle: Empowering Sequential Recommendation with Large Language Models. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3726302.3730059>

1 Introduction

With the explosive growth of the interactions with web applications and platforms [55–57, 68], research on sequential recommender system (SRS) [14, 20, 22, 25, 49, 67] has garnered increasing attention. These models aim to capture the sequential dependencies in user behavior sequences [63, 64], modeling both long-term and short-term preferences [35]. However, most existing SR methods rely exclusively on tabular data or ID-based features as inputs [36, 54, 69], which often leads to challenges such as the cold-start problem [45], resulting in suboptimal performance, particularly for less frequent users, items, and scenarios.

To address the limitations of conventional sequential recommendation (SR) models, recent efforts have drawn inspiration from the success of large language models (LLMs) in understanding semantics and processing natural language [60]. Several approaches have sought to enhance SR by leveraging LLM capabilities [33]. On the one hand, methods like TALLRec [1] and LC-Rec [70] explicitly adapt LLMs for recommendation by instruction tuning. They formulate the sequential recommendation task as text generation, *i.e.*, predicting the title of the next item based on a user's historical interactions. Afterward, LLaRA [30], and CoLLM [66] also leverage LLM as recommender and they propose to concatenate the token embedding with collaborative embedding, which enables LLM to comprehend collaborative information.

On the other hand, models such as CTRL [28] and Flip [52] propose aligning LLM embeddings with collaborative embeddings using contrastive learning, employing an InfoNCE alignment loss based on non-characteristic cosine or linear kernels. Besides, Taobao [45] and AlignRec [34] propose to pre-train multi-modal representations and incorporate them into recommendation. Despite these advances, three practical challenges remain to be addressed:

- **High Inference Latency of LLM-based Recommenders.** LLM-based recommenders [1, 30, 53, 66] rely on the complex architecture of large language models to make predictions, which is impractical for real-world recommendation systems, where the inference of hundreds of items must be completed within hundreds of milliseconds [19, 38, 41].
- **Inability to Capture All Statistics of Data Distribution.** Many existing works employ non-characteristic kernels in their alignment loss [28, 52]. However, research on Reproducing Kernel Hilbert Space [13, 39] has demonstrated that non-characteristic kernels fail to capture all statistical aspects of the data distribution, potentially leading to suboptimal performance. Evidence of the superiority of characteristic kernels in recommendation will be presented in Sec. 4.5.
- **Catastrophic Forgetting in Alignment.** Catastrophic forgetting is a key challenge in multi-modal learning [26, 27]. Our comprehensive empirical evaluation in Sec. 4.6.1 shows that, although alignment helps transfer textual embeddings into the collaborative space, it results in catastrophic forgetting of the collaborative embeddings.

To address the aforementioned challenges, we propose a three-phase framework—Pre-Train, Align, and Disentangle—to empower sequential recommendation with LLMs. This framework consists of (1) pre-training LLM and recommendation models, (2) alignment between textual and collaborative with MK-MMD, and (3) supervised fine-tuning for recommendation. Specifically, during the pre-training phase, we employ a pre-trained LLM, *e.g.*, Llama3 [9], alongside a sequential recommendation model, *e.g.*, SASRec [22]. The textual and collaborative embeddings are then obtained from these two models, respectively.

In the alignment phase, we propose a novel characteristic recommendation anchored alignment loss, which integrates a characteristic alignment loss with a Binary Cross Entropy (BCE) loss based on the recommendation label. The characteristic alignment loss ensures that all statistical properties of the distribution are accounted for during alignment, while the BCE loss mitigates catastrophic forgetting of the collaborative embeddings during the alignment.

In the subsequent disentangle phase, inspired by recent advances in recommendation for handling the distinct characteristics of tasks [48], domains [31], and modalities [45], we incorporate two modality-specific embeddings and experts, in addition to the alignment expert, resulting in a triple-expert architecture. Specifically, we employ an LLM expert and a recommendation expert, which use the LLM embeddings and collaborative embeddings as inputs, respectively. The key contributions of this paper are as follows:

- We identify the limitations of existing SRS works in capturing distribution statistics and catastrophic forgetting.
- We propose a novel three-phase framework—Pre-train, Align, and Disentangle (PAD)—featuring a recommendation-anchored characteristic alignment loss and a triple-expert architecture.

- We conduct extensive experiments on three public datasets, demonstrating the effectiveness of PAD, especially on cold items. Furthermore, we provide a comprehensive study regarding the kernels, catastrophic forgetting and tools to measure alignment.

2 Preliminary

In this section, we first illustrate the problem formulation and introduce maximum mean discrepancy.

2.1 Problem Formulation

First we provide the formulation and notations. In our problem setting, we obtain a semantic domain $\mathcal{D}_{\text{text}} = \{(\{\mathbf{h}_i^s\}, \mathbf{x}_i^s, y_i)\}_{i=1}^n$ and a collaborative domain $\mathcal{D}_{\text{rec}} = \{(\{\mathbf{h}_i^c\}, \mathbf{x}_i^c, y_i)\}_{i=1}^n$ with n samples each. Specifically, $\{\mathbf{h}_i^s\}$, $\{\mathbf{h}_i^c\}$, \mathbf{x}_i^s , \mathbf{x}_i^c , and y_i denotes behavioral item sequence in semantic embedding space, behavioral item sequence in collaborative embedding space, target item in semantic embedding space, target item in collaborative embedding space, and true label. The probability distributions characterized by these two domains are P and Q .

Given a user's historical interaction sequence with length l consisting of item ID, it is first sorted by timestamps in an ascending order and mapped into collaborative embedding $\{\mathbf{h}_i^c\}, i = 1 \dots l$. Then sequential recommender system (SRS) f_θ usually takes it as input and outputs user embedding, which is multiplied by target item embedding \mathbf{x}_i^c through dot product to obtain the prediction logit. Finally, the binary cross entropy (BCE) loss is often adopted.

$$\min_{\theta} \mathcal{L} = \frac{1}{n} \sum_{i=1}^n \text{BCE}(f_\theta(\{\mathbf{h}_i^c\}, \mathbf{x}_i^c), y_i) \quad (1)$$

2.2 Maximum Mean Discrepancy

Kernel function $k(\cdot, \cdot)$ characterizes how to measure similarity between samples. The idea of *kernel mean* or *mean embedding* is to represent the distribution in the reproducing kernel Hilbert space (RKHS) [39]. Specifically, considering a symmetric, positive-definite kernel k and its corresponding unique RKHS \mathcal{H}_k , each distribution $P(X)$ is mapped into \mathcal{H}_k through $\mu_P \triangleq \mathbb{E}[k(X, \cdot)] = \mathbb{E}[\varphi(X)]$ where φ is the feature map.

Furthermore, the idea of Maximum Mean Discrepancy (MMD) is to represent the distance between distributions as the distance between kernel mean of features [44]. One can define the distance between probability distributions P and Q in RKHS \mathcal{H}_k as:

$$D_k(P, Q) \triangleq \|\mu_P - \mu_Q\|_{\mathcal{H}_k}, \quad (2)$$

where $\|\cdot\|_{\mathcal{H}_k}$ is the norm of \mathcal{H}_k and this metric is known as MMD [44]. As illustrated by previous works [16], the choice of the kernel will affect the power of the two-sample test, whose null hypothesis is $P = Q$. If the positive definite kernel k is characteristic, *i.e.*, the mapping $P \mapsto \mu_P \in \mathcal{H}_k$ is injective, the kernel mean is proven to preserve all information of the distribution $P(X)$ [12]. Besides, the following multi-kernel MMD (MK-MMD) [15] is introduced to improve test power where k is a kernel function in the function space. Formally,

$$\mathcal{K} \triangleq \{k = \sum_{u=1}^m \beta_u k_u, \sum_{u=1}^m \beta_u = d, \beta_u \geq 0, \forall u\} \quad (3)$$

for some $d \geq 0$ where $\{k_u\}_{u=1}^m$ are positive definite single kernels.

