

From IDs to Semantics: A Generative Framework for Cross-Domain Recommendation with Adaptive Semantic Tokenization

Peiyu Hu^{*1,2}, Wayne Lu^{*1,2}, Jia Wang^{1,2†}

¹Xi’an Jiaotong–Liverpool University, Suzhou, China

²University of Liverpool, Liverpool, United Kingdom
{peiyuhu30, lu.wayne0603}@gmail.com, Jia.Wang02@xjtlu.edu.cn

Abstract

Cross-domain recommendation (CDR) is crucial for improving recommendation accuracy and generalization, yet traditional methods are often hindered by the reliance on shared user/item IDs, which are unavailable in most real-world scenarios. Consequently, many efforts have focused on learning disentangled representations through multi-domain joint training to bridge the domain gaps. Recent Large Language Model (LLM)-based approaches show promise, they still face critical challenges, including: (1) the **item ID tokenization dilemma**, which leads to vocabulary explosion and fails to capture high-order collaborative knowledge; and (2) **insufficient domain-specific modeling** for the complex evolution of user interests and item semantics. To address these limitations, we propose **GenCDR**, a novel **Generative Cross-Domain Recommendation** framework. GenCDR first employs a **Domain-adaptive Tokenization** module, which generates disentangled semantic IDs for items by dynamically routing between a universal encoder and domain-specific adapters. Symmetrically, a **Cross-domain Autoregressive Recommendation** module models user preferences by fusing universal and domain-specific interests. Finally, a **Domain-aware Prefix-tree** enables efficient and accurate generation. Extensive experiments on multiple real-world datasets demonstrate that GenCDR significantly outperforms state-of-the-art baselines.

Code — <https://github.com/hupeiyu21/GenCDR>

Introduction

Recommender systems have become indispensable tools for navigating the vast amount of information in modern online services, including e-commerce, social media, and content streaming (Chen et al. 2024a; Liu et al. 2024a; Li et al. 2025d). In real-world scenarios, users often interact across multiple, heterogeneous domains, creating rich behavioral data. Effectively leveraging these interactions for Cross-Domain Recommendation (CDR) has thus emerged as a critical challenge for improving recommendation accuracy and generalization (Li et al. 2023; Xiang et al. 2025). However, the majority of existing CDR methods heavily rely

^{*}These authors contributed equally.

[†]Corresponding author.

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

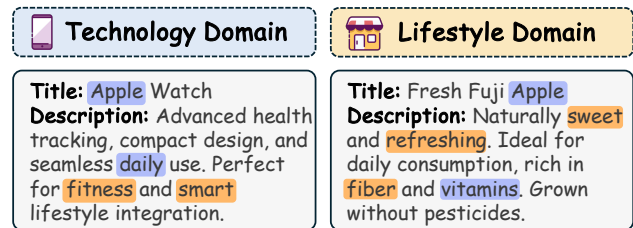


Figure 1: An “Apple” across Lifestyle vs. Technology domains. Blue: shared semantics (e.g., daily use); Orange: domain-specific attributes (e.g., sweet, vitamins for fresh apple; fitness, smart for Apple Watch).

on shared user or item identifiers (IDs) as the bridge for knowledge transfer (Cao et al. 2022; Ma et al. 2024; Guo et al. 2025). This assumption often does not hold in practice, as many cross-domain scenarios, such as online content platforms and offline services, lack a strict alignment of user or item IDs, creating a significant bottleneck for traditional ID-based approaches (Liu et al. 2025c).

The recent advancements in Large Language Models (LLMs) have provided a promising new paradigm for CDR, owing to their powerful capabilities in semantic understanding and sequence generation (Wu et al. 2024; Liu et al. 2025c). Current research in this area can be broadly categorized into two main streams. The first stream utilizes LLMs as powerful feature enhancers, leveraging their world knowledge to generate richer representations for users and items, which are then fed into traditional recommendation models (Liu et al. 2025b; Zhang et al. 2025b). The second stream treats the LLM as the core recommender, reformulating the CDR task as a natural language problem solved via prompting or parameter-efficient fine-tuning (Liu et al. 2025c).

Despite this progress, current LLM-based CDR frameworks still face two fundamental challenges. The first is an **Item Tokenization Gap**, as traditional item indexing methods are ill-suited for LLMs in multi-domain scenarios (Wu et al. 2024). The second is a **Domain Personalization Gap**, where existing methods struggle to effectively disentangle and model the dynamic interplay between universal interests and their domain-specific expressions (Liu et al. 2025c; Zhang et al. 2023). For instance, as illustrated in Figure 1,

an ‘‘Apple Watch’’ in the technology domain and a fresh ‘‘Apple’’ in the lifestyle domain share a common semantic concept (Apple), but also possess highly distinct, domain-specific attributes (e.g., ‘health’, ‘fitness’ vs. ‘sweet’, ‘vitamins’).

To address these challenges, we propose **GenCDR**, a novel **Generative Cross-Domain Recommendation framework based on Large Language Models**. Our work is motivated by a key insight: raw semantic information (e.g., text descriptions) is inherently transferable across domains, whereas item IDs are not. Inspired by the success of generative models in single-domain recommendation, we introduce the concept of discrete *semantic IDs (SIDs)* to GenCDR, directly tackling the item tokenization dilemma (Challenge a). Furthermore, to address the lack of domain personalization (Challenge b), GenCDR features two core modules: a **Domain-adaptive Tokenization** module and a symmetric **Cross-Domain Autoregressive Recommendation** module. These are designed to systematically disentangle and dynamically fuse universal and domain-specific knowledge at both the item and user levels, respectively.

The main contributions of this paper are summarized as follows:

- We propose a novel generative cross-domain recommendation framework, GenCDR. To the best of our knowledge, this is the first work to introduce the generative semantic ID paradigm into LLM-based cross-domain recommendation, effectively resolving the long-standing item tokenization dilemma.
- We systematically design a **Domain-adaptive Tokenization** module that dynamically disentangles and precisely models the universal and item-wise domain-specific knowledge at the semantic level.
- We design a symmetric and collaborative **Cross-Domain Autoregressive Recommendation** module that effectively disentangles and fuses universal and user-wise domain-specific interests during the recommendation process.
- We propose a **Domain-aware Prefix-tree** based decoding strategy to ensure efficient and accurate generation in cross-domain scenarios.
- Extensive experiments on multiple real-world cross-domain datasets demonstrate that GenCDR significantly outperforms existing state-of-the-art methods in terms of both accuracy and generalization.

