

Align³GR: Unified Multi-Level Alignment for LLM-based Generative Recommendation

Wencai Ye*, Mingjie Sun*, Shuhang Chen*, Wenjin Wu†, Peng Jiang

Kuaishou Technology, China

{yewencai, sunmingjie, chenshuhang, wuwenjin, jiangpeng}@kuaishou.com

Abstract

Large Language Models (LLMs) demonstrate significant advantages in leveraging structured world knowledge and multi-step reasoning capabilities. However, fundamental challenges arise when transforming LLMs into real-world recommender systems due to semantic and behavioral misalignment. To bridge this gap, we propose **Align³GR**, a novel framework that unifies token-level, behavior modeling-level, and preference-level alignment. Our approach introduces: Dual tokenization fusing user-item semantic and collaborative signals. Enhanced behavior modeling with bidirectional semantic alignment. Progressive DPO strategy combining self-play (SP-DPO) and real-world feedback (RF-DPO) for dynamic preference adaptation. Experiments show Align³GR outperforms the SOTA baseline by +17.8% in Recall@10 and +20.2% in NDCG@10 on the public dataset, with significant gains in online A/B tests and full-scale deployment on an industrial large-scale recommendation platform.

Introduction

Recommender systems (RS) are essential infrastructures in modern digital platforms such as e-commerce (Li and Karahanna 2015), video streaming (Lubos, Felfernig, and Tautschnig 2023), and social media (Narayanan 2023). With the rapid advancement of large language models (LLMs), researchers have explored two primary paradigms for integrating LLMs into recommender systems. Initially, LLMs have been employed to enhance traditional discriminative recommenders, for example, by providing improved content and user understanding (Wu et al. 2023), or enabling query rewriting and reasoning (Ye et al. 2023). More recently, LLMs have been developed as standalone generative recommenders, where the model directly outputs recommended items in an end-to-end manner (Rajput et al. 2023; Zheng et al. 2024). This shift from augmentation to full replacement brings forth a fundamental challenge: *How can LLMs be truly transformed into recommender systems?*

The key is effectively aligning the fundamental gap between the pre-trained LLMs and the personalized recommendation system (Wang et al. 2025b): On the one hand,

*These authors contributed equally.

†Corresponding Author.

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

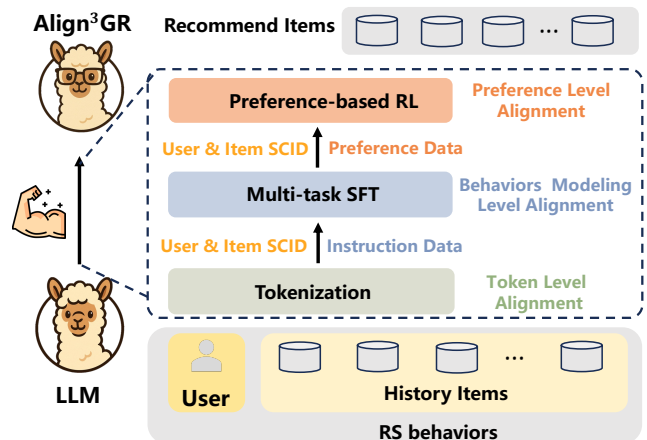


Figure 1: LLM-to-Recommendation Alignment Pipeline.

language modeling is primarily concerned with the semantic information with next-token prediction (NTP); on the other hand, recommender systems tend to model users' implicit preferences based on their interaction behavior information (Chen et al. 2024a). Recent research efforts (Li et al. 2025; Wang et al. 2024b) have attempted to bridge this gap. We systematically summarized them into a unified alignment pipeline spanning three indispensable levels: tokenization, multi-task supervised fine-tuning (SFT), and preference-based reinforcement learning (RL).

First, tokenization serves as the foundation by transforming user and item information into compact and expressive token representations (Jia et al. 2025), and EAGER (Wang et al. 2024b) and DAS (Ye et al. 2025) have incorporated collaborative signals at this stage to better bridge the gap between language modeling and downstream NTP tasks that require user-item interaction modeling. Second, the objective of the SFT stage is to enable general-purpose LLMs to comprehend the data structures and user behavior patterns inherent to recommendation tasks, thereby equipping the model with initial recommendation capabilities (Hong et al. 2025). Third, preference-based RL is essential for aligning model outputs with real user interests and business objectives, thereby enhancing the personalization of recommendations (Deng et al. 2025).