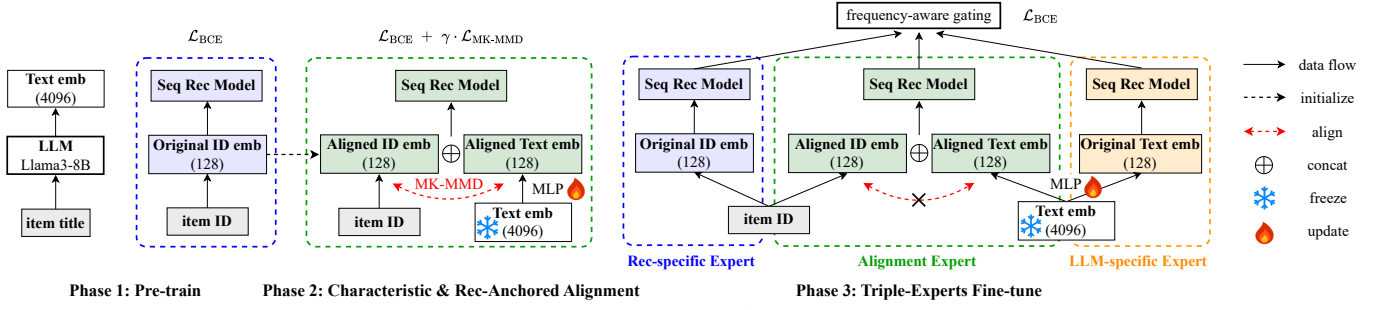


Figure 1: Overall framework of PAD. The number in parentheses (128 and 4096) denotes the embedding dimension. The prediction logit is calculated by multiplying the sequence embedding (output of recommendation model) and target item embedding. For simplicity the multiplication operation is omitted.

3 Method

In this section, we first briefly present the three phases in Sec. 3.1, and then provide a detailed description of the align and disentangle phase in Sec. 3.2 and Sec. 3.3, respectively.

3.1 Overview

The overview process of our model is depicted in Fig. 1, which consists of the following three phases:

Phase 1: LLM & Recommendation Pre-train. First, the textual embedding is generated from the textual information of items like titles and descriptions, and then it is frozen as fixed semantic knowledge. We adopt LLM2Vec [2] which transforms the primary large language models (LLMs) with decoder-only structure like Llama3 into powerful text encoders through fine-tuning. Next, the recommendation (ID) expert is pre-trained only on item ID to capture collaborative information using a SASRec model.

Phase 2: Characteristic & Rec-Anchored Alignment. After obtaining the textual and collaborative embeddings from the pre-trained LLM and recommendation model, we will align these two embeddings with each training sample. In particular, we build an alignment expert, which takes both textual and collaborative embedding as inputs, and adopts MK-MMD [12] as the alignment loss with characteristic kernels, thus being able to preserve all information about the distribution [13]. In addition, we introduce a Binary Cross Entropy loss regarding the recommendation label so as to present the collaborative embeddings from catastrophic forgetting [26] during alignment.

Phase 3: Collaborative Fine-tuning. In this phase, besides the alignment expert, we introduce two modality-specific experts, *i.e.*, one takes only the textual embedding as input, while the other takes the pre-trained collaborative embedding as input. These three experts are combined via a Mixture-of-Experts (MoE) architecture through a frequency-aware gating mechanism and then fine-tuned by the collaborative supervision signals.

3.2 Characteristic & Rec-Anchored Alignment

Existing methods [28] mainly adopt non-characteristic kernels and alignment loss is usually the only optimization objective. Nonetheless, they would fail to grasp all statistics of the data distribution and suffer catastrophic forgetting which will be detailed in Sec. 4.4. Therefore, to tackle these limitations, we employ an alignment

expert to align the textual embeddings towards the collaborative space. Specifically, we adopt MK-MMD [16, 44] with characteristic kernels as the alignment loss since it's able to capture all information about the distribution [12, 39]. Besides, we introduce an auxiliary recommendation Binary Cross Entropy loss to conduct Rec-Anchored alignment. It could enable the collaborative embeddings keep as much recommendation semantics as possible after the alignment, also avoiding catastrophic forgetting. The collaborative embeddings are not frozen during the alignment stage. By contrast, the pre-trained text embeddings remain frozen while the MLPs for dimension reduction are learnable. The overall loss function is:

$$\mathcal{L} = \mathcal{L}_{\text{REC}} + \gamma \cdot \mathcal{L}_{\text{MK-MMD}} \quad (4)$$

$$\mathcal{L}_{\text{MK-MMD}} = D_k^2(\{\mathbf{h}_i^s\}_a, \{\mathbf{h}_i^c\}_a) \quad (5)$$

$$\mathcal{L}_{\text{REC}} = \frac{1}{n} \sum_{i=1}^n \text{BCE} \left(f_{\theta} \left(\{\mathbf{h}_i^s\}_a, \{\mathbf{h}_i^c\}_a, \mathbf{x}_{i_a}^s, \mathbf{x}_{i_a}^c \right), y_i \right) \quad (6)$$

$$\{\mathbf{h}_i^s\}_a = f_{\text{MLP}}(\text{SG}(\{\mathbf{h}_{\text{text}}\}), \mathbf{w}) \quad (7)$$

where γ is the hyper-parameter. The subscript a denotes the newly aligned embedding. k is the combination of multiple characteristic kernels. f_{θ} is the recommender system taking both collaborative and text data as input. $\{\mathbf{h}_{\text{text}}\}$ denotes the textual embedding from LLM2Vec. SG denotes the stop gradient operation on the textual embedding. f_{MLP} is a learnable fully connected network with parameter \mathbf{w} reducing the dimension of textual embedding. Other notations have been illustrated at the beginning of Sec. 2.

3.3 Collaborative Fine-tuning

Existing works [8, 61] simply append the aligned embeddings to the input of recommendation model and regard it as additional features. Nevertheless, these embeddings may not be fully comprehended and exploited by the model. Therefore, in this stage, in order to fully utilize the aligned representation to enhance the performance of recommendation tasks, we propose a triple-experts architecture with four embedding tables (*i.e.*, two for each modality). Specifically, it consists of an alignment expert, an LLM-specific expert, and a recommendation-specific expert. Besides, existing works [30, 66] neglect the impact of item frequency on the credibility of different modality information. Therefore, we propose to fuse the output of these three experts via a frequency-aware gating.

Triple-Experts Architecture. To fully exploit different modality data and alleviate catastrophic forgetting, our model consists of an

aligned expert f_{align} , which takes the aligned $\{\mathbf{h}_i^s\}_a$ and $\{\mathbf{h}_i^c\}_a$ as input, and two modality-specific experts f_{LLM} and f_{id} , which take only the original textual embedding $\{\mathbf{h}_i^s\}$ or pre-trained collaborative embedding $\{\mathbf{h}_i^c\}$ as input. The MK-MMD alignment loss is removed in the alignment expert in this phase and \mathcal{L}_{REC} only is the loss function. Besides, the textual embeddings in both the alignment and LLM-specific experts are frozen, while the collaborative embeddings in the alignment and recommendation-specific experts are being updated. Finally, the output of each expert is obtained:

$$o_{\text{id}} = f_{\text{id}}(\{\mathbf{h}_i^c\}) \quad (8)$$

$$o_{\text{align}} = f_{\text{align}}(f_{\text{MLP}}(\text{SG}(\{\mathbf{h}_i^s\}_a), \mathbf{w}), \{\mathbf{h}_i^c\}_a) \quad (9)$$

$$o_{\text{LLM}} = f_{\text{LLM}}(f_{\text{MLP}}(\text{SG}(\{\mathbf{h}_i^s\}), \mathbf{w}')) \quad (10)$$

Such a design shares spirits with the Multi-Embedding paradigm [17, 31, 48] in the sense that for each modality we have two embedding tables, one for the alignment expert and the other one for the recommendation- or LLM-specific expert.

Frequency-aware Gating. We propose to fuse the output of the three experts based on the target item frequency in an adaptive manner. To be specific, we first divide all items into B buckets based on their frequency. Next, given the target item i and its corresponding bucket ID $b(i)$, a gating network g is learned which takes $b(i)$ and the expert embedding as input to generate the probability for each expert. Finally, we fuse the output of each expert with a frequency-aware gating mechanism to generate the final logit:

$$\begin{aligned} \Phi &= g_{\text{id}}(b(i), \{\mathbf{h}_i^c\}) \cdot o_{\text{id}} \\ &+ g_{\text{align}}(b(i), \{\mathbf{h}_i^s\}_a, \{\mathbf{h}_i^c\}_a) \cdot o_{\text{align}} \\ &+ g_{\text{LLM}}(b(i), \{\mathbf{h}_i^s\}) \cdot o_{\text{LLM}} \end{aligned} \quad (11)$$

where g_{id} , g_{align} , and g_{LLM} denote the probability of recommendation-specific, alignment, and LLM-specific expert, respectively.