Related Work

Cross-Domain Sequential Recommendation This task seeks to model a user’s evolving interests across multiple domains by transferring knowledge from their diverse interaction sequences (Chen et al. 2024b). Mainstream approaches often rely on collaborative item IDs, using architectures like gating mechanisms, attention modules (Kang and McAuley 2018; Lu and Yin 2025; Cui et al. 2025), or Graph Neural Networks (GNNs) (Liu et al. 2024b; Li and Lu 2024; Cao et al. 2022) to fuse and transfer knowledge, frequently enhancing the representations with contrastive learning objectives (Ma et al. 2024). Recognizing the limitation of purely

ID-based signals, a more recent trend has started to incorporate richer semantic information by leveraging features from pre-trained language models (Liu et al. 2025c; Li et al. 2022). However, how to effectively integrate these semantics into a unified generative framework, while explicitly disentangling shared and domain-specific knowledge, remains a significant and open challenge.

Generative Recommendation. The paradigm of generative recommendation recasts the task from ranking to an autoregressive sequence generation problem, where Transformer-based models predict sequences of semantic item IDs (Petrov and Macdonald 2023; Hou et al. 2025). The construction of these IDs is a critical research area, with key approaches including content-based tokenization via vector quantization (Li et al. 2025a) (e.g., RQ-VAE (Rajput et al. 2023)), structure-aware methods using hierarchical clustering (Si et al. 2024), and embedding collaborative signals directly into the tokenization process (Mo et al. 2024). However, these techniques have been developed almost exclusively for single-domain datasets (Zheng et al. 2025), leaving their application to complex, multi-domain environments as an open research question.

Large Language Models for Recommendation Large Language Models (LLMs) are integrated into recommender systems in two main ways: either as auxiliary components that enhance traditional models by providing rich semantic features or data augmentation (Sun et al. 2024; Yin et al. 2025; Zhang et al. 2025a), or as core generative engines that reformulate recommendation as a task of autoregressively predicting item IDs (Rajput et al. 2023; Zheng et al. 2024; Lin et al. 2024). Fine-tuning on recommendation datasets, often with parameter-efficient techniques (PEFT) like LoRA (Hu et al. 2022), is a crucial step to align these models for recommendation tasks (Bao et al. 2023; Liu et al. 2025a; Zhang et al. 2023). However, existing work has predominantly focused on single-domain applications, leaving the challenge of effective knowledge transfer and representation across heterogeneous domains largely unaddressed.

Problem Formulation

Let \mathcal{U} be the set of users, \mathcal{D} be the set of domains, and \mathcal{I}_d be the item set for each domain $d \in \mathcal{D}$. The total item set is $\mathcal{I} = \bigcup_{d \in \mathcal{D}} \mathcal{I}_d$. For each user $u \in \mathcal{U}$, their historical interactions in a domain d are represented as a chronological sequence $S_d^u = (i_1, \dots, i_t)$, where $i_k \in \mathcal{I}_d$. The user’s complete cross-domain historical profile, \mathcal{H}^u , is the collection of all their single-domain sequences: $\mathcal{H}^u = \{S_d^u \mid d \in \mathcal{D}_u\}$, where $\mathcal{D}_u \subseteq \mathcal{D}$ is the set of domains user u has interacted with.

The task of Cross-Domain Sequential Recommendation (CDSR) is to predict the next item i_{target}^u that a user u is most likely to interact with in a target domain $d_t \in \mathcal{D}$, given their entire historical profile \mathcal{H}^u . The objective is to learn a generative model parameterized by θ that maximizes the log-likelihood of the held-out target items:

$$\mathcal{L} = \sum_{u \in \mathcal{U}} \log P(i_{\text{target}}^u | \mathcal{H}^u; \theta). \quad (1)$$

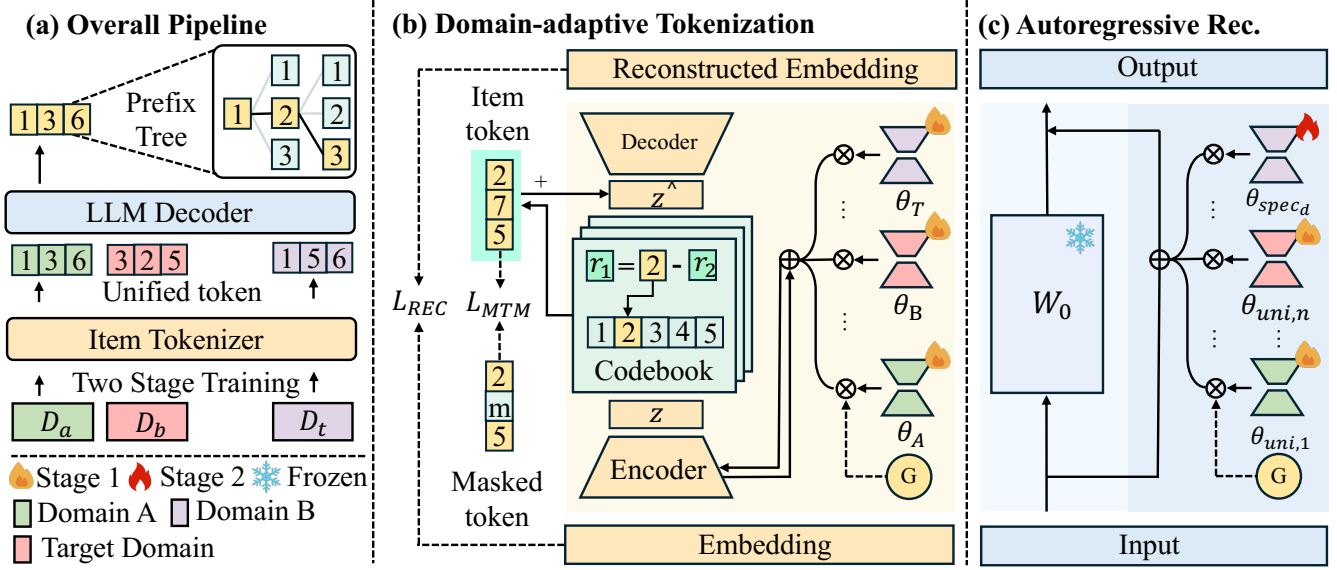


Figure 2: The architecture of our GenCDR framework. (a) The two-stage pipeline comprising the tokenization and recommendation modules. (b) The detailed structure of the Domain-adaptive Tokenization module, featuring a hierarchical adapter system with dynamic routing. (c) The symmetric architecture of the Cross-Domain Autoregressive Recommendation module.

Methodology

To address the key challenges of item tokenization and domain personalization in cross-domain recommendation, we propose **GenCDR**, a novel generative framework. As illustrated in Figure 2, our framework is composed of three core modules: a **Domain-adaptive Tokenization** module, a **Cross-Domain Autoregressive Recommendation** module, and a **Domain-aware Prefix-tree** for efficient inference.