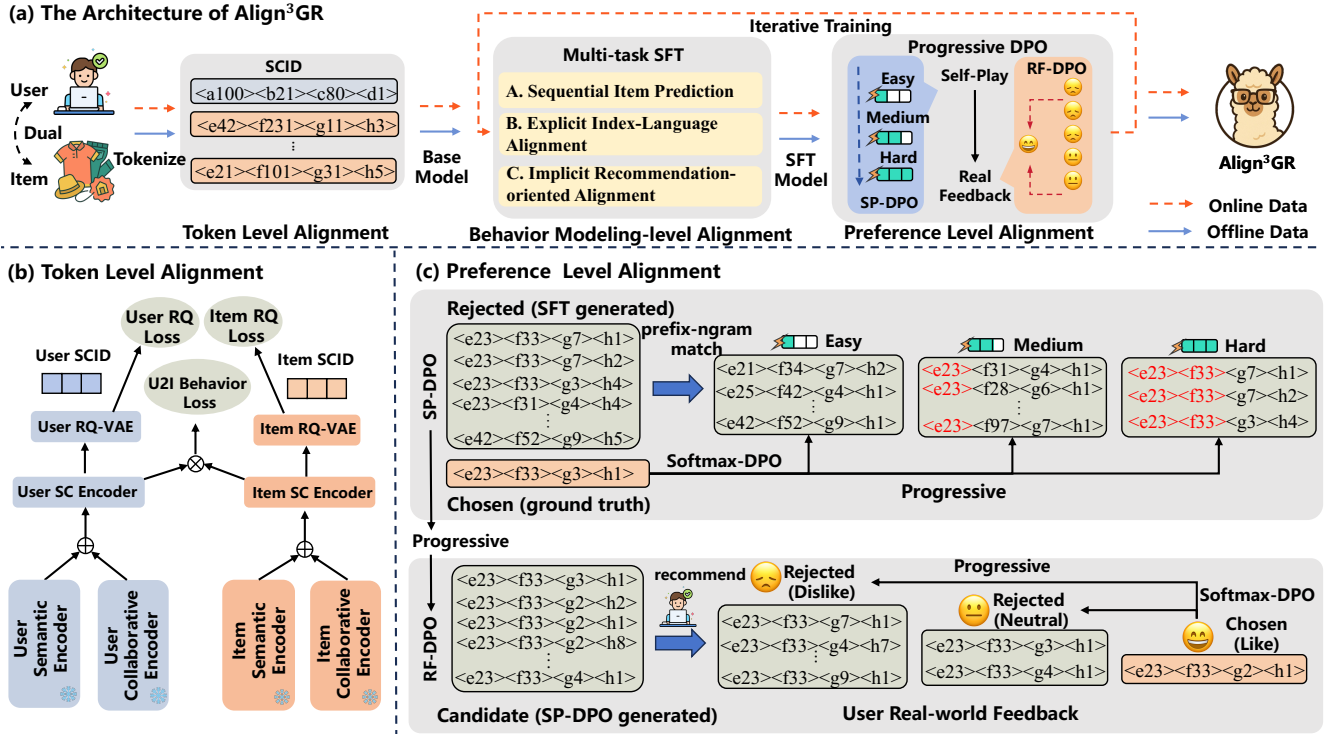


Figure 2: (a) The architecture of **Align³GR**, a unified multi-level alignment framework for generative recommendation, which integrates hierarchical dual SCID, multi-task SFT, and progressive DPO. (b) Token-level alignment is achieved through user-item dual SC encoders and RQ-VAEs. (c) Preference-level alignment is accomplished via progressive SP-DPO and RF-DPO.

Despite their merits, previous methods often treat user and item information independently during tokenization, neglecting the collaborative and semantic dependencies between them. Such isolated modeling fails to capture the mutual influences critical for comprehensive preference learning and has been shown to result in suboptimal recommendation performance (Zhang et al. 2025). On the other hand, aligning LLMs with real-world user preferences and business objectives presents additional challenges (Wang et al. 2025a; Tan et al. 2025). Although recent preference-based RL methods like DPO (Srivastava et al. 2025; Furuta et al. 2024) attempt to incorporate user preference signals into the RS, they typically rely on the collection of offline data without any progressive learning mechanisms. In real-world recommendation scenarios, user behavioral preferences are inherently complex and ambiguous compared to explicit labels, increasing the difficulty for models to learn directly.

To address these challenges, we propose **Align³GR**, a unified multi-level alignment framework that systematically integrates token-level, behavior modeling-level, and preference-level alignment. Specifically, our approach first introduces a hierarchical tokenization scheme, enabling compact and jointly optimized Semantic-Collaborative ID (SCID) for both users and items, which lays a solid foundation for subsequent SFT and preference-based RL, ensuring effective alignment throughout the entire recommendation pipeline. Next, during behavior modeling-level alignment,

we further introduce two additional tasks. We first incorporate the user’s SCID into the main sequence generation task, enriching the input representation and enabling the model to leverage more comprehensive user features. Then, we introduce a bidirectional alignment task between the user’s SCID and their semantic information, which explicitly grounds the SCID tokens in their real-world semantic meanings. Furthermore, inspired by curriculum learning (Liao et al. 2024), we introduce a progressive DPO strategy (easy to hard) that enables the model to improve continually by learning from preference pairs of increasing difficulty, leading to smoother convergence and more stable training. Building on this, we adopt the self-play DPO (SP-DPO), which involves interacting with itself to generate diverse training data, thereby performing cost-efficient adversarial bias correction on SFT models (Gao et al. 2025) and alleviating the exploration bottleneck of traditional RL. However, the lack of real-world feedback during self-play limits the model’s ability to generalize to real-world scenarios. Hence, we incorporate real-world feedback DPO (RF-DPO), progressively aligning the model with actual user interests and business objectives. Finally, we conduct comprehensive experiments on both public and industrial datasets to validate the effectiveness of our approach. On the Instruments dataset, **Align³GR** outperforms the SOTA baseline by 17.8% in Recall@10 and 20.2% in NDCG@10. Moreover, the gains are further corroborated by real-world business metrics in industrial set-

tings. The core contributions of this work are as follows:

1) We propose Align³GR, a unified alignment framework that jointly optimizes token-level, behavior modeling-level, and preference-level alignment to bridge the gap between LLMs and RS.

2) We design a dual tokenization scheme that enables semantic and collaborative fusion at the input level. In addition, we develop a progressive DPO strategy, including SP-DPO and RF-DPO, adapted for dynamic scenario preference alignment in RS.

3) Comprehensive experiments on benchmark and industrial datasets demonstrate consistent performance gains across all settings, with Align³GR yielding significant improvements in offline metrics and online A/B testing spanning diverse recommendation scenarios.

Related Works

Generative Recommendation

Generative recommendation (Wang et al. 2023) reframes classical retrieval as a sequence generation task, enabling models to benefit from complete semantic context and dynamically adaptable output formats. Generative approaches are different from standard embedding-based design (e.g., two-tower models with ANN search (Briskorn and Dientknecht 2019)) in that they autoregressively generate item tokens (e.g., IDs, titles, semantic IDs), enabling dynamic, context-aware recommendations. The three key technical directions are as follows: First, sequence generation retrieval (e.g., DSI (Tay et al. 2022; Chen et al. 2023) and GENRE (Si et al. 2023)) transforms retrieval into autoregressive user context token generation. Second, breakthroughs in indexing or tokenization, such as RQ-VAE (Lee et al. 2022), hierarchical k-means (Qi et al. 2017), and PQ (Jegou, Douze, and Schmid 2010), transform content embeddings into short discrete tokens, powering generative recommendation models like TIGER (Rajput et al. 2023), LC-Rec (Zheng et al. 2024), LETTER (Wang et al. 2024a), and EAGER-LLM (Hong et al. 2025).

