LLM-Aligned Geographic Item Tokenization