4 Experiments

We conduct extensive experiments on three public datasets and answer the following six research questions:

- **RQ1:** How does PAD perform, compared with other SOTA baseline methods?
- **RQ2:** What is the effect of recommendation anchoring in the alignment loss?
- **RQ3:** How do the characteristic alignment losses perform, compared with non-characteristic ones?
- **RQ4:** How do the modality-specific embedding and experts contribute to the performance enhancement?
- **RQ5:** What is the impact of each expert (ID, alignment, and LLM-specific expert) and the proposed frequency-aware fusion?
- **RQ6:** Do large language models (LLMs) have more powerful encoding capability than pre-trained language models (PLMs)?
- **RQ7:** Is PAD compatible with other sequential recommendation models as a model-agnostic paradigm?

4.1 Experimental Settings

4.1.1 Datasets. Our experiments are conducted on three datasets: MIND [59], Electronics, and Prime Pantry where the last two are two categories from Amazon [40]. Their statistics are summarized

Table 1: The statistics of three public datasets: MIND, Electronics, and Prime Pantry category of Amazon.

Dataset	Users	Items	Interactions
MIND	630,235	79,707	9,667,540
Electronics	598,307	423,191	5,137,265
Prime Pantry	12,871	8,347	115,004

in Tab. 1. The detailed data preprocess and split procedure is demonstrated in Appendix. A.

- **MIND**² is collected from the Microsoft news website. It contains abundant text information with news title, abstract, body, etc. Given the historical click events of user, the task is to predict whether this user would click the target news.
- **Amazon**³ is collected from the e-commerce platform Amazon in which users rate items from 1 to 5. The sequential recommendation task is to predict whether a user will give a rating higher than 3 to the target item.

4.1.2 Evaluation Metrics. We choose the widely-used top-k Hit Ratio (HR@k) and top-k normalized Discounted Cumulative Gain (nDCG@k) with $k = 10$ for evaluation on the whole item set. The values reported below are averaged over all users.

4.1.3 Baselines. The following representative baseline methods are chosen for comparison.

- **SASRec** [23] denotes the primary SASRec model, which is only trained by pure ID of items in the collaborative space.
- **Hybrid** denotes Hybrid Encoding adopted in CoLLM [66] and LLaRA [30]. It directly concatenates the collaborative with text embedding and feeds into the subsequent SRS.
- **MoRec** [62] simply adopts an MLP to reduce the dimension of text embedding and inputs it into SRS.
- **CTRL** [28] first pre-trains the model parameters by aligning the collaborative with text embedding using contrastive learning, then fine-tunes on the recommendation task.
- **MAKE** [45] adopts a two-experts structure. It first pre-trains the LLM-specific expert in the same way as MoRec, then fuse the output of ID and LLM-specific expert to generate prediction.
- **DisCo** [8] splits the embeddings into chunks and incorporates sufficiency and disentanglement loss to explicitly preserve the task-related and unique information.
- **SMEM** is a novel baseline we propose, which adopts Shared and Modality-specific EMbeddings. It is similar to Shared and Task-specific EMbeddings (STEM) [48] in multi-task recommendation.

4.1.4 Implementation Details. We implement PAD with multiple Gaussian kernels:

$$k_u(x, x') = \exp\left(-\frac{\|x - x'\|^2}{2\sigma^2}\right) \quad (12)$$

where $k_u \in \mathcal{K}$ in (3). We take $m = 5$ and $\sigma = \{-3, -2, -1, 0, 1\}$ in the following experiments. Besides, our experiments are conducted on Tesla V100 GPUs and all the results shown are averaged over 3 runs. Detailed experimental settings are provided in Appendix. A.

²<https://msnews.github.io/>

³https://cseweb.ucsd.edu/~jmcauley/datasets/amazon_v2/

Table 2: Overall performance comparison on MIND, Electronics, and Prime Pantry dataset. Boldface denotes the highest value while underline indicates the second best result. ‘Impr.’ indicates our improvement against the original SASRec. ★ represents statistical significance with p -value < 0.05 in t -test compared with the best baseline.

Datasets	Metric	SASRec	Hybrid	MoRec	CTRL	MAKE	DisCo	SMEM	Ours	Impr.
MIND	HR@10	16.8437	17.3385	14.1235	18.1318	16.4317	17.3367	<u>18.4319</u>	18.6703★	10.84%
	nDCG@10	9.0520	9.3321	7.7090	9.7996	8.9684	9.4131	<u>10.0003</u>	10.1515★	12.15%
Electronics	HR@10	1.6754	1.5673	0.9562	1.9468	0.9495	1.8148	<u>2.3492</u>	2.4804★	48.05%
	nDCG@10	0.7938	0.8385	0.4721	0.9370	0.4611	0.9025	<u>1.4287</u>	1.5373★	93.67%
Prime Pantry	HR@10	2.6338	3.1000	3.1699	2.8048	3.3408	2.8358	<u>3.4341</u>	3.8303★	45.43%
	nDCG@10	1.2926	1.5606	1.6013	1.3700	1.6379	1.3451	<u>1.7440</u>	1.9104★	47.80%




Figure 2: Comparison of the original SASRec and PAD on the Mind (left), Electronics (mid) and Prime Pantry (right) datasets. Warm, median, cold denote target items with high to low frequency on the test set. The y-axis denotes the HR@10.

4.2 Overall Performance (RQ1)

To answer **RQ1**, we verify PAD’s effectiveness by comparing it with various baseline methods introduced in Sec. 4.1.3. The overall performance on three public datasets is shown in Tab. 2. In summary, *on all three datasets, our proposed PAD achieves state-of-the-art performance, beating the best performing baseline SMEM by 1.51%, 7.60%, and 9.54% on nDCG@10 on each dataset.*

Moreover, in Fig. 2 we also provide the results on different subsets of items, *i.e.*, on warm, median, and cold items. We observe that PAD surpasses SASRec on all the item subsets on each dataset, and the improvement on cold items is more significant. For example, on Electronics dataset, our method achieves 63.36%, 202.02% and 462.61% relative performance lift on nDCG@10 on the warm, median and cold items. The performance lift on the cold items are 7.3 times of that on the warm items. *This validates that our method can mitigate the cold-start problem with the LLM knowledge effectively.*

We further illustrate the effectiveness of our method on the cold items with the following analysis. Denote $\mathcal{P}_{\text{Top-10\%}}^{\text{ID}}$ and $\mathcal{P}_{\text{Bottom-10\%}}^{\text{ID}}$ as the group of item pairs with the top-10% and bottom-10% distance based on collaborative embeddings. We present the distribution of these item pairs under the distance distribution regarding the collaborative embeddings and textual embeddings in each model in Fig. 3. We observe that all existing methods, including SASRec, SMEM, and CTRL, can not differentiate the distribution of $\mathcal{P}_{\text{Top-10\%}}^{\text{ID}}$ and $\mathcal{P}_{\text{Bottom-10\%}}^{\text{ID}}$ well. In contrast, *our method succeeds in learning a generally larger textual distance for those item pairs with top collaborative distance and learning a smaller textual distance for those with bottom collaborative distance.* Consequently, we conclude with the following result:

Result 1. Our proposed PAD greatly improves the performance on the cold items with better alignment of textual embeddings towards the collaborative space, leading to SOTA performance on three datasets.

4.3 Discrepancy Measurement

In this section, we propose a novel metric to quantify to what extent are the aligned embedding space deviates from the original one. Specifically, rather than the deviation of the aligned and original embedding for each ID itself, *we are more interested in the discrepancy of embedding distance* [48]. To this end, we employ the Kendall’s tau [24] between the embedding distance distribution of the original embeddings and that of the aligned embeddings.