Preference Alignment of LLMs

Alignment of LLMs with human preferences has been extensively studied in NLP, primarily through Reinforcement Learning from Human Feedback (RLHF) (Bai et al. 2022) and Direct Preference Optimization (DPO) (Rafailov et al. 2023). While RLHF guides policy updates via reward models built from human annotations, it suffers from instability and high computational costs. DPO addresses these limitations by directly optimizing model parameters on preference data, inspiring variants such as IPO (Yang, Tan, and Li 2025), cDPO (Furuta et al. 2024), rDPO (Qian et al. 2025), and Softmax-DPO (Chen et al. 2024b) to tackle noise robustness and unbiased learning. However, applying these techniques to recommender systems remains challenging due to sparse, implicit preference data, which limits static offline optimization and adaptation to dynamic user interests and business objectives (Deng et al. 2025). Recent advances, including curriculum learning (Liao et al. 2024) and self-play (Wu et al. 2024; Gao et al. 2025), have improved pref-

erence alignment. Notably, progressive strategies enhance robustness by organizing training from easy to hard preference pairs. Our work builds on these progressive strategies to bridge static preference optimization and the dynamic demands of large-scale recommender systems.

Methodology

Overview of the Proposed Framework

In this work, we propose Align³GR—a unified multi-level alignment framework that systematically bridges the gap between LLMs and recommender systems. As illustrated in Figure 2, our framework consists of three tightly aligned stages: First, at the token level, we introduce a novel tokenization scheme based on a user-item dual learning strategy, enabling dual-track fusion of user/item-specific semantic-collaborative features to generate hierarchical discrete SCID, while maintaining a compact token space and minimizing computational overhead during both training and inference. Second, at the behavior modeling level, we design an enhanced multi-task SFT based on LC-Rec (Zheng et al. 2024) by not only incorporating the SCID of users but also explicitly aligning these SCIDs with their semantic information through the user alignment task during explicit index-language alignment. Finally, at the preference alignment level, we draw inspiration from curriculum learning to propose a progressive DPO strategy (easy to hard), combining self-play DPO (SP-DPO) for continual self-improvement with real-world feedback DPO (RF-DPO), while simultaneously addressing preference learning from both generative performance and real user feedback perspectives. Together, these three levels form a coherent alignment pipeline that enables our LLM-based GR model to achieve both high-quality personalization and robust adaptation in large-scale dynamic recommendation.

Token-level Alignment: Dual SCID Tokenization

Current tokenization methods for generative recommendation primarily encode items while overlooking user structure modeling. Although some incorporate user representations, these are seldom co-optimized with item embeddings, resulting in suboptimal user-item alignment and representations that lack critical collaborative signals for recommendation. We argue that a truly effective tokenization scheme should jointly encode both users and items, leveraging their respective semantic and collaborative signals, and must leverage a unified, co-optimized framework to learn mutually aligned, expressive representations.

Dual SCID Tokenization addresses these gaps by integrating semantic and collaborative features for users and items within a unified dual learning framework (Figure 2b). Specifically, we first extract both semantic features (e.g., profiles or descriptions) and collaborative features (e.g., behavioral patterns) for users and items. These features are processed by dedicated encoders: a frozen semantic encoder (initialized with LMs like T5 (Ni et al. 2021)) captures textual representations, while a frozen collaborative encoder (e.g., DIN (Zhou et al. 2018)) models behavioral dynamics. The resulting semantic and collaborative embeddings are

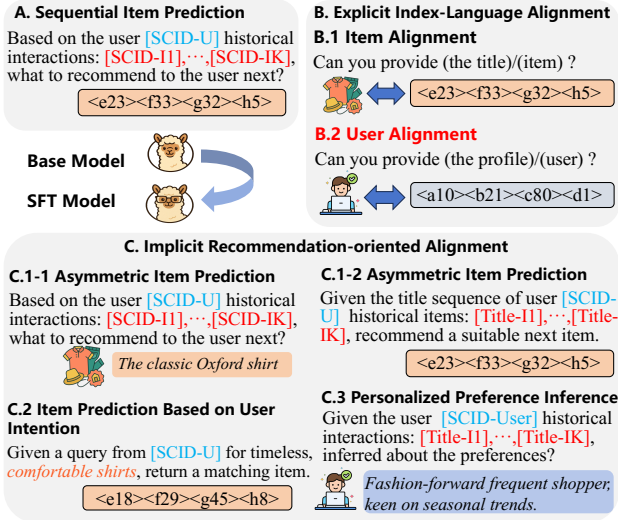


Figure 3: Behavior Modeling-level Alignment.

then concatenated and passed through a hybrid Semantic-Collaborative (SC) Encoder (e.g., MLP), which integrates both information types to produce unified SC embeddings (\mathbf{SC}_u , \mathbf{SC}_i). Finally, we quantize these unified embeddings into SCID using RQ-VAE (Lee et al. 2022). The training objective consists of two main components. First, at the embedding level, we optimize a sampled-softmax user-item behavior loss to enhance alignment between user and item SC embeddings. \mathcal{L}_{U2I} , defined as follows:

$$\mathcal{L}_{U2I} = -\frac{1}{|\mathcal{B}|} \sum_{(u, i^+) \in \mathcal{B}} \left[\log \frac{\exp(\mathbf{u}_u^\top \mathbf{v}_{i^+})}{\exp(\mathbf{u}_u^\top \mathbf{v}_{i^+}) + \sum_{j \in \mathcal{N}_u} \exp(\mathbf{u}_u^\top \mathbf{v}_j)} \right] \quad (1)$$

where \mathcal{B} is the batch size, and \mathbf{u} and \mathbf{v} represent \mathbf{SC}_u and \mathbf{SC}_i , the positive examples are user-item interaction behaviors. \mathcal{N}_u is the set of negative samples, which are randomly sampled within the batch. Second, the overall joint loss combines the U2I behavior loss and quantization losses from user-specific and item-specific RQ-VAE as:

$$\mathcal{L} = \alpha \cdot \mathcal{L}_{U2I} + \gamma \cdot (\mathcal{L}_{\text{User RQ}} + \mathcal{L}_{\text{Item RQ}}) \quad (2)$$

where $\mathcal{L}_{\text{User RQ}}$ and $\mathcal{L}_{\text{Item RQ}}$ denote the reconstruction and quantization losses for user and item embeddings, and α, γ are trade-off hyperparameters. In practice, we first set $\alpha = 1, \gamma = 0$ and optimize \mathcal{L}_{U2I} to stabilize behavior alignment and ensure thorough learning of the SC Encoder (monitored via AUC). Once stabilized, we switch to $\alpha = 0.1, \gamma = 1$ to focus on optimizing the quantization losses. During inference, user and item modules are deployed separately, each generating their own SCID for downstream use.