Domain-adaptive Tokenization

To empower large language models (LLMs) with the ability to process items from diverse domains, we introduce a *Domain-adaptive Tokenization* module. This module generates unified SIDs for items, balancing domain-agnostic universal semantics with domain-specific discriminative features to yield expressive representations for generative recommendation tasks. The SIDs are designed to exhibit three critical properties: (i) *Semantic Richness*: capturing comprehensive item semantics; and (ii) *Semantic Similarity*: ensuring similar items across domains share comparable IDs.

Domain-Universal Semantic Token Generation To establish a unified semantic foundation for knowledge transfer, we propose a Universal Discrete Semantic Encoder based on a Residual-Quantized Variational Autoencoder (RQ-VAE) framework (Lee et al. 2022). The RQ-VAE, consisting of an encoder E , a decoder D , and M codebooks, is pre-trained on the textual features of all items. It converts an item’s feature embedding \mathbf{x} into a sequence of discrete codes $\mathbf{c} = (c_0, \dots, c_{M-1})$. The model encodes \mathbf{x} to a latent representation $\mathbf{z} = E(\mathbf{x})$, with initial residual $\mathbf{r}_0 = \mathbf{z}$. For each level $d = 0$ to $M - 1$, \mathbf{r}_d is quantized to the nearest codebook vector \mathbf{e}_{c_d} from codebook \mathcal{C}_d , where $c_d = \arg \min_k \|\mathbf{r}_d - \mathbf{e}_k\|^2$,

and the next residual is $\mathbf{r}_{d+1} = \mathbf{r}_d - \mathbf{e}_{c_d}$. The quantized $\hat{\mathbf{z}} := \sum_{d=0}^{M-1} \mathbf{e}_{c_d}$ is decoded to $\hat{\mathbf{x}} = D(\hat{\mathbf{z}})$.

The model is optimized via a joint objective function. This objective includes a standard reconstruction loss $\mathcal{L}_{\text{REC}} = \|\mathbf{x} - \hat{\mathbf{x}}\|^2$ and a quantization loss \mathcal{L}_Q that aligns the encoder’s output with the codebook vectors using commitment terms (Van Den Oord, Vinyals et al. 2017):

$$\mathcal{L}_Q = \sum_{d=0}^{M-1} \|\text{sg}(\mathbf{r}_d) - \mathbf{e}_{c_d}\|^2 + \beta \|\mathbf{r}_d - \text{sg}(\mathbf{e}_{c_d})\|^2. \quad (2)$$

To further ensure the codes are contextually coherent, we add a Masked Token Modeling (MTM) loss, which trains the model to predict masked codes c_i from their surrounding context S_{masked} :

$$\mathcal{L}_{\text{MTM}} = -\mathbb{E}_{\mathbf{x} \sim \mathcal{X}, I_{\text{mask}}} \left[\sum_{i \in I_{\text{mask}}} \log P(c_i | S_{\text{masked}}; \theta_{\text{ctx}}) \right]. \quad (3)$$

The total pre-training loss, $\mathcal{L}_{\text{pretrain}} = \mathcal{L}_{\text{REC}} + \mu \mathcal{L}_Q + \lambda \mathcal{L}_{\text{MTM}}$, guides the model to learn universal semantic tokens that are both representative and contextually aware. Upon completion, the universal encoder and codebooks are frozen.

Domain-specific Semantic Token Adapters While the universal encoder establishes a domain-agnostic semantic foundation, it may not fully capture domain-specific discriminative features, such as visual aesthetics in videos or narrative styles in books. To address this, we introduce domain-specific semantic adapters that refine universal representations in a parameter-efficient manner, enhancing their relevance for each domain.

We leverage Low-Rank Adaptation (LoRA) (Hu et al. 2022) to achieve this. For each domain $d \in \mathcal{D}$, a lightweight

LoRA module is introduced, comprising low-rank matrices $B_d \in \mathbb{R}^{d_{\text{out}} \times r}$ and $A_d \in \mathbb{R}^{r \times d_{\text{in}}}$, where $r \ll \min(d_{\text{in}}, d_{\text{out}})$. These matrices augment the frozen weights $W_0 \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ of the universal encoder E , modifying the forward pass as:

$$h_{\text{out}} = W_0 h_{\text{in}} + B_d A_d h_{\text{in}}. \quad (4)$$

Denoting the adapted encoder as E_{θ_d} with trainable parameters $\theta_d = \{B_d, A_d\}$ for domain d , we fine-tune θ_d in a second training phase. For each item embedding \mathbf{x} from domain d , we minimize a self-supervised reconstruction loss:

$$\mathcal{L}_{\text{adapter}} = \mathbb{E}_{\mathbf{x} \sim \mathcal{X}_d} [\|\mathbf{x} - D(Q(E_{\theta_d}(\mathbf{x})))\|_2^2], \quad (5)$$

where Q and decoder D remain frozen. This approach ensures domain-specific refinements with minimal additional parameters, enabling efficient adaptation to diverse domains.

Item-level Dynamic Semantic Routing Network To effectively integrate universal and domain-specific representations, we propose an Item-level Dynamic Semantic Routing Network that adaptively balances these representations on a per-item basis. This approach mitigates the risk of negative transfer inherent in static fusion strategies by dynamically determining the contribution of general cross-domain and domain-specific semantics for each item.

The routing network, denoted R_ϕ with parameters ϕ , is a lightweight neural network (e.g., a multi-layer perceptron) that takes an item’s embedding \mathbf{x} as input and produces a gating weight $\alpha \in [0, 1]$. For an item from domain d , we compute two latent representations prior to quantization: the universal representation $\mathbf{z}_{\text{uni}} = E(\mathbf{x})$ from the frozen universal encoder, and the domain-specific representation $\mathbf{z}_{\text{spec}} = E_{\theta_d}(\mathbf{x})$ from the adapted encoder. The router calculates:

$$\alpha = \sigma(R_\phi(\mathbf{x})), \quad (6)$$

$$\mathbf{z}_{\text{fused}} = (1 - \alpha) \cdot \mathbf{z}_{\text{uni}} + \alpha \cdot \mathbf{z}_{\text{spec}}, \quad (7)$$

where $\sigma(\cdot)$ is the sigmoid function. The fused representation $\mathbf{z}_{\text{fused}}$ is then quantized and decoded.