Specifically, in sequential recommendation, the embedding distance between each behavior and target item pair is calculated. Afterward, Kendall’s tau is adopted to measure the degree of concordance between these two variables. For example, given the user’s historical interaction sequence $\{i_1, i_2, i_3\}$, the SRS usually takes $\{i_1\}$ as the behavior item sequence to predict the target item i_2 and takes $\{i_1, i_2\}$ as the behavior item sequence to predict the target item i_3 . We map each item i_s into embedding \mathbf{e}_s and calculate the distance $\langle \cdot, \cdot \rangle$ (like Euclidean distance) between each behavior and target item embedding pair, which includes $\{\langle \mathbf{e}_1, \mathbf{e}_2 \rangle, \langle \mathbf{e}_1, \mathbf{e}_3 \rangle, \langle \mathbf{e}_2, \mathbf{e}_3 \rangle\}$. For another model with embeddings \mathbf{e}'_s for i_s , its distance variable is $\{\langle \mathbf{e}'_1, \mathbf{e}'_2 \rangle, \langle \mathbf{e}'_1, \mathbf{e}'_3 \rangle, \langle \mathbf{e}'_2, \mathbf{e}'_3 \rangle\}$. Any pair of samples is called concordant if the sort order agrees, *e.g.*, both $\langle \mathbf{e}_1, \mathbf{e}_2 \rangle < \langle \mathbf{e}_1, \mathbf{e}_3 \rangle$ and $\langle \mathbf{e}'_1, \mathbf{e}'_2 \rangle < \langle \mathbf{e}'_1, \mathbf{e}'_3 \rangle$ hold or $\langle \mathbf{e}_1, \mathbf{e}_2 \rangle > \langle \mathbf{e}_1, \mathbf{e}_3 \rangle$ and $\langle \mathbf{e}'_1, \mathbf{e}'_2 \rangle > \langle \mathbf{e}'_1, \mathbf{e}'_3 \rangle$ hold. The Kendall’s tau is defined as

$$\text{KT} = \frac{\#(\text{concordant pairs}) - \#(\text{discordant pairs})}{\#(\text{pairs})} \quad (13)$$

where $\#$ denotes the number. Notably, this method of measurement is also applicable to the circumstance where the dimensions of embedding are inconsistent between models or modalities, and it can be easily extended to other binary relation like user-user and user-item distance.

4.4 Rec-Anchored Alignment Loss (RQ2)

We investigate the effect of recommendation anchoring in avoiding catastrophic forgetting by comparing the following three alignment losses in the align phase, namely: (1) **No Alignment**, *i.e.*, **SMEM**, which doesn’t involve any alignment loss but simply concat the textual and collaborative embeddings within the alignment expert, (2) **Non-Anchored Alignment**, with only the alignment loss, *i.e.*, $D_k^2(\mathcal{D}_{\text{text}}, \mathcal{D}_{\text{rec}})$, (3) **Rec-Anchored Alignment**, which combines the alignment loss with a BCE loss: $\mathcal{L}_{\text{rec}} + \gamma \cdot D_k^2(\mathcal{D}_{\text{text}}, \mathcal{D}_{\text{rec}})$ and updates both collaborative and text embedding, and (4) **Rec-Anchored and Frozen Alignment**, which freezes the collaborative embeddings upon the Rec-Anchored Alignment: $\mathcal{L}_{\text{rec}} + \gamma \cdot D_k^2(\mathcal{D}_{\text{text}}, \text{SG}(\mathcal{D}_{\text{rec}}))$, where SG denotes stop gradient operation.

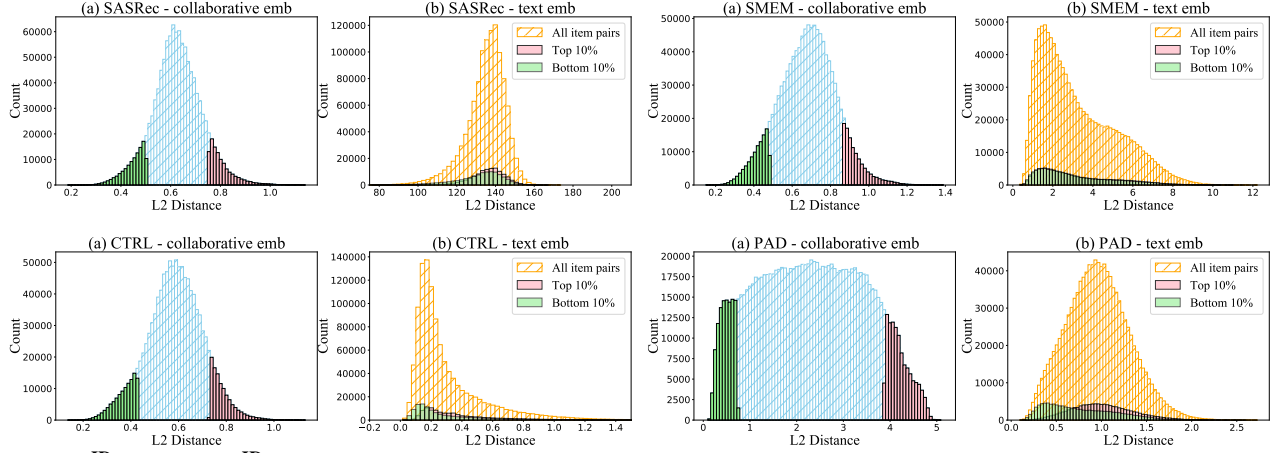


Figure 3: $\mathcal{P}_{\text{Top-10\%}}^{\text{ID}}$ and $\mathcal{P}_{\text{Bottom-10\%}}^{\text{ID}}$ under the distance distribution regarding the collaborative and textual embeddings in original SASRec, SMEM, CTRL, and our proposed PAD on cold items.




Figure 4: Comparison of anchored and non-anchored alignment losses on the Mind (left), Electronics (mid) and Prime Pantry (right) datasets. The y-axis denotes the HR@10.

This can totally avoid catastrophic forgetting of the collaborative embeddings since they are frozen. However, this may hurt the alignment of the textual embeddings towards the collaborative space.

The results are shown in Fig. 4, and we can observe that: 1) Our proposed Rec-Anchored Alignment performs the best among all losses; 2) Rec-Anchored losses, including Rec-Anchored as well as Rec-Anchored and Frozen, perform in general better than those non-anchored method, indicating that recommendation anchoring can effectively avoid catastrophic forgetting; 3) The Non-Anchored Alignment sometimes performs worse than Non-Alignment, indicating that simple alignment without anchoring even hurt the performance; 4) Rec-Anchored Alignment performs better than Rec-Anchored and Frozen Alignment, indicating the brute force of freezing collaborative embeddings hurts the alignment of the textual embeddings;

We further illustrate such forgetting of Non-Anchored as the violet lines in Fig. 6(a) with the tool introduced in Sec. 4.3, which show those cold item IDs (*i.e.*, with bucket 7-10) has their embedding distance distribution more diverged from the original embeddings.

Result 2. Our proposed Rec-Anchored Alignment loss avoids the catastrophic forgetting of collaborative embeddings, and succeeds in aligning textual embeddings towards the collaborative space, achieving the best performance.

4.5 Characteristic Alignment Loss (RQ3)

To verify the advantages of characteristic kernels over non-characteristic ones in maximum mean discrepancy, apart from Gaussian kernels, we also experiment on the following three kernels and Info_NCE




Figure 5: (a) Comparison of alignment losses and (b) ablation study comparing different model variants on MIND dataset.

loss (usually adopted by contrastive learning), in which only Laplacian kernels are characteristic on \mathbb{R}^d .

- **Linear kernel**

$$k(x, x') = \langle x, x' \rangle = x^\top x' \quad (14)$$

- **Cosine kernel**

$$k(x, x') = 1 - \left\langle \frac{x}{\|x\|}, \frac{x'}{\|x'\|} \right\rangle \quad (15)$$

- **Info_NCE**

$$\mathcal{L}_{CL} = - \sum_{i=1}^n \log \frac{\exp(k(x_i, x'_i) / \tau)}{\exp(k(x_i, x'_i) / \tau) + \sum_{j \neq i} \exp(k(x_i, x'_j) / \tau)} \quad (16)$$

where τ is the temperature coefficient and cosine kernel is usually adopted to measure the sample similarity.

- **Laplacian kernel**

$$k(x, x') = \exp\left(-\frac{\|x - x'\|_1}{\sigma^2}\right) \quad (17)$$

Their results are shown in Fig. 5(a), and we can see that characteristic kernels (*i.e.*, Laplacian, and Gaussian) indeed perform better than non-characteristic ones. Notably, even if Info_NCE explicitly pulls positive samples closer and negative samples farther thus achieving better results than pure cosine kernel, it is still inferior to characteristic kernels because the mean embedding of its RKHS is not rich enough. Therefore, we conclude that:

Result 3. Characteristic kernels perform better than non-characteristic kernels in aligning textual embeddings towards collaborative space by maximum mean discrepancy.