This design not only compresses the token space but also enables more effective downstream multi-task SFT and preference-based RL by preserving collaborative relationships throughout the model pipeline.

Behavior Modeling-level Alignment: Multi-task SFT

After obtaining quantized SCIDs for each user and item through alignment tokenization, we continue to enhance the

generative and semantic alignment capabilities of the LLM in the new token space. Following previous work (Zheng et al. 2024), we first expand the LLM vocabulary to include the SCID tokens of the user and the item, thus avoiding out-of-vocabulary issues and ensuring smooth integration with autoregressive generation tasks.

We formulate a multi-task SFT framework based on LC-Rec, which includes multiple tasks: Sequential Item Prediction, Asymmetric Item Prediction, Item Prediction Based on User Intention, and Personalized Preference Inference. These tasks aim to enhance the model’s ability to capture sequential dependencies, understand implicit user preferences, and align user behavior with items diversely and adaptively. However, LC-Rec remains limited in capturing user-item collaborative and semantic relationships. To address this, we propose two key enhancements: First, we inject the user’s SCID token into all task prompts, as shown in Figure 3, to ensure richer contextual alignment. Second, we introduce bi-directional alignment objectives (B.2): one task predicts a user’s SCID token from their profile text (text \rightarrow SCID), while the other reconstructs the user profile text from a given SCID token (SCID \rightarrow text).

Compared with prior works, our SFT design enriches user modeling by directly incorporating SCID tokens and explicitly aligns structured and semantic information through bi-directional tasks, providing a stronger foundation for downstream preference optimization.

Preference-level Alignment: Progressive DPO with Self-Play and Real-world Feedback

Although previous stages enable the model to have preliminary recommendation capability, simple preference optimization after SFT is insufficient for continual improvement or robust business alignment, due to the limited coverage of annotated preference data, which fails to capture the full complexity of real recommendation scenarios. To address this, we introduce progressive DPO with self-play (SP-DPO) and real-world feedback (RF-DPO). Specifically, SP-DPO first leverages self-play to acquire basic generative abilities, by generating diverse and informative data, thereby mitigating data sparsity and exploration limitations. Then RF-DPO utilizes real-world feedback to constrain the model toward real recommendation tasks, forming a synergistic, progressive learning strategy. The progressive DPO is based on the Softmax-DPO (Chen et al. 2024b) by constructing training samples containing multiple rejected responses, which is initialize as SFT model. The training objective for each stage is formally defined as:

$$\mathcal{L}(\pi_\theta^i, \pi_{\text{ref}}^i) = -\mathbb{E}_{(x, y_w^i, Y_i^i) \sim \mathcal{D}^i} \left[\log \sigma \left(-\log \sum_{y_i^i \in Y_i^i} \exp \left(\beta \log \frac{\pi_\theta^i(y_i^i | x)}{\pi_{\text{ref}}^i(y_i^i | x)} - \beta \log \frac{\pi_\theta^i(y_w^i | x)}{\pi_{\text{ref}}^i(y_w^i | x)} \right) \right) \right] \quad (3)$$

where π_θ^i denotes the current policy at stage i , π_{ref}^i is the reference policy, x is the prompt, y_w^i is the chosen response,

Y_l^i is the set of rejected responses, \mathcal{D}^i is the progressive training set at stage i and β is a hyperparameter, and $\sigma(\cdot)$ is the sigmoid function. The fine-tuned model at each stage becomes the reference model for the next stage ($\pi_\theta^i \rightarrow \pi_{\text{ref}}^{i+1}$), enabling the model to progressively adapt to more nuanced preference distinctions.

Progressive SP-DPO leverages self-play mechanisms to enhance the model’s generative capability by comparing its generated outputs with the ground truth. Specifically, considering the hierarchical nature of SCID, we divide progressive SP-DPO learning into three stages: Easy, Medium, and Hard, using a prefix-ngram match metric (Zheng et al. 2025) (the same prefix exhibits similar semantic and collaborative information). For the easy stage, the chosen and rejected SCID responses of preference data are completely different, with no shared prefix-ngram, making them easy to distinguish. For the medium and hard stages, the prefix-ngram overlap between chosen and rejected SCID responses progressively increases, increasing the difficulty of discrimination, but they remain non-identical. These three-stage preference data, coupled with real user behavior sequences, progressively serve as training data \mathcal{D}^i for preference learning, as shown in Eq (3). Alternatively, the prefix-ngram match metric can be extended to SCID vector-similarity metric, for a softer sample construction strategy.

Progressive RF-DPO captures authentic user feedback as preference data for alignment by recommending its own generated results to users. Feedback is categorized into three levels: disliked, neutral, and liked. Aligned with progressive learning, training occurs in stages: an easy stage uses strongly disliked items as negatives (liked as positives), while a hard stage uses neutral items as harder negatives (liked remains positive). This staged approach systematically strengthens preference learning. In industrial recommendation settings, levels are defined by user behavior: disliked (explicit negative, e.g., dislike), neutral (implicit negative, e.g., impression without click), liked (positive, e.g., like or purchase). For public datasets (e.g., Amazon reviews), we use an LLM-based sentiment model (ecomgpt (Li et al. 2023)) to score reviews, mapping scores to levels: disliked (1), neutral (2-3), and liked (4-5). Integrating this fine-grained feedback enables RF-DPO to better align with user interests and business goals, improving recommendation relevance.