To promote disentangled representations and prevent overfitting, we regularize the router using the Variational Information Bottleneck (VIB) principle (Alemi et al. 2016). The VIB loss minimizes the information the router extracts from \mathbf{x} , ensuring only essential features influence the routing decision. This is enforced via a KL-divergence term:

$$\mathcal{L}_{\text{VIB}} = D_{\text{KL}}(q(\mathbf{z}_r | \mathbf{x}) \| p(\mathbf{z}_r)), \quad (8)$$

where $q(\mathbf{z}_r | \mathbf{x})$ is the router’s internal representation distribution, and $p(\mathbf{z}_r)$ is a prior (e.g., standard normal). This loss is incorporated into the second-phase training objective, enabling a balanced fusion of shared and domain-specific knowledge.

Cross-Domain Autoregressive Recommendation

Leveraging the unified SIDs produced by the Domain-adaptive Tokenization module, this component models intricate temporal patterns in user interaction sequences to enable personalized cross-domain recommendations. We introduce a parameter-efficient, two-phase fine-tuning strategy that acknowledges the multifaceted nature of user interests (e.g., brand preferences or category affinities). In

the initial phase, a mixture of diverse LoRA adapters is trained on aggregated data from all domains to capture transferable, domain-agnostic interest patterns. The subsequent phase fine-tunes domain-specific LoRA adapters, with a dynamic routing network facilitating adaptive fusion of universal and specialized knowledge during inference.

Universal Interest Modeling Network To model the diverse facets of user interests across domains, we develop a Universal Interest Modeling Network. This is achieved by enhancing a pre-trained large language model (LLM) with a mixture of multiple Low-Rank Adaptation (LoRA) adapters (Li et al. 2025c,b). This collection of adapters is trained jointly to capture distinct, transferable behavioral patterns. The parameters of the i -th universal expert are denoted as $\theta_{\text{uni},i}$. The complete set of these parameters, $\Theta_{\text{uni}} = \{\theta_{\text{uni},1}, \dots, \theta_{\text{uni},N}\}$, represents all trainable weights for the universal module G .

The input to this network consists of sequences of cross-domain SIDs, $S^u = (c_1^u, c_2^u, \dots, c_t^u)$. In the initial fine-tuning phase, we optimize the universal parameters Θ_{uni} using a standard autoregressive objective, predicting the next semantic ID given the preceding sequence. The training loss is defined as:

$$\mathcal{L}_{\text{uni}} = - \sum_{u \in \mathcal{U}} \sum_{k=1}^{|S^u|-1} \log P(c_{k+1}^u | c_{\leq k}^u; \theta_{\text{LLM}}, \Theta_{\text{uni}}), \quad (9)$$

where θ_{LLM} denotes the frozen LLM parameters. After this phase, both θ_{LLM} and the universal parameter set Θ_{uni} are fixed, forming the Universal Interest Modeling Network that serves as the basis for domain-specific adaptation.

Domain-specific Interest Adaptation While the Universal Interest Modeling Network captures general user preferences, domain-specific nuances require tailored modeling. To address this, we introduce a second fine-tuning phase to train domain-specific LoRA adapters, enabling the model to adapt to the unique characteristics of each domain.

For each domain $d \in \mathcal{D}$, we augment the frozen model with a dedicated, trainable LoRA adapter, denoted θ_{spec_d} . During this phase, both the base LLM parameters θ_{LLM} and the universal parameters Θ_{uni} remain fixed. The training focuses solely on domain-specific interaction sequences S_d^u from users $u \in \mathcal{U}_d$. We optimize θ_{spec_d} by minimizing the autoregressive loss:

$$\mathcal{L}_{\text{spec}_d} = - \sum_{u \in \mathcal{U}_d} \sum_{k=1}^{|S_d^u|-1} \log P(c_{k+1}^u | c_{\leq k}^u; \theta_{\text{LLM}}, \Theta_{\text{uni}}, \theta_{\text{spec}_d}). \quad (10)$$

This approach enables the model to efficiently learn domain-specific interest patterns, setting the stage for dynamic integration during inference.

User-level Dynamic Interest Routing Network Symmetrically to the item-level router, we employ a VIB-regularized User-level Dynamic Interest Routing Network to prevent negative transfer during inference. This lightweight gate takes the user’s history representation \mathbf{h}_t as input to compute a dynamic weight $\gamma \in [0, 1]$. This weight fuses the

probability distributions from the universal model (P_{uni}) and the domain-adapted model (P_{spec}) as follows:

$$P_{\text{final}}(i | S^u) = (1-\gamma) \cdot P_{\text{uni}}(i | S^u) + \gamma \cdot P_{\text{spec}}(i | S^u). \quad (11)$$

Here, P_{uni} is the output distribution from the frozen universal network (parameterized by Θ_{uni}), while P_{spec} is from the network augmented with domain-specific adapters (parameterized by Θ_{uni} and θ_{spec_d}). The VIB regularization on the router ensures the fusion logic is efficient and robust.

Inference – Domain-aware Prefix-tree To ensure efficient and valid semantic ID generation, we propose a Domain-aware Prefix-tree mechanism that mitigates the limitations of standard autoregressive decoding, such as computational inefficiency and invalid ID outputs. For each domain $d \in \mathcal{D}$, we construct an offline prefix tree T_d encoding all valid semantic ID sequences produced by the Domain-adaptive Tokenization module. During inference, given a target domain d_t , the corresponding tree T_{d_t} guides the generation process. At each decoding step k , the tree identifies a valid subset of next codes $V_{\text{valid}}(s_{k-1}) \subset C_k$ based on the current prefix s_{k-1} . The LLM’s predictions are constrained to this subset using a masked softmax:

$$P(c_k | s_{k-1}, T_{d_t}) = \frac{\exp(z_k)}{\sum_{c' \in V_{\text{valid}}(s_{k-1})} \exp(z_{c'})}, \quad (12)$$

where z_k are the LLM’s logits. This approach ensures valid sequence generation while significantly reducing computational overhead, enhancing the efficiency of the recommendation process.

Experiments

In this section, we conduct extensive experiments on several public datasets to evaluate our proposed model. The experiments are designed to answer the following key Research Questions (RQs):

- **(Effectiveness) RQ1:** How does our proposed model perform against state-of-the-art single-domain and cross-domain recommendation baselines?
- **(Ablation) RQ2:** What is the contribution of each key component in our framework?
- **(Analysis) RQ3:** Can our framework learn visually separable representations for universal and domain-specific knowledge?
- **(Sensitivity) RQ4:** How does our model’s performance change with respect to key hyper-parameter settings?
- **(Efficiency) RQ5:** How efficient is our GenCDR framework in terms of training cost and inference scalability?