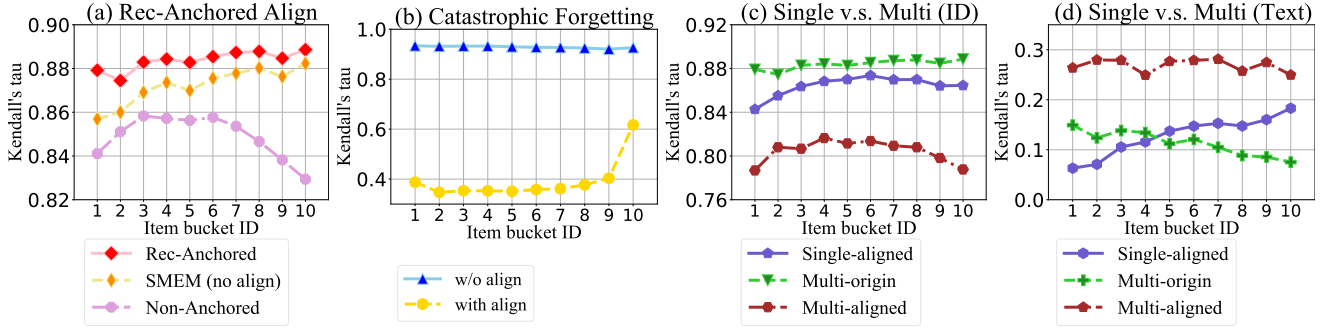


Figure 6: Kendall’s tau between the Euclidean distance (same as L2 distance) of (a) collaborative embedding of different anchored alignment losses (b) collaborative embedding of single-expert structure with or without alignment. (c) collaborative embedding of single-expert and triple-experts (d) text embedding of single-expert and triple-experts. They all use SASRec as backbone on MIND dataset. Item bucket ID from 1 to 10 denotes warm to cold.

4.6 Towards Triple-Experts Architecture (RQ4)

4.6.1 Catastrophic Forgetting of Alignment. Many existing LLM4Rec methods [28] first align the collaborative embeddings of recommendation with the LLM, and then fine-tune these embeddings in the downstream recommendation task. However, we wonder whether there is catastrophic forgetting of the alignment on the collaborative embeddings. That is, when there is only one single recommendation embedding table, whether the alignment leads to catastrophic damage to these embeddings, which *CANNOT be recovered by the supervised fine-tuning*.

To validate this, we implement three methods, which all have only one single collaborative embedding table but differ in: the first method only employs the Binary Cross Entropy (BCE) loss upon the collaborative signal for supervised learning, without any alignment with the LLM, named SE-BCE in short; the second method employs BCE loss for both pre-training and fine-tuning, named SE-BCE-Tune; and the third method employs both BCE loss and the alignment loss for pre-training, then conducts a supervised fine-tuning using the collaborative labels, named SE-BCE-Align-Tune.

We measure the Kendall’s tau of (SE-BCE, SE-BCE-Tune) and (SE-BCE, SE-BCE-Align-Tune) on the collaborative embedding distance of the methods and the results are shown in Fig. 6(b). We divide the items into 10 buckets based on their frequency in the training set. The bucket ID range from 1 to 10, denoting from warm (high frequency) to cold (low frequency) items. We observe that SE-BCE-Align-Tune (*i.e.*, the ‘with align’ line in Fig. 6(b)) has a much smaller value than the SE-BCE-Tune (*i.e.*, the ‘w/o align’ line in Fig. 6(b)) on all buckets. In particular, Kendall’s tau on buckets 1 to 9 is less than 0.5, indicating that more than 50% of the partial orders are disturbed. Only the 10-th bucket gets Kendall’s tau of 0.77.

Result 4. With only one set of collaborative embeddings, the alignment loss leads to catastrophic forgetting on the collaborative embeddings.

The catastrophic forgetting of collaborative embeddings makes it hard to rely only on the aligned collaborative embeddings in the downstream recommendation tasks. To this end, we propose to incorporate two modality-specific experts in addition to the alignment expert, leading to a Triple-Experts architecture below.

4.6.2 Triple-Experts Architecture. To illustrate the effect and disentanglement capability of triple-experts architecture, on MIND dataset and SASRec as the backbone model, we implement SE-BCE, SE-BCE-Align-Tune, and TE-BCE-Align-Tune whose training objective is the same as SE-BCE-Align-Tune but implemented on triple-experts with two embedding tables for both ID and text embedding. We measure the Kendall’s tau of (SE-BCE, SE-BCE-Align-Tune) (*i.e.*, the ‘Single-aligned’ line in Fig. 6(c) and (d)) and (SE-BCE, TE-BCE-Align-Tune) (*i.e.*, the ‘Multi-origin’ and ‘Multi-aligned’ line in Fig. 6(c) and (d)). The results of ID and text embedding are shown in Fig. 6(c) and (d), where ‘Single-aligned’ denotes the ID or text embedding of SE-BCE-Align-Tune. ‘Multi-origin’ and ‘Multi-aligned’ denote two sets of ID or text embeddings of triple-experts. We have the following observations.

First, the Kendall’s tau between the distance of pre-trained and updated collaborative embedding is generally larger than that between the distance of pre-trained and updated text embedding. This is because the text data act as auxiliary information of recommendation, thus the aligned space is more closer to the original collaborative space than the semantic space. Second, the Kendall’s tau of at least one set of embedding in triple-experts (yellow and purple line) is consistently larger than that of embedding in single-expert structure (red line) on all item buckets. Third, TE-BCE-Align-Tune surpasses SE-BCE-Align-Tune by 3.59% on nDCG@10 on MIND dataset. Therefore, we can conclude that too much deviation of the aligned collaborative embedding distribution from the original one has negative effect on the recommendation performance. A possible reason is some unique information (*e.g.*, the rank of distance between all behavior and item embedding pairs) in the original embedding distribution is lost. These results indicate that compared with single-expert, triple-experts architecture could retain more information of the original embedding, thus being capable of alleviating the catastrophic forgetting and plays a vital role of disentanglement leading to performance enhancement.

4.7 Ablation Study (RQ5)

We conduct an ablation study on the MIND dataset to investigate the effect of each expert and the frequency-aware gating. The results are shown in Fig. 5(b). First, we remove each expert in the last phase and observe the performance deteriorates when we remove

Table 3: Compatibility experiments on MIND, Electronics, and Prime Pantry dataset. ‘Origin’ denotes the model trained on ID only and ‘Impr.’ indicates the improvement against the original backbone model.

Datasets	GRU4Rec						Caser					
	Origin		+PAD		Impr.		Origin		+PAD		Impr.	
	HR@10	nDCG@10	HR@10	nDCG@10	HR@10	nDCG@10	HR@10	nDCG@10	HR@10	nDCG@10	HR@10	nDCG@10
MIND	13.3969	6.9098	16.5762	8.8475	23.73%	28.04%	8.6879	3.7981	9.0569	4.2696	4.25%	12.41%
Electronics	0.8437	0.4104	1.1016	0.5898	30.57%	43.73%	0.5746	0.2355	0.6351	0.2501	10.53%	6.20%
Prime Pantry	1.7248	0.8227	2.5484	1.2307	47.75%	49.59%	0.6526	0.4136	1.4063	0.5507	115.49%	33.16%

any of the three experts. Interestingly, removing the recommendation (ID) expert leads to the largest performance drop. This indicates that the alignment loss leads to inevitable information loss of the collaborative embeddings, even with our Rec-Anchored loss. Therefore, the recommendation-specific tower is essential to retain the collaborative signals.

Second, we try to retain only one of these three expert and compare their performance. Surprisingly, the single alignment model achieves the best performance. This is possibly due to the fact that, on the one hand, it keeps the collaborative semantics through the Rec-Anchored alignment loss, and on the other hand, it boosts the performance on the cold items with the help of LLM semantics.

Finally, the model variant removing the frequency-aware gating fuses all experts with the learned weights on all items equally, leading to a decrease of 3.06% in HR@10. Therefore, from the perspective of target item frequency, this gating acts as a bridge connecting the collaborative, semantic, and alignment space adaptively considering their different credibility.