Our progressive DPO framework offers several key advantages. By adopting a progression from easy to hard stages and leveraging both self-play and real-world feedback, the model continually enhances its ability to discern and generalize user preferences, overcoming the “preference ceiling” of static data. The key theoretical insight is that well-designed curricula can provide smoother interpolation between task distributions, leading to more efficient learning compared to direct approaches.

Experiments

Experimental Settings

Datasets. We conduct experiments on three real-world sequential recommendation datasets from diverse domains: (1)

Instruments, a subset of the Amazon review corpus focusing on user interactions with musical equipment; (2) Beauty, also from the Amazon review datasets, contains extensive user behaviors related to beauty products; and (3) Yelp, comprising user-business interactions from the Yelp challenge dataset. For fair comparison, we preprocess the data following standard protocols (Zheng et al. 2024; Rajput et al. 2023; Wang et al. 2024a), filtering users and items with fewer than five interactions and applying the leave-one-out strategy for splitting into training, validation, and test sets. We restrict each user’s history length to a maximum of 20 for all sequential models. Finally, we deploy Align³GR online and conduct A/B tests to further validate its performance on an industrial advertising recommendation platform.

Baselines. We compare our method with strong baselines from both conventional and generative recommender systems, covering various tokenization and alignment strategies: (1) **MF** (Mehta and Rana 2017) (Matrix Factorization); (2) **Caser** (Tang and Wang 2018); (3) **HGN**(Ma, Kang, and Liu 2019); (4) **BERT4Rec** (Sun et al. 2019); (5) **LightGCN** (He et al. 2020); (6) **SASRec** (Kang and McAuley 2018); (7) **BIGRec** (Bao et al. 2025), an LLM-based GR using item titles as textual identifiers; (8) **P5-SemID** (Wang et al. 2024a), leveraging item metadata for semantic identifiers; (9) **P5-CID** (Wang et al. 2024a), incorporating collaborative signals via clustering for LLM-based models; (10) **TIGER** (Rajput et al. 2023), which applies codebook-based quantized identifiers; (11) **LC-Rec** (Zheng et al. 2024), enhancing codebook tokenization with auxiliary alignment tasks; (12) **LETTER** (Wang et al. 2024a), a learnable tokenizer for generative recommendation; (13) **EAGER-LLM** (Hong et al. 2025), which further models user-item collaborative signals for token-level alignment. All baselines are implemented or adapted using open-source code where available.

Evaluation Metrics. We evaluate model performance using standard top- K metrics: Recall@ K ($R@K$) and NDCG@ K ($N@K$) for $K \in \{5, 10\}$, following previous work (Rajput et al. 2023; Zheng et al. 2024; Hong et al. 2025). During training, we limit each user’s historical interaction sequence to the most recent 20 items. For generative methods utilizing beam search, we follow EAGER-LLM and consistently set the beam width to 20.

Implementation Details. Our method is instantiated on Llama2-7B (Touvron et al. 2023) as the backbone LLM, with LoRA-based parameter-efficient fine-tuning (Hu et al. 2022). For item tokenization, we use a 3-level RQ-VAE, each codebook containing 256 embeddings of dimension 32. SCID representations for both users and items are incorporated into the model vocabulary to prevent OOV. Training is performed for 20,000 steps using the AdamW optimizer with a batch size of 1024. The learning rate is selected from the set $1e-3$, $5e-4$, $1e-4$ based on validation performance. All experiments are conducted on 4 NVIDIA RTX A800 GPUs. Hyperparameters such as α and β are tuned on the validation set. For each sample of Softmax-DPO, the number chosen is set to 1, and the number rejected is set to 20. During evaluation, we report the average results over five runs with different random seeds.

Model	Instruments				Beauty				Yelp			
	R@5	R@10	N@5	N@10	R@5	R@10	N@5	N@10	R@5	R@10	N@5	N@10
Traditional Recommendation Methods												
MF	0.0479	0.0735	0.0330	0.0412	0.0294	0.0474	0.0145	0.0191	0.0220	0.0381	0.0138	0.0190
LightGCN	0.0794	0.1000	0.0662	0.0728	0.0305	0.0511	0.0194	0.0260	0.0248	0.0407	0.0156	0.0207
Sequential Recommendation Methods												
Caser	0.0543	0.0710	0.0355	0.0409	0.0205	0.0347	0.0131	0.0176	0.0150	0.0263	0.0099	0.0134
HGN	0.0813	0.1048	0.0668	0.0774	0.0325	0.0512	0.0206	0.0266	0.0186	0.0326	0.0115	0.0159
Bert4Rec	0.0671	0.0822	0.0560	0.0608	0.0203	0.0347	0.0124	0.0170	0.0186	0.0291	0.0115	0.0159
SASRec	0.0751	0.0947	0.0627	0.0690	0.0380	0.0588	0.0246	0.0313	0.0183	0.0296	0.0116	0.0152
BigRec	0.0513	0.0576	0.0470	0.0491	0.0243	0.0299	0.0181	0.0198	0.0154	0.0169	0.0137	0.0142
Generative and LLM-based Recommendation Methods												
P5-SemID	0.0775	0.0964	0.0669	0.0730	0.0393	0.0584	0.0273	0.0335	0.0202	0.0324	0.0131	0.0170
P5-CID	0.0809	0.0987	0.0695	0.0751	0.0404	0.0597	0.0284	0.0347	0.0219	0.0347	0.0140	0.0181
TIGER	0.0870	0.1058	0.0737	0.0797	0.0395	0.0610	0.0262	0.0331	0.0253	0.0407	0.0164	0.0213
LETTER-TIGER	0.0909	0.1122	0.0763	0.0831	0.0431	0.0672	0.0286	0.0364	0.0277	0.0426	0.0184	0.0231
LC-Rec	0.0824	0.1006	0.0712	0.0772	0.0443	0.0642	0.0311	0.0374	0.0230	0.0359	0.0158	0.0199
LETTER-LC-Rec	0.0913	0.1115	0.0789	0.0854	0.0505	0.0703	0.0355	0.0418	0.0255	0.0393	0.0168	0.0211
EAGER-LLM	0.0991	0.1224	0.0851	0.0926	0.0548	0.0830	0.0369	0.0459	0.0373	0.0569	0.0251	0.0315
Align ³ GR	0.1103	0.1442	0.0947	0.1113	0.0627	0.0994	0.0434	0.0529	0.0425	0.0679	0.0299	0.0403
Improvement	+11.3%	+17.8%	+11.3%	+20.2%	+14.4%	+19.8%	+17.6%	+15.3%	+13.9%	+19.3%	+19.1%	+27.9%

Table 1: Overall performance comparison across public datasets.