Experimental Setup

Datasets. We experiment on three cross-domain dataset pairs, each reflecting a distinct real-world scenario: **Sports-Clothing** (Leisure), **Phones-Electronics** (Technology), and **Books-Movies** (Entertainment). The first two pairs are derived from the public Amazon product review dataset (McAuley et al. 2015), while the third is collected from Douban (Zhu et al. 2020).

Dataset	#Users	#Items	#Interactions	Sparsity	Overlap
Sports	35,598	18,357	296,337	99.95%	1.73%
Clothing	39,387	23,033	278,677	99.97%	(704)
Phones	27,879	10,429	194,439	99.93%	0.55%
Electronics	192,403	63,001	1,689,188	99.99%	(404)
Books	1,713	8,601	104,295	99.29%	7.48%
Movies	2,628	20,964	1,249,016	97.73%	(2,058)

Table 1: Statistics of the datasets used in our experiments. Item overlap and sparsity are computed after merging.

Following (Rajput et al. 2023; Deldjoo et al. 2024), we treat users’ historical reviews as interactions arranged chronologically. We use the leave-last-out evaluation protocol (Kang and McAuley 2018; Zhao et al. 2022), where the last item is for testing, and the second-to-last for validation. Table 1 shows dataset statistics.

Baselines. To verify the effectiveness of our approach, we compare it with a comprehensive set of state-of-the-art baselines. These fall into three main categories: (1) **Single-domain Sequential Recommendation (SDSR)**, including BERT4Rec (Sun et al. 2019), SASRec (Kang and McAuley 2018) and STOSA (Fan et al. 2022); (2) **Generative Recommendation Systems (GRS)**, featuring VQ-Rec (Hou et al. 2023), TIGER (Rajput et al. 2023), and HSTU (Zhai et al. 2024); and (3) **Cross-domain Sequential Recommendation (CDSR)**, with models like C2DSR (Cao et al. 2022), TriCDR (Ma et al. 2024), and LLM4CDSR (Liu et al. 2025c). We follow each baseline’s original experimental setup: single-domain models are evaluated within their domains, while multi-domain methods are trained across domains to reflect their intended design.

Evaluation Metrics. Following the standard practice in sequential recommendation literature (Kang and McAuley 2018; Rajput et al. 2023), we adopt Recall@ K and NDCG@ K as our evaluation metrics, with K set to 5 and 10. For each model, the checkpoint that achieves the best Recall@10 on the validation set is selected for the final testing phase.

Implementation Details. Our framework is implemented in PyTorch with Hugging Face PEFT for LoRA-based fine-tuning. We use Qwen2.5-7B as the backbone model, optimized with AdamW ($\text{lr}=5 \times 10^{-5}$, $\text{batch}=8$). All models are trained for 10 epochs on NVIDIA H200 GPUs with mixed precision (FP16).

Overall Performance (RQ1)

The overall performance comparison of our proposed GenCDR against all baseline models is summarized in Table 2. The results show that our proposed GenCDR consistently and significantly outperforms all baseline models, demonstrating its overall superiority in the cross-domain sequential recommendation task.

Analyzing the baselines offers insight into this improvement. Cross-domain (CDSR) models generally outperform single-domain (SDSR) models, confirming the benefit of cross-domain signals. Generative (GenRec) models also sur-

Dataset		Metric	SDSR			GenRec			CDSR			Ours
Scene	Domain		Bert4Rec	SASRec	STOSA	VQ-Rec	TIGER	HSTU	C2DSR	TriCDR	LLM4CDSR	GenCDR
Leisure	Sports	R@5	0.0188	0.0197	0.0236	0.0261	<u>0.0267</u>	0.0254	0.0265	0.0266	0.0263	0.0274
		N@5	0.0121	0.0126	0.0162	0.0238	<u>0.0244</u>	0.0241	0.0253	0.0255	<u>0.0257</u>	0.0261
		R@10	0.0325	0.0334	0.0346	0.0389	<u>0.0397</u>	0.0381	0.0395	0.0396	0.0398	0.0403
		N@10	0.0169	0.0173	<u>0.0283</u>	0.0281	0.0287	0.0277	0.0258	0.0259	0.0260	0.0262
	Clothing	R@5	0.0128	0.0132	0.0162	0.0171	0.0173	0.0175	0.0172	0.0174	<u>0.0176</u>	0.0181
		N@5	0.0078	0.0081	0.0119	0.0129	0.0125	0.0132	0.0158	0.0161	<u>0.0163</u>	0.0167
		R@10	0.0219	0.0227	0.0223	0.0248	0.0241	0.0253	0.0255	0.0258	<u>0.0261</u>	0.0265
		N@10	0.0105	0.0108	0.0135	0.0170	0.0167	0.0174	0.0191	0.0194	<u>0.0196</u>	0.0203
Technology	Phones	R@5	0.0331	0.0345	0.0415	0.0411	0.0423	0.0415	0.0428	0.0434	<u>0.0431</u>	0.0436
		N@5	0.0215	0.0224	0.0283	0.0308	0.0315	0.0327	0.0392	0.0396	<u>0.0401</u>	0.0411
		R@10	0.0524	0.0537	0.0618	0.0607	0.0613	0.0615	0.0589	0.0593	0.0614	0.0621
		N@10	0.0278	0.0287	0.0346	0.0399	0.0406	0.0425	0.0493	0.0505	<u>0.0506</u>	0.0512
	Electronics	R@5	0.0179	0.0186	0.0213	0.0219	0.0228	0.0232	0.0235	<u>0.0238</u>	0.0237	0.0241
		N@5	0.0118	0.0122	0.0148	0.0211	0.0214	0.0226	0.0229	<u>0.0231</u>	0.0230	0.0235
		R@10	0.0276	0.0285	0.0315	0.0318	0.0322	0.0328	0.0336	<u>0.0339</u>	0.0338	0.0342
		N@10	0.0149	0.0154	0.0172	0.0262	0.0269	0.0271	0.0278	<u>0.0280</u>	0.0279	0.0283
Entertainment	Books	R@5	0.0089	0.0093	0.0142	0.0175	0.0172	<u>0.0181</u>	0.0152	0.0155	0.0161	0.0192
		N@5	0.0071	0.0076	0.0117	0.0178	0.0177	<u>0.0180</u>	0.0143	0.0148	0.0153	0.0187
		R@10	0.0176	0.0182	0.0219	0.0224	0.0221	<u>0.0230</u>	0.0205	0.0211	0.0216	0.0237
		N@10	0.0158	0.0164	0.0165	0.0201	0.0198	<u>0.0206</u>	0.0182	0.0185	0.0189	0.0212
	Movies	R@5	0.1503	0.1542	0.1562	0.1680	0.1652	<u>0.1682</u>	0.1588	0.1601	0.1613	0.1713
		N@5	0.1015	0.1047	0.1063	0.1182	0.1156	<u>0.1189</u>	0.1092	0.1105	0.1149	0.1215
		R@10	0.1798	0.1825	0.1753	0.1922	0.1893	<u>0.1931</u>	0.1854	0.1865	0.1878	0.1971
		N@10	0.1211	0.1265	0.1223	0.1261	0.1255	<u>0.1268</u>	0.1203	0.1217	0.1225	0.1275

Table 2: Overall performance comparison on all datasets. R@K and N@K denote Recall and NDCG at cutoff K. Best results are in **bold**, and the best baseline results are underlined. The t -tests showed significant performance improvements ($p \leq 0.05$).

pass SDSR baselines but still lag behind specialized CDSR methods. This performance gap highlights that directly applying existing generative models to cross-domain scenarios is insufficient. GenCDR is designed to close this gap by tightly coupling the generative paradigm with cross-domain knowledge transfer, achieving state-of-the-art results.