4.8 LLM2Vec v.s. BERT (RQ6)

To compare the performance of LLM and pre-train language model (PLM) like BERT, we replace LLM2Vec with BERT-Base as the text encoder of items. It achieves the performance of HR@10 = 18.3989 and nDCG@10 = 9.9810 on MIND, which is inferior to LLM2Vec achieving HR@10 = 18.6703 and nDCG@10 = 10.1515. We speculate the reasons are: 1) More parameters of Llama3-8B bring in more powerful encoding capability than BERT-Base with 110 million parameters only. 2) Compared with BERT, LLM2Vec is trained with more abundant text data on wiki corpus and adopts a different Unsupervised contrastive training (SimCSE) objective.

4.9 Compatibility (RQ7)

To validate the compatibility of PAD, apart from SASRec [23], we also experiment on two typical sequential recommendation models as backbones including GRU4Rec [21] and Caser [50]. As shown in Tab. 3, PAD significantly enhance the performance of the original model on all three datasets, indicating that it acts as a powerful model-agnostic enhancement paradigm.

5 Related Work

This section summarizes the related work on sequential recommendation, large language model, and multi-modal recommendation.

5.1 Sequential Recommendation

Sequential recommendation (SR) focuses on modeling the sequential dependency over the interaction sequence to capture user preference. Inspired by the success of self-attention in natural language

processing [51], SASRec [23] incorporates it into SR and afterward many works explore target-attention recommendation in a similar manner [4, 10, 11, 42, 46, 71–73]. Caser [50] and GRU4Rec [21] adopts convolution and recurrent networks for sequential modeling. Nevertheless, these models mainly takes ID as the only input thus suffering from the cold-start problem. By contrast, our proposed PAD paradigm is capable of enhancing them as backbones by capturing complex correlation of different modality data through our designed alignment and triple-experts structure.

5.2 Large Language Model & Multi-Modal Rec

In the recommendation community, some works explicitly adopt alignment between tabular and textual data which could achieve information gain. For example, CTRL [28] leverages contrastive learning, and FLIP [52] uses data masking and reconstruction pre-training with contrastive learning. Afterward, some works like DisCo [8] and DaRec [61] conduct disentanglement to capture the specific information. However, they simply separate the embeddings into different parts which may not be fully learned by the common single-expert structure. By contrast, our PAD adopts triple-experts in a frequency-aware manner to adaptively fuse different information, which achieves better disentanglement ability.

Meanwhile, numerous studies [6, 7, 18, 29, 32, 43, 47, 49, 65] focused on enhancing recommender systems by incorporating multi-modal content. Recently, researchers have explored capturing user fine-grained preferences across different modalities. For instance, MMGCN [58] attempts to model user preferences using modal-specific user-item bipartite graphs. M3SRec [3] proposes a novel multi-modal mixture-of-experts (MoE) fusion network. Taobao [45] proposes a two-stage paradigm to enhance recommendation with multi-modality recommendation. M3CSR [5] proposes a multi-modality framework for to tackle cold-start in short-video recommendation. However, the existing methods either use a non-characteristic alignment loss or don’t employ a disentangled expert, while our method has both of these two critical parts.

6 Conclusion

In this paper, we propose a novel Pre-train, Align, and Disentangle (PAD) paradigm to empower sequential recommendation with large language models. We propose a characteristic recommendation-anchored loss to better align textual embeddings towards the collaborative space and avoid catastrophic forgetting. We employ a triple-experts architecture, consisting of aligned and modality-specific experts with disentangled embeddings, which is fine-tuned in a frequency-aware manner. Comprehensive experiments on three public datasets validate the effectiveness of our proposed method. Notably, the framework can be extended to multi-modal modeling.

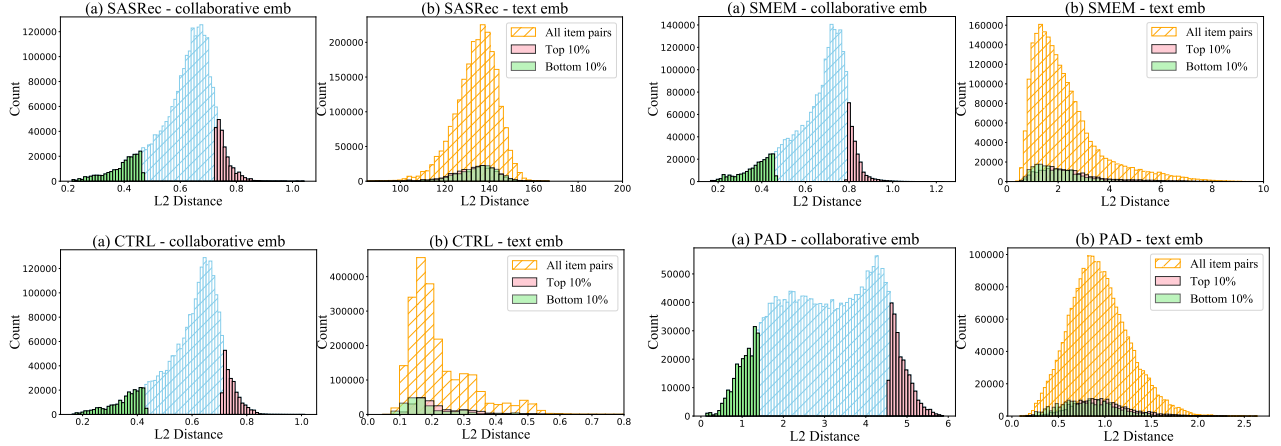


Figure 7: $\mathcal{P}_{\text{Top-10\%}}^{\text{ID}}$ and $\mathcal{P}_{\text{Bottom-10\%}}^{\text{ID}}$ under the distance distribution regarding the collaborative and text embeddings in the original SASRec, SMEM, CTRL, and PAD on warm items.

ACKNOWLEDGEMENT

This research was partially supported by Tencent Rhino-Bird Focused Research Program, Research Impact Fund (No.R1015-23), and Collaborative Research Fund (No.C1043-24GF).

A Experimental Settings

For LLM2Vec we choose the Llama3-8B as the base model with the unsupervised-trained LoRA weights. We leverage the news title on MIND and item title on Amazon to generate text embedding. Suppose l is defined as the length of user’s historical interaction sequence. For all the three datasets, we select the users with no less than 5 interactions and extract the latest 23 items interacted. The task is to predict whether the user would click the $(l - 2)$ -th, $(l - 1)$ -th, and l -th item on training, validation, and test set, respectively, given the previous interacted items. Besides, AdamW [37] is adopted as optimizer. Batch size is 16, the early stop epoch is set to 10, and dropout probability is 0.1. L_2 regularization is adopted with weight of 0.1. The weight γ is searched and set to 0.2.

B Pseudo-code

The general procedure of PAD is given in Alg. 1, which consists of three phases: (1) LLM & Recommendation Pre-train (from Line 1 to 6), (2) Alignment with MK-MMD (from Line 7 to 14), and (3) Recommendation Supervised Fine-tuning (from Line 15 to 23).

C Advanced Analysis

Denote $\mathcal{P}_{\text{Top-10\%}}^{\text{ID}}$ and $\mathcal{P}_{\text{Bottom-10\%}}^{\text{ID}}$ as the group of item pairs with the top-10% and bottom-10% distance based on collaborative ID embeddings. Apart from the results on cold items in Fig. 3, we present the distribution of these item pairs under the distance distribution regarding the collaborative and textual embeddings learned by each model on warm items in Fig. 7. It indicates that the distribution of $\mathcal{P}_{\text{Top-10\%}}^{\text{ID}}$ and $\mathcal{P}_{\text{Bottom-10\%}}^{\text{ID}}$ can not be differentiated in textual modality probably because all these recommendation models mainly rely on collaborative information for warm items.

Algorithm 1: PAD paradigm for sequential recommender systems (SRS)

Input: User set \mathcal{U} ; item set \mathcal{I} ; historical interaction sequence; item text embedding; true label y ;

Output: A well trained SRS with three experts.