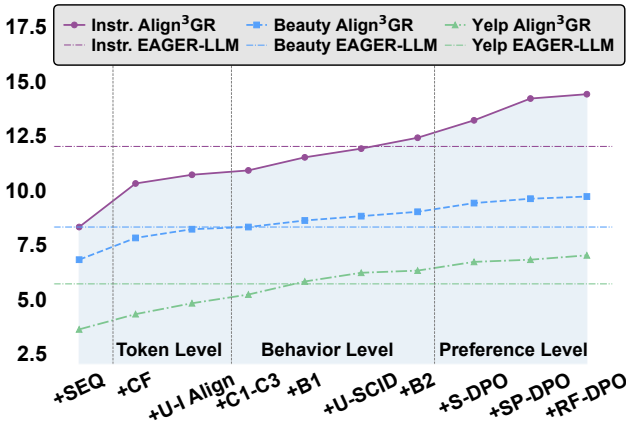


Figure 4: Recall@10 (%) under incremental alignment configurations; “+SEQ” denotes using Item-SCID as tokens for the sequence task, “U-SCID” denotes User-SCID, “S-DPO” denotes Softmax-DPO, “U-I Align” denotes U-I alignment, and the “+” indicates cumulative module addition.

Offline Performance

Table 1 reports the overall offline performance of all methods on three benchmark datasets. From the results, we have the following observations:

First, Align³GR consistently achieves the best or competitive results across all datasets and evaluation metrics. Compared to the strongest generative baselines such as LETTER and LC-Rec, our method yields substantial improvements with $p < 0.05$, especially on the Instruments dataset, where it surpasses EAGER-LLM by 17.8% in Recall@10 and 20.2% in NDCG@10. This demonstrates the advantage of our multi-level alignment strategy in effectively capturing complex user preferences and collaborative relationships.

Second, We visualize the incremental improvements

	Baseline	TIGER	Align ³ GR
Recall@100	0.218	0.229	0.242
Revenue (Improve.)	-	0.555%↑	1.432%↑

Table 2: Performance comparison in industrial scenario.

under incremental alignment configurations. Figure 4 reveals a clear upward trajectory. Notably, replacing item-side semantic IDs with dual learning-based item-side SCIDs yields a sharp performance boost in token-level alignment, underscoring the value of modeling collaborative semantics at the token level. The preference-level alignment stage yields the most substantial improvement, underscoring the effectiveness of progressive DPO with self-play and real-world feedback. Throughout all stages, our method significantly outperforms the SOTA baseline (EAGER-LLM), confirming the superiority of multi-level alignment in bridging the gap between LLMs and RS.

Finally, these results underscore the effectiveness of Align³GR. By jointly optimizing user and item representations at the token, behavior, and preference levels, Align³GR is able to capture collaborative signals and dynamic user intents comprehensively. The significant performance gains across offline evaluations confirm that each level of alignment, together with our unified tokenization and adaptive preference optimization strategies, is essential for achieving robust, scalable generative recommendation.

Online A/B Test

We further validate the practical efficacy of Align³GR using an internal industrial dataset and online A/B test. On a large-scale industrial advertising recommendation platform, we allocated 10% of the traffic (approximately 40+ million users) over multiple weeks for A/B testing. As Table 2 demonstrates, Align³GR surpasses both an industrial two-tower retrieval baseline and the generative TIGER

Tokenization	CF	U-I Alignment	Recall@10	NDCG@10
Item	✗	✗	0.1322	0.0978
Item	✓	✗	0.1346	0.0991
Dual	✗	✗	0.1390	0.1032
Dual	✗	✓	0.1426	0.1083
Dual	✓	✗	0.1428	0.1091
Dual	✓	✓	0.1442	0.1113

Table 3: Ablation study on dual SCID tokenization. Specifically, “Single” denotes using only item-side tokenization without user-item joint encoding, while “Dual” activates both user and item tokenization. The “CF” column indicates whether collaborative features are included, and “U-I Alignment” specifies if user and item embeddings are jointly optimized by U2I behavior loss.

model in online retrieval performance (recall@100) while significantly improving critical business metrics. Specifically, Align³GR achieves a statistically significant **+1.432%** revenue improvement in all advertising scenarios under full-scale deployment. These results confirm that Align³GR’s multi-level alignment translates offline advantages into measurable business value in production environments at scale.

Ablation Study

To thoroughly investigate the effect of each alignment level, we conduct comprehensive ablation studies on the Instruments dataset. *To ensure fairness, when ablations are performed on one alignment level, the configurations of the remaining levels are fixed to their best settings.* For example, when studying token-level alignment, we adopt the full setting of SFT tasks and the progressive RF-DPO strategy.

Effect of Dual SCID Tokenization. Table 3 presents the ablation results on the Instruments dataset to evaluate the effectiveness of Dual SCID Tokenization. First, simply switching from single-sided (item-only) to dual-sided tokenization yields a significant performance boost, demonstrating the necessity of jointly modeling both user and item token representations. Second, incorporating CF on top of semantic inputs further enhances performance under both single and dual settings, highlighting the importance of integrating collaborative signals during representation learning. Third, enabling U-I alignment via the U2I behavior loss leads to consistent gains, especially when combined with dual tokenization and CF. This confirms the efficacy of our joint optimization strategy in bridging the gap between user and item SC embeddings.