Ablation Study (RQ2)

To dissect the contribution of each key design choice in our GenCDR framework, we conducted a thorough ablation study. The results, summarized in Table 3, unequivocally demonstrate that each component plays an integral role.

- **Impact of Contextual Code Modeling.** Removing the MTM loss (**w/o MTM**) degrades performance, confirming that learning the contextual "grammar" of the semantic codes is crucial, beyond simple reconstruction.
- **Impact of Item-specific Adaptation.** Removing the item-specific adapter (**w/o Item Adapter**) degrades performance, validating the need for domain-specific item semantics.
- **Impact of the Specific Expert.** Removing the domain-specific expert (**w/o Specific Expert**) significantly hurts performance, proving its crucial role in capturing fine-grained user preferences.
- **Impact of the Universal Experts.** Removing all N universal experts (**w/o Universal Experts**) causes a sharp performance drop, confirming that a shared cross-domain knowledge foundation is indispensable.

- **Impact of the MoE Gate.** Replacing the trainable MoE gate with simple averaging (**w/o MoE Gate**) hurts performance, highlighting the importance of a dynamic, context-aware selection of experts over naive fusion.
- **Impact of Constrained Decoding.** Removing the prefix-tree constraint (**w/o Prefix Tree**) leads to a consistent performance drop, as it guarantees the generation of valid item IDs and prevents "hallucinated" recommendations.

In-depth Analysis (RQ3)

To visualize the learned representations, we plot the final item embeddings ($\mathbf{z}_{\text{fused}}$) using t-SNE in Figure 3. With only universal adapters (b), domain embeddings mix, whereas with domain-specific adapters (c), they form clear clusters, highlighting the benefit of domain-specific adaptation.

Hyper-parameter Analysis (RQ4)

We analyze the sensitivity of key hyperparameters on the **Cloth** dataset in Figure 4. Universal experts (N), LoRA rank (r), and alpha (α) exhibit clear optima (e.g., $N = 4$, $r = 64$), with larger values causing overfitting. A small LoRA dropout (0.05) provides effective regularization. Overall, the results show a stable trade-off between capacity and generalization.

Analysis of Efficiency (RQ5)

Training Efficiency. As shown in Figure 5, our LoRA-based fine-tuning greatly reduces trainable parameters, training time, and GPU memory compared to full fine-tuning.

Category	Variant	Phones	Electronics	Sports	Clothing
Full Model	GenCDR	0.0512	0.0283	0.0262	0.0203
Tokenization	w/o MTM	0.0483 (↓5.7%)	0.0267 (↓5.7%)	0.0245 (↓6.5%)	0.0190 (↓6.4%)
	w/o Adapter	0.0466 (↓9.0%)	0.0255 (↓9.9%)	0.0238 (↓9.2%)	0.0183 (↓9.9%)
Autoregressive Rec.	w/o Specific Expert	0.0448 (↓12.5%)	0.0245 (↓13.4%)	0.0226 (↓13.7%)	0.0173 (↓14.8%)
	w/o Universal Experts	0.0425 (↓17.0%)	0.0232 (↓18.0%)	0.0212 (↓19.1%)	0.0162 (↓20.2%)
	w/o MoE Gate (Avg.)	0.0475 (↓7.2%)	0.0262 (↓7.4%)	0.0242 (↓7.6%)	0.0186 (↓8.4%)
Inference Strategy	w/o Prefix Tree	0.0498 (↓2.7%)	0.0274 (↓3.2%)	0.0255 (↓2.7%)	0.0198 (↓2.5%)

Table 3: Ablation study on GenCDR components across four datasets (NDCG@10). Values in parentheses denote the drop in performance compared to the full model.

Inference Efficiency and Scalability. Furthermore, Figure 6 demonstrates GenCDR’s superior scalability, as its inference cost remains constant regardless of the item pool size due to our prefix-tree constrained generative architecture.

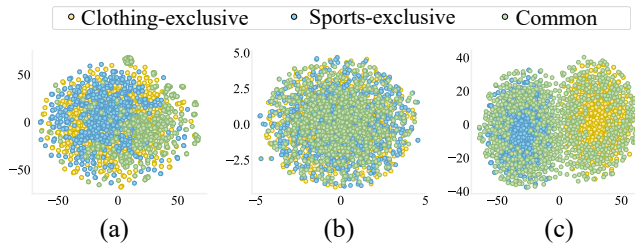


Figure 3: t-SNE visualization of item embeddings in three different settings. (a) Original Item Embeddings, (b) Shared LoRA Embeddings, (c) Domain-Specific LoRA Embeddings.

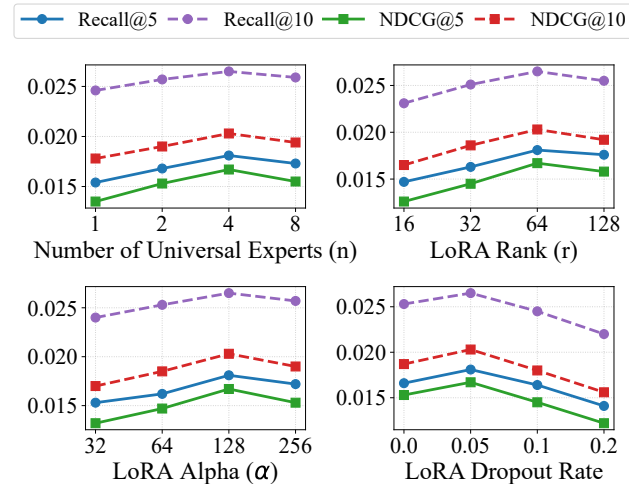


Figure 4: Sensitivity of LoRA fine-tuning to key hyper-parameters on the Cloth dataset.