Phase 1: Pre-train

```

1 while not converge do
2   Sample a mini-batch data from  $\mathcal{U}$ ;
3   Calculate the prediction result of sequential rec model;
4   Calculate the BCE loss;
5   Take the gradient and update sequential rec model;
6 end
Phase 2: Rec-Anchored Alignment
7 Load pre-trained ID embedding
8 while not converge do
9   Sample a mini-batch data from  $\mathcal{U}$ ;
10  Retrieve text embedding of behavior items;
11  Calculate the prediction result of alignment expert;
12  Calculate the BCE loss plus MMD;
13  Take the gradient and update alignment expert;
14 end
Phase 3: Triple-Experts Fine-tune
15 Load parameters of pre-trained Rec-specific & Alignment expert
16 while not converge do
17   Sample a mini-batch data from  $\mathcal{U}$ ;
18   Retrieve text embedding of behavior items;
19   Calculate the prediction result of all three experts;
20   Fuse the results via frequency-aware gating;
21   Calculate the BCE loss;
22   Take the gradient and update all parameters;
23 end

```

References

- [1] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1007–1014.
- [2] Parishad BehnamGhader, Vaibhav Adlakha, Marius Mosbach, Dzmitry Bahdanau, Nicolas Chapados, and Siva Reddy. 2024. Llm2vec: Large language models are secretly powerful text encoders. *arXiv preprint arXiv:2404.05961* (2024).
- [3] Shuqing Bian, Xingyu Pan, Wayne Xin Zhao, Jinpeng Wang, Chuyuan Wang, and Ji-Rong Wen. 2023. Multi-modal mixture of experts representation learning for sequential recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 110–119.
- [4] Jianxin Chang, Chenbin Zhang, Zhiyi Fu, Xiaoxue Zang, Lin Guan, Jing Lu, Yiqun Hui, Dewei Leng, Yanan Niu, Yang Song, et al. 2023. TWIN: Two-stage interest network for lifelong user behavior modeling in CTR prediction at kuaishou. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3785–3794.
- [5] Gaode Chen, Ruina Sun, Yuezhan Jiang, Jiangxia Cao, Qi Zhang, Jingjian Lin, Han Li, Kun Gai, and Xinghua Zhang. 2024. A Multi-modal Modeling Framework for Cold-start Short-video Recommendation. In *Proceedings of the 18th ACM Conference on Recommender Systems*. 391–400.
- [6] Jingyuan Chen, Hanwang Zhang, Xiangnan He, Liqiang Nie, Wei Liu, and Tat-Seng Chua. 2017. Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*. 335–344.
- [7] Xu Chen, Hanxiong Chen, Hongteng Xu, Yongfeng Zhang, Yixin Cao, Zheng Qin, and Hongyuan Zha. 2019. Personalized fashion recommendation with visual explanations based on multimodal attention network: Towards visually explainable recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 765–774.
- [8] Kounianhua Du, Jizheng Chen, Jianghao Lin, Yunxia Xi, Hangyu Wang, Xinyi Dai, Bo Chen, Ruiming Tang, and Weinan Zhang. 2024. DisCo: Towards Harmonious Disentanglement and Collaboration between Tabular and Semantic Space for Recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 666–676.
- [9] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
- [10] Ningya Feng, Junwei Pan, Jialong Wu, Baixu Chen, Ximei Wang, Qian Li, Xian Hu, Jie Jiang, and Mingsheng Long. 2024. Long-Sequence Recommendation Models Need Decoupled Embeddings. *arXiv preprint arXiv:2410.02604* (2024).
- [11] Yufei Feng, Fuyu Lv, Weichen Shen, Menghan Wang, Fei Sun, Yu Zhu, and Keping Yang. 2019. Deep session interest network for click-through rate prediction. In *IJCAI*.
- [12] Kenji Fukumizu, Francis R Bach, and Michael I Jordan. 2004. Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research* 5, Jan (2004), 73–99.
- [13] Kenji Fukumizu, Arthur Gretton, Bernhard Schölkopf, and Bharath K Sriperumbudur. 2008. Characteristic kernels on groups and semigroups. *Advances in neural information processing systems* 21 (2008).
- [14] Jingtong Gao, Xiangyu Zhao, Muyang Li, Minghao Zhao, Runze Wu, Ruocheng Guo, Yiding Liu, and Dawei Yin. 2024. SMLP4Rec: an Efficient all-MLP architecture for sequential recommendations. *ACM Transactions on Information Systems* 42, 3 (2024), 1–23.
- [15] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A kernel two-sample test. *The Journal of Machine Learning Research* 13, 1 (2012), 723–773.
- [16] Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K Sriperumbudur. 2012. Optimal kernel choice for large-scale two-sample tests. *Advances in neural information processing systems* 25 (2012).
- [17] Xingzhuo Guo, Junwei Pan, Ximei Wang, Baixu Chen, Jie Jiang, and Mingsheng Long. 2024. On the Embedding Collapse when Scaling up Recommendation Models. *ICML* (2024).
- [18] Ruining He and Julian McAuley. 2016. VBPR: visual bayesian personalized ranking from implicit feedback. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 30.
- [19] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, et al. 2014. Practical lessons from predicting clicks on ads at facebook. In *International Workshop on Data Mining for Online Advertising (ADKDD)*. 1–9.
- [20] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).
- [21] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2016. SESSION-BASED RECOMMENDATIONS WITH RECURRENT NEURAL NETWORKS. In *ICLR*.
- [22] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. IEEE, 197–206.
- [23] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [24] Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika* 30, 1-2 (1938), 81–93.
- [25] Chengxi Li, Yejing Wang, Qidong Liu, Xiangyu Zhao, Wanyu Wang, Yiqi Wang, Lixin Zou, Wenqi Fan, and Qing Li. 2023. STRec: Sparse transformer for sequential recommendations. In *Proceedings of the 17th ACM conference on recommender systems*. 101–111.
- [26] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*. PMLR, 19730–19742.
- [27] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*. PMLR, 12888–12900.
- [28] Xiangyang Li, Bo Chen, Lu Hou, and Ruiming Tang. 2023. CTRL: Connect Collaborative and Language Model for CTR Prediction. *arXiv preprint arXiv:2306.02841* (2023).
- [29] Jiahao Liang, Xiangyu Zhao, Muyang Li, Zijian Zhang, Wanyu Wang, Haochen Liu, and Zitao Liu. 2023. Mmmlp: Multi-modal multilayer perceptron for sequential recommendations. In *Proceedings of the ACM Web Conference 2023*. 1109–1117.
- [30] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2023. Llara: Aligning large language models with sequential recommenders. *arXiv preprint arXiv:2312.02445* (2023).
- [31] Zhutian Lin, Junwei Pan, Haibin Yu, Xi Xiao, Ximei Wang, Zhixiang Feng, Shifeng Wen, Shudong Huang, Lei Xiao, and Jie Jiang. 2024. Crocodile: Cross Experts Covariance for Disentangled Learning in Multi-Domain Recommendation. *arXiv preprint arXiv:2405.12706* (2024).
- [32] Fan Liu, Zhiyong Cheng, Changchang Sun, Yinglong Wang, Liqiang Nie, and Mohan Kankanhalli. 2019. User diverse preference modeling by multimodal attentive metric learning. In *Proceedings of the 27th ACM international conference on multimedia*. 1526–1534.
- [33] Qidong Liu, Xian Wu, Yejing Wang, Zijian Zhang, Feng Tian, Yefeng Zheng, and Xiangyu Zhao. 2024. Llm-esr: Large language models enhancement for long-tailed sequential recommendation. *Advances in Neural Information Processing Systems* 37 (2024), 26701–26727.
- [34] Yifan Liu, Kangning Zhang, Xiangyuan Ren, Yanhua Huang, Jiarui Jin, Yingjie Qin, Ruilong Su, Ruiwen Xu, and Weinan Zhang. 2024. An Aligning and Training Framework for Multimodal Recommendations. *arXiv preprint arXiv:2403.12384* (2024).
- [35] Ziru Liu, Shuchang Liu, Zijian Zhang, Qingpeng Cai, Xiangyu Zhao, Kesen Zhao, Lantao Hu, Peng Jiang, and Kun Gai. 2024. Sequential recommendation for optimizing both immediate feedback and long-term retention. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1872–1882.
- [36] Ziru Liu, Jiejie Tian, Qingpeng Cai, Xiangyu Zhao, Jingtong Gao, Shuchang Liu, Dayou Chen, Tonghao He, Dong Zheng, Peng Jiang, et al. 2023. Multi-task recommendations with reinforcement learning. In *Proceedings of the ACM web conference 2023*. 1273–1282.
- [37] Ilya Loshchilov, Frank Hutter, et al. 2017. Fixing weight decay regularization in adam. *arXiv preprint arXiv:1711.05101* (2017).
- [38] H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, et al. 2013. Ad click prediction: a view from the trenches. In *ACM SIGKDD International conference on Knowledge Discovery & Data Mining (KDD)*. 1222–1230.
- [39] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, Bernhard Schölkopf, et al. 2017. Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning* 10, 1-2 (2017), 1–141.
- [40] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proc. of EMNLP*.
- [41] Junwei Pan, Wei Xue, Ximei Wang, Haibin Yu, Xun Liu, Shijie Quan, Xueming Qiu, Dapeng Liu, Lei Xiao, and Jie Jiang. 2024. Ads Recommendation in a Collapsed and Entangled World. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5566–5577.
- [42] Qi Pi, Guorui Zhou, Yujing Zhang, Zhe Wang, Lejian Ren, Ying Fan, Xiaoqiang Zhu, and Kun Gai. 2020. Search-based user interest modeling with lifelong sequential behavior data for click-through rate prediction. In *ACM International Conference on Information & Knowledge Management (CIKM)*. 2685–2692.
- [43] Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Tran, Jonah Samost, et al. 2024. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems* 36 (2024).
- [44] Dino Sejdinovic, Bharath Sriperumbudur, Arthur Gretton, and Kenji Fukumizu. 2013. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The annals of statistics* (2013), 2263–2291.