Effect of Behavior Modeling-level Alignment Tasks. Table 4 presents the ablation results for multi-task SFT on the Instruments dataset. We begin with the SEQ task as the backbone, using the item SCIDs from users’ historical behaviors as the sequence. Introducing user SCID yields additional gains, indicating that structured and informative token representations help the LLM better capture user-item interaction semantics. Furthermore, the inclusion of user-side alignment (B_2), which introduces bidirectional supervision between user profiles and SCID, brings the most significant

Methods	Recall@5	Recall@10	NDCG@5	NDCG@10
SEQ	0.1042	0.1329	0.0867	0.0982
+ $C_1 - C_3$	0.1046	0.1344	0.0881	0.0988
+ B_1	0.1054	0.1399	0.0908	0.1045
+ User SCID	0.1091	0.1417	0.0937	0.1051
+ B_2	0.1103	0.1442	0.0959	0.1113

Table 4: Ablation study of various semantic alignment tasks.

DPO Variant	Self-Play	Progressive	Recall@10	NDCG@10
Softmax-DPO	✗	✗	0.1295	0.0972
SP-DPO	✓	✗	0.1356	0.1033
SP-DPO	✓	✓	0.1396	0.1042
RF-DPO	✓	✗	0.1414	0.1049
RF-DPO	✓	✓	0.1442	0.1113

Table 5: Ablation study on preference-level alignment strategies. Using naive Softmax-DPO method as the baseline, chosen: the ground truth next item SCID; rejected: randomly sample 20 generated item SCIDs. The results of RF-DPO are based on progressive SP-DPO, as shown in Figure 2. The “Progressive” denotes the easy-to-hard learning.

performance boost across all metrics, highlighting the critical role of exposing LLMs to structured user semantics. These findings support our hypothesis that effective alignment requires supervision at both the token level and the semantic level, enabling the model to build stronger correspondence between language and recommendation signals.

Effect of Preference-Level Alignment Tasks. To assess the impact of our proposed preference-level alignment strategies, we perform an ablation study on SP-DPO and RF-DPO using the Instruments dataset, as shown in Table 5. Starting from a well-trained SFT model, we first examine the DPO baseline and then progressively incorporate Self-Play and Progressive strategies within each framework. In the SP-DPO branch, adding Self-Play improves Recall@10 from 0.1295 to 0.1356, and further gains are observed when applying progressive learning. In the RF-DPO branch, progressive training consistently outperforms the static variant, culminating in the best overall performance. These results demonstrate the complementary benefits of progressive optimization and real feedback integration in aligning LLMs with user preference signals.

Conclusion

We propose a unified multi-level alignment framework Align³GR that bridges the gap between LLMs and personalized recommendations through dual SCID tokenization, multi-task behavior modeling, and progressive preference optimization. Experiments on real-world datasets show our method outperforms strong baselines. Results highlight the importance of hierarchical alignment for LLM-based GRs, enabling more robust and adaptive personalization.

References

- Bai, Y.; Jones, A.; Ndousse, K.; Askell, A.; Chen, A.; Das-Sarma, N.; Drain, D.; Fort, S.; Ganguli, D.; Henighan, T.; et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Bao, K.; Zhang, J.; Wang, W.; Zhang, Y.; Yang, Z.; Luo, Y.; Chen, C.; Feng, F.; and Tian, Q. 2025. A bi-step grounding paradigm for large language models in recommendation systems. *ACM Transactions on Recommender Systems*, 3(4): 1–27.
- Briskorn, D.; and Dienstknecht, M. 2019. Mixed-integer programming models for tower crane selection and positioning with respect to mutual interference. *European Journal of Operational Research*, 273(1): 160–174.
- Chen, J.; Liu, Z.; Huang, X.; Wu, C.; Liu, Q.; Jiang, G.; Pu, Y.; Lei, Y.; Chen, X.; Wang, X.; et al. 2024a. When large language models meet personalization: Perspectives of challenges and opportunities. *World Wide Web*, 27(4): 42.
- Chen, X.; Liu, Y.; He, B.; Sun, L.; and Sun, Y. 2023. Understanding differential search index for text retrieval. *arXiv preprint arXiv:2305.02073*.
- Chen, Y.; Tan, J.; Zhang, A.; Yang, Z.; Sheng, L.; Zhang, E.; Wang, X.; and Chua, T.-S. 2024b. On softmax direct preference optimization for recommendation. *Advances in Neural Information Processing Systems*, 37: 27463–27489.
- Deng, J.; Wang, S.; Cai, K.; Ren, L.; Hu, Q.; Ding, W.; Luo, Q.; and Zhou, G. 2025. Onerec: Unifying retrieve and rank with generative recommender and iterative preference alignment. *arXiv preprint arXiv:2502.18965*.
- Furuta, H.; Lee, K.-H.; Gu, S. S.; Matsuo, Y.; Faust, A.; Zen, H.; and Gur, I. 2024. Geometric-averaged preference optimization for soft preference labels. *Advances in Neural Information Processing Systems*, 37: 57076–57114.
- Gao, C.; Chen, R.; Yuan, S.; Huang, K.; Yu, Y.; and He, X. 2025. Sprec: Self-play to debias llm-based recommendation. In *Proceedings of the ACM on Web Conference 2025*, 5075–5084.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. Lightgen: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 639–648.
- Hong, M.; Xia, Y.; Wang, Z.; Zhu, J.; Wang, Y.; Cai, S.; Yang, X.; Dai, Q.; Dong, Z.; Zhang, Z.; et al. 2025. EAGER-LLM: Enhancing Large Language Models as Recommenders through Exogenous Behavior-Semantic Integration. In *Proceedings of the ACM on Web Conference 2025*, 2754–2762.
- 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.
- Jegou, H.; Douze, M.; and Schmid, C. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1): 117–128.
- Jia, J.; Gao, J.; Xue, B.; Wang, J.; Cai, Q.; Chen, Q.; Zhao, X.; Jiang, P.; and Gai, K. 2025. From principles to applications: A comprehensive survey of discrete tokenizers in generation, comprehension, recommendation, and information retrieval. *arXiv preprint arXiv:2502.12448*.
- 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, G.; Zhang, X.; Zhang, Y.; Yin, Y.; Yin, G.; and Lin, W. 2025. 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, S. S.; and Karahanna, E. 2015. Online recommendation systems in a B2C E-commerce context: a review and future directions. *Journal of the association for information systems*, 16(2): 2.
- Li, Y.; Ma, S.; Wang, X.; Huang, S.; Jiang, C.; Zheng, H.-T.; Xie, P.; Huang, F.; and Jiang, Y. 2023. EcomGPT: Instruction-tuning Large Language Models with Chain-of-Task Tasks for E-commerce. *arXiv:2308.06966*.
- Liao, J.; Li, S.; Yang, Z.; Wu, J.; Yuan, Y.; Wang, X.; and He, X. 2024. Llara: Large language-recommendation assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1785–1795.
- Lubos, S.; Felfernig, A.; and Tautschnig, M. 2023. An overview of video recommender systems: state-of-the-art and research issues. *Frontiers in big Data*, 6: 1281614.
- Ma, C.; Kang, P.; and Liu, X. 2019. Hierarchical gating networks for sequential recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 825–833.
- Mehta, R.; and Rana, K. 2017. A review on matrix factorization techniques in recommender systems. In *2017 2nd International Conference on Communication Systems, Computing and IT Applications (CSCITA)*, 269–274. IEEE.
- Narayanan, A. 2023. Understanding social media recommendation algorithms.
- Ni, J.; Abrego, G. H.; Constant, N.; Ma, J.; Hall, K. B.; Cer, D.; and Yang, Y. 2021. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. *arXiv preprint arXiv:2108.08877*.
- Qi, J.; Yu, Y.; Wang, L.; Liu, J.; and Wang, Y. 2017. An effective and efficient hierarchical K-means clustering algorithm. *International Journal of Distributed Sensor Networks*, 13(8): 1550147717728627.
- Qian, W.; Wang, C.; Peng, H.; Tan, Z.; Li, H.; and Zeng, A. 2025. RDPO: Real Data Preference Optimization for Physics Consistency Video Generation. *arXiv preprint arXiv:2506.18655*.

- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741.
- 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. 2023. Generative retrieval with semantic tree-structured item identifiers via contrastive learning. *arXiv preprint arXiv:2309.13375*.
- Srivastava, P.; Nalli, S. S.; Deshpande, A.; and Sharma, A. 2025. Outlier-Aware Preference Optimization for Large Language Models. In *ICLR 2025 Workshop on Bidirectional Human-AI Alignment*.
- 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.
- Tan, Z.; Zeng, Z.; Zeng, Q.; Wu, Z.; Liu, Z.; Mo, F.; and Jiang, M. 2025. Can Large Language Models Understand Preferences in Personalized Recommendation? *arXiv preprint arXiv:2501.13391*.
- Tang, J.; and Wang, K. 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.
- Tay, Y.; Tran, V.; Dehghani, M.; Ni, J.; Bahri, D.; Mehta, H.; Qin, Z.; Hui, K.; Zhao, Z.; Gupta, J.; et al. 2022. Transformer memory as a differentiable search index. *Advances in Neural Information Processing Systems*, 35: 21831–21843.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, J.; Liu, Y.; Sun, Y.; Ma, X.; Wang, Y.; Ma, H.; Su, Z.; Chen, M.; Gao, M.; Dalal, O.; et al. 2025a. User Feedback Alignment for LLM-powered Exploration in Large-scale Recommendation Systems. *arXiv preprint arXiv:2504.05522*.
- Wang, W.; Bao, H.; Lin, X.; Zhang, J.; Li, Y.; Feng, F.; Ng, S.-K.; and Chua, T.-S. 2024a. Learnable item tokenization for generative recommendation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2400–2409.
- Wang, W.; Lin, X.; Feng, F.; He, X.; and Chua, T.-S. 2023. Generative recommendation: Towards next-generation recommender paradigm. *arXiv preprint arXiv:2304.03516*.
- Wang, Y.; Xun, J.; Hong, M.; Zhu, J.; Jin, T.; Lin, W.; Li, H.; Li, L.; Xia, Y.; Zhao, Z.; et al. 2024b. EAGER: Two-Stream Generative Recommender with Behavior-Semantic Collaboration. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3245–3254.
- Wang, Z.; Lin, J.; Yang, X.; Liu, Y.; Feng, S.; Wang, D.; and Zhang, Y. 2025b. Enhancing LLM-based Recommendation through Semantic-Aligned Collaborative Knowledge. *arXiv preprint arXiv:2504.10107*.
- Wu, L.; Zheng, Z.; Qiu, Z.; and 2023. A Survey on Large Language Models for Recommendation. *CoRR*, abs/2305.19860.
- Wu, Y.; Sun, Z.; Yuan, H.; Ji, K.; Yang, Y.; and Gu, Q. 2024. Self-play preference optimization for language model alignment. *arXiv preprint arXiv:2405.00675*.
- Yang, X.; Tan, Z.; and Li, H. 2025. Ipo: Iterative preference optimization for text-to-video generation. *arXiv preprint arXiv:2502.02088*.
- Ye, F.; Fang, M.; Li, S.; and Yilmaz, E. 2023. Large Language Model-Aided Informative Query Rewriting. In *EMNLP Findings*.
- Ye, W.; Sun, M.; Shi, S.; Wang, P.; Wu, W.; and Jiang, P. 2025. DAS: Dual-Aligned Semantic IDs Empowered Industrial Recommender System. *arXiv preprint arXiv:2508.10584*.
- Zhang, Y.; Feng, F.; Zhang, J.; Bao, K.; Wang, Q.; and He, X. 2025. Collm: Integrating collaborative embeddings into large language models for recommendation. *IEEE Transactions on Knowledge and Data Engineering*.
- 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, C.; Huang, M.; Pedchenko, D.; Rangadurai, K.; Wang, S.; Nahum, G.; Lei, J.; Yang, Y.; Liu, T.; Luo, Z.; et al. 2025. Enhancing Embedding Representation Stability in Recommendation Systems with Semantic ID. *arXiv preprint arXiv:2504.02137*.
- Zhou, G.; Zhu, X.; Song, C.; Fan, Y.; Zhu, H.; Ma, X.; Yan, Y.; Jin, J.; Li, H.; and Gai, K. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1059–1068.