Conclusion

In this paper, we addressed the critical challenges of item tokenization and domain personalization in LLM-based cross-domain recommendation. We proposed **GenCDR**, a novel generative framework that systematically tackles these issues. Our approach introduces a **Domain-adaptive Tokenization** module to create hybrid SIDs and a symmetric **Cross-Domain Autoregressive Recommendation** module to dynamically model user interests by fusing universal and specific knowledge. Furthermore, a prefix-tree mechanism ensures efficient and valid inference. Experiments show that GenCDR achieves superior effectiveness and scalability, and future work will explore incorporating multimodal features for richer representations.

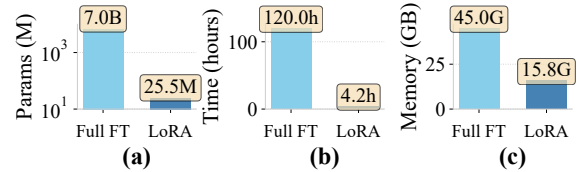


Figure 5: Comparison of training efficiency using the Qwen2.5-7B model. The plots show (a) trainable parameters (log scale), (b) training time, and (c) peak GPU memory for our LoRA-based GenCDR versus a Full Fine-Tuning (Full FT) version.

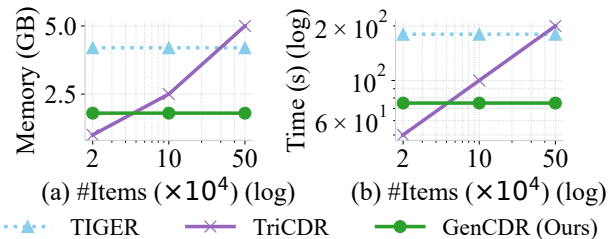


Figure 6: Comparison of runtime memory and inference time w.r.t. the item pool size for TriCDR, TIGER, and GenCDR (Qwen2.5-0.5B).

Acknowledgements

This research was supported in part by the National Natural Science Foundation of China (Grant No. 72401232), the Natural Science Foundation of the Jiangsu Higher Educational Institution of China (Grant No. 23KJB520037), and the XJTLU Research Development Fund under RDF-21-01-053.

References

- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Bao, K.; Zhang, J.; Zhang, Y.; Wang, W.; Feng, F.; and He, X. 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.
- Cao, J.; Cong, X.; Sheng, J.; Liu, T.; and Wang, B. 2022. Contrastive cross-domain sequential recommendation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 138–147.
- Chen, C.; Zhang, Y.; Li, Y.; Wang, J.; Qi, L.; Xu, X.; Zheng, X.; and Yin, J. 2024a. Post-training attribute unlearning in recommender systems. *ACM Transactions on Information Systems*, 43(1): 1–28.
- Chen, S.; Xu, Z.; Pan, W.; Yang, Q.; and Ming, Z. 2024b. A survey on cross-domain sequential recommendation. *arXiv preprint arXiv:2401.04971*.
- Cui, X.; Lu, W.; Tong, Y.; Li, Y.; and Zhao, Z. 2025. Multi-Modal Multi-Behavior Sequential Recommendation with Conditional Diffusion-Based Feature Denoising. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1593–1602.
- Deldjoo, Y.; He, Z.; McAuley, J.; Korikov, A.; Sanner, S.; Ramisa, A.; Vidal, R.; Sathiamoorthy, M.; Kasirzadeh, A.; and Milano, S. 2024. A review of modern recommender systems using generative models (gen-recsys). In *Proceedings of the 30th ACM SIGKDD conference on Knowledge Discovery and Data Mining*, 6448–6458.
- Fan, Z.; Liu, Z.; Wang, Y.; Wang, A.; Nazari, Z.; Zheng, L.; Peng, H.; and Yu, P. S. 2022. Sequential recommendation via stochastic self-attention. In *Proceedings of the ACM web conference 2022*, 2036–2047.
- Guo, J.; Zhang, J.; Wu, F.; Lu, H.; Wang, Q.; Yang, W.; Lim, E. G.; and Lu, D. 2025. HiGarment: Cross-modal Harmony Based Diffusion Model for Flat Sketch to Realistic Garment Image. *arXiv preprint arXiv:2505.23186*.
- Hou, Y.; He, Z.; McAuley, J.; and Zhao, W. X. 2023. Learning vector-quantized item representation for transferable sequential recommenders. In *Proceedings of the ACM Web Conference 2023*, 1162–1171.
- Hou, Y.; Ni, J.; He, Z.; Sachdeva, N.; Kang, W.-C.; Chi, E. H.; McAuley, J.; and Cheng, D. Z. 2025. ActionPiece: Contextually Tokenizing Action Sequences for Generative Recommendation. *arXiv preprint arXiv:2502.13581*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Kang, W.-C.; and McAuley, J. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, 197–206. IEEE.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W.-S. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11523–11532.
- Li, C.; Zhao, M.; Zhang, H.; Yu, C.; Cheng, L.; Shu, G.; Kong, B.; and Niu, D. 2022. RecGURU: Adversarial learning of generalized user representations for cross-domain recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 571–581.
- Li, G.; Zhang, X.; Zhang, Y.; Yin, Y.; Yin, G.; and Lin, W. 2025a. Semantic convergence: Harmonizing recommender systems via two-stage alignment and behavioral semantic tokenization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12040–12048.
- Li, Y.; and Lu, W. 2024. MHHCR: Multi-behavior Heterogeneous Hypergraph Contrastive Recommendation. In *International Conference on Web Information Systems Engineering*, 91–102. Springer.
- Li, Y.; Xiong, H.; Kong, L.; Zhang, R.; Xu, F.; Chen, G.; and Li, M. 2023. Mhrr: Moocs recommender service with meta hierarchical reinforced ranking. *IEEE Transactions on Services Computing*, 16(6): 4467–4480.
- Li, Y.; Zhang, H.; Zhang, Y.; Cai, H.; Cai, M.; Wang, S.; Xiong, H.; Kong, L.; Yin, D.; and Chen, L. 2025b. Rankexpert: A mixture of textual-and-behavioral experts for multi-objective learning-to-rank in web search. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, 4578–4589.
- Li, Y.; Zhang, H.; Zhang, Y.; Ma, X.; Ye, W.; Song, N.; Wang, S.; Xiong, H.; Yin, D.; and Chen, L. 2025c. M2oERank: Multi-Objective Mixture-of-Experts Enhanced Ranking for Satisfaction-Oriented Web Search. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, 4441–4454. IEEE.
- Li, Y.; Zhang, Y.; Liu, W.; Feng, X.; Han, Z.; Chen, C.; and Yan, C. 2025d. Multi-Objective Unlearning in Recommender Systems via Preference Guided Pareto Exploration. *IEEE Transactions on Services Computing*.
- Lin, X.; Wang, W.; Li, Y.; Feng, F.; Ng, S.-K.; and Chua, T.-S. 2024. Bridging items and language: A transition paradigm for large language model-based recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1816–1826.
- Liu, J.; Yan, X.; Li, D.; Zhang, G.; Gu, H.; Zhang, P.; Lu, T.; Shang, L.; and Gu, N. 2025a. Enhancing LLM-Based Recommendations Through Personalized Reasoning. *arXiv e-prints*, arXiv–2502.
- Liu, Q.; Wu, X.; Wang, W.; Wang, Y.; Zhu, Y.; Zhao, X.; Tian, F.; and Zheng, Y. 2025b. Llmemb: Large language model can be a good embedding generator for sequential

- recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12183–12191.
- Liu, Q.; Zhao, X.; Wang, Y.; Zhang, Z.; Zhong, H.; Chen, C.; Li, X.; Huang, W.; and Tian, F. 2025c. Bridge the Domains: Large Language Models Enhanced Cross-domain Sequential Recommendation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1582–1592.
- Liu, R.; Chen, H.; Bei, Y.; Shen, Q.; Zhong, F.; Wang, S.; and Wang, J. 2024a. Fine Tuning Out-of-Vocabulary Item Recommendation with User Sequence Imagination. *Advances in Neural Information Processing Systems*, 37: 8930–8955.
- Liu, Z.; Hua, Y.; Chen, M.; Zhang, Y.; Chen, Z.; Liang, L.; Chen, H.; and Zhang, W. 2024b. UniHR: Hierarchical Representation Learning for Unified Knowledge Graph Link Prediction. *arXiv preprint arXiv:2411.07019*.
- Lu, W.; and Yin, L. 2025. DMMD4SR: Diffusion Model-based Multi-level Multimodal Denoising for Sequential Recommendation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 6363–6372.
- Ma, H.; Xie, R.; Meng, L.; Chen, X.; Zhang, X.; Lin, L.; and Zhou, J. 2024. Triple sequence learning for cross-domain recommendation. *ACM Transactions on Information Systems*, 42(4): 1–29.
- McAuley, J.; Targett, C.; Shi, Q.; and Van Den Hengel, A. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, 43–52.
- Mo, M.; Lu, W.; Xie, Q.; Lv, X.; Xiao, Z.; Yang, H.; and Zhang, Y. 2024. MIN: Multi-stage Interactive Network for Multimodal Recommendation. In *International Conference on Web Information Systems Engineering*, 191–205. Springer.
- Petrov, A. V.; and Macdonald, C. 2023. Generative sequential recommendation with gptrec. *arXiv preprint arXiv:2306.11114*.
- Rajput, S.; Mehta, N.; Singh, A.; Hulikal Keshavan, R.; Vu, T.; Heldt, L.; Hong, L.; Tay, Y.; Tran, V.; Samost, J.; et al. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems*, 36: 10299–10315.
- Si, Z.; Sun, Z.; Chen, J.; Chen, G.; Zang, X.; Zheng, K.; Song, Y.; Zhang, X.; Xu, J.; and Gai, K. 2024. Generative retrieval with semantic tree-structured identifiers and contrastive learning. In *Proceedings of the 2024 Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, 154–163.
- Sun, F.; Liu, J.; Wu, J.; Pei, C.; Lin, X.; Ou, W.; and Jiang, P. 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.
- Sun, Z.; Si, Z.; Zang, X.; Zheng, K.; Song, Y.; Zhang, X.; and Xu, J. 2024. Large language models enhanced collaborative filtering. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2178–2188.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Wu, L.; Zheng, Z.; Qiu, Z.; Wang, H.; Gu, H.; Shen, T.; Qin, C.; Zhu, C.; Zhu, H.; Liu, Q.; et al. 2024. A survey on large language models for recommendation. *World Wide Web*, 27(5): 60.
- Xiang, Y.; Fan, L.; Yin, C.; Kong, M.; and Ji, C. 2025. Harnessing Light for Cold-Start Recommendations via Epistemic Uncertainty. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM)*, 5361–5365.
- Yin, J.; Zeng, Z.; Li, M.; Yan, H.; Li, C.; Han, W.; Zhang, J.; Liu, R.; Sun, H.; Deng, W.; et al. 2025. Unleash LLMs Potential for Sequential Recommendation by Coordinating Dual Dynamic Index Mechanism. In *Proceedings of the ACM on Web Conference 2025*, 216–227.
- Zhai, J.; Liao, L.; Liu, X.; Wang, Y.; Li, R.; Cao, X.; Gao, L.; Gong, Z.; Gu, F.; He, M.; et al. 2024. Actions speak louder than words: Trillion-parameter sequential transducers for generative recommendations. *arXiv preprint arXiv:2402.17152*.
- Zhang, G.; Yuan, G.; Cheng, D.; Liu, L.; Li, J.; and Zhang, S. 2025a. Mitigating propensity bias of large language models for recommender systems. *ACM Transactions on Information Systems*, 43(6): 1–26.
- Zhang, J.; Guo, J.; Yuan, J.; Lu, H.; Zhou, Y.; Wu, F.; Wang, Q.; and Lu, D. 2025b. LLM-Driven Completeness and Consistency Evaluation for Cultural Heritage Data Augmentation in Cross-Modal Retrieval. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 19418–19428.
- Zhang, Q.; Chen, M.; Bukharin, A.; Karampatziakis, N.; He, P.; Cheng, Y.; Chen, W.; and Zhao, T. 2023. Adalora: Adaptive budget allocation for parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.10512*.
- Zhao, W. X.; Lin, Z.; Feng, Z.; Wang, P.; and Wen, J.-R. 2022. A Revisiting Study of Appropriate Offline Evaluation for Top-N Recommendation Algorithms. *ACM Trans. Inf. Syst.*, 41(2).
- Zheng, B.; Hou, Y.; Lu, H.; Chen, Y.; Zhao, W. X.; Chen, M.; and Wen, J.-R. 2024. Adapting large language models by integrating collaborative semantics for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 1435–1448. IEEE.
- Zheng, B.; Lu, H.; Chen, Y.; Zhao, W. X.; and Wen, J.-R. 2025. Universal Item Tokenization for Transferable Generative Recommendation. *arXiv preprint arXiv:2504.04405*.
- Zhu, F.; Wang, Y.; Chen, C.; Liu, G.; and Zheng, X. 2020. A Graphical and Attentional Framework for Dual-Target Cross-Domain Recommendation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, 3001–3008.