- [45] Xiang-Rong Sheng, Feifan Yang, Litong Gong, Biao Wang, Zhangming Chan, Yujing Zhang, Yueyao Cheng, Yong-Nan Zhu, Tiezheng Ge, Han Zhu, et al. 2024. Enhancing Taobao Display Advertising with Multimodal Representations: Challenges, Approaches and Insights. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. 4858–4865.
- [46] Zihua Si, Lin Guan, ZhongXiang Sun, Xiaoxue Zang, Jing Lu, Yiqun Hui, Xingchao Cao, Zeyu Yang, Yichen Zheng, Dewei Leng, et al. 2024. TWIN V2: Scaling Ultra-Long User Behavior Sequence Modeling for Enhanced CTR Prediction at Kuaishou. *arXiv preprint arXiv:2407.16357* (2024).
- [47] Anima Singh, Trung Vu, Nikhil Mehta, Raghunandan Keshavan, Maheswaran Sathiamoorthy, Yilin Zheng, Lichan Hong, Lukasz Heldt, Li Wei, Devansh Tandon, et al. 2024. Better Generalization with Semantic IDs: A Case Study in Ranking for Recommendations. In *Proceedings of the 18th ACM Conference on Recommender Systems*. 1039–1044.
- [48] Liangcai Su, Junwei Pan, Ximei Wang, Xi Xiao, Shijie Quan, Xihua Chen, and Jie Jiang. 2024. STEM: Unleashing the Power of Embeddings for Multi-task Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 9002–9010.
- [49] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.
- [50] Jiayi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the eleventh ACM international conference on web search and data mining*. 565–573.
- [51] A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* (2017).
- [52] Hangyu Wang, Jianghao Lin, Xiangyang Li, Bo Chen, Chenxu Zhu, Ruiming Tang, Weinan Zhang, and Yong Yu. 2023. FLIP: Towards Fine-grained Alignment between ID-based Models and Pretrained Language Models for CTR Prediction. *arXiv e-prints* (2023), arXiv–2310.
- [53] Hanbing Wang, Xiaorui Liu, Wenqi Fan, Xiangyu Zhao, Venkataramana Kini, Devendra Yadav, Fei Wang, Zhen Wen, Jiliang Tang, and Hui Liu. 2024. Rethinking large language model architectures for sequential recommendations. *arXiv preprint arXiv:2402.09543* (2024).
- [54] Yuhao Wang, Ha Tsz Lam, Yi Wong, Ziru Liu, Xiangyu Zhao, Yichao Wang, Bo Chen, Huifeng Guo, and Ruiming Tang. 2023. Multi-task deep recommender systems: A survey. *arXiv preprint arXiv:2302.03525* (2023).
- [55] Yuhao Wang, Ziru Liu, Yichao Wang, Xiangyu Zhao, Bo Chen, Huifeng Guo, and Ruiming Tang. 2024. Diff-MSR: A diffusion model enhanced paradigm for cold-start multi-scenario recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*. 779–787.
- [56] Yuhao Wang, Yichao Wang, Zichuan Fu, Xiangyang Li, Wanyu Wang, Yuyang Ye, Xiangyu Zhao, Huifeng Guo, and Ruiming Tang. 2024. Llm4msr: An llm-enhanced paradigm for multi-scenario recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. 2472–2481.
- [57] Yuhao Wang, Xiangyu Zhao, Bo Chen, Qidong Liu, Huifeng Guo, Huanshuo Liu, Yichao Wang, Rui Zhang, and Ruiming Tang. 2023. PLATE: A prompt-enhanced paradigm for multi-scenario recommendations. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1498–1507.
- [58] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In *Proceedings of the 27th ACM international conference on multimedia*. 1437–1445.
- [59] Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020. Mind: A large-scale dataset for news recommendation. In *Proceedings of the 58th annual meeting of the association for computational linguistics*. 3597–3606.
- [60] Derong Xu, Ziheng Zhang, Zhenxi Lin, Xian Wu, Zhihong Zhu, Tong Xu, Xiangyu Zhao, Yefeng Zheng, and Enhong Chen. 2024. Multi-perspective Improvement of Knowledge Graph Completion with Large Language Models. In *LREC/COLING*.
- [61] Xihong Yang, Heming Jing, Zixing Zhang, Jindong Wang, Huakang Niu, Shuaiqiang Wang, Yu Lu, Junfeng Wang, Dawei Yin, Xinwang Liu, et al. 2024. DaRec: A Disentangled Alignment Framework for Large Language Model and Recommender System. *arXiv preprint arXiv:2408.08231* (2024).
- [62] Zheng Yuan, Fajie Yuan, Yu Song, Youhua Li, Junchen Fu, Fei Yang, Yunzhu Pan, and Yongxin Ni. 2023. Where to go next for recommender systems? id- vs. modality-based recommender models revisited. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2639–2649.
- [63] Chi Zhang, Yantong Du, Xiangyu Zhao, Qilong Han, Rui Chen, and Li Li. 2022. Hierarchical item inconsistency signal learning for sequence denoising in sequential recommendation. In *Proceedings of the 31st ACM international conference on information & knowledge management*. 2508–2518.
- [64] Chi Zhang, Qilong Han, Rui Chen, Xiangyu Zhao, Peng Tang, and Hongtao Song. 2024. Ssdrec: self-augmented sequence denoising for sequential recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 803–815.
- [65] Taolin Zhang, Junwei Pan, Jinpeng Wang, Yaohua Zha, Bin Chen, Shengshui Luo, Yuan Wang, Ming Yue, Jie Jiang, and Shu-Tao Xia. 2024. Towards Scalable Semantic Representation for Recommendation. *arXiv preprint arXiv:2410.09560* (2024).
- [66] Yang Zhang, Fuli Feng, Jizhi Zhang, Keqin Bao, Qifan Wang, and Xiangnan He. 2023. Collm: Integrating collaborative embeddings into large language models for recommendation. *arXiv preprint arXiv:2310.19488* (2023).
- [67] Kesen Zhao, Xiangyu Zhao, Zijian Zhang, and Muyang Li. 2022. Mae4rec: Storage-saving transformer for sequential recommendations. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 2681–2690.
- [68] Xiangyu Zhao, Long Xia, Liang Zhang, Zhuoye Ding, Dawei Yin, and Jiliang Tang. 2018. Deep reinforcement learning for page-wise recommendations. In *Proceedings of the 12th ACM conference on recommender systems*. 95–103.
- [69] Xiangyu Zhao, Liang Zhang, Zhuoye Ding, Long Xia, Jiliang Tang, and Dawei Yin. 2018. Recommendations with negative feedback via pairwise deep reinforcement learning. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1040–1048.
- [70] Bowen Zheng, Yupeng Hou, Hongyu Lu, Yu Chen, Wayne Xin Zhao, Ming Chen, and Ji-Rong Wen. 2024. Adapting large language models by integrating collaborative semantics for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 1435–1448.
- [71] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *AAAI*. 5941–5948.
- [72] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *SIGKDD*. 1059–1068.
- [73] Haolin Zhou, Junwei Pan, Xinyi Zhou, Xihua Chen, Jie Jiang, Xiaofeng Gao, and Guihai Chen. 2024. Temporal Interest Network for User Response Prediction. In *Companion Proceedings of the ACM on Web Conference 2024*. 413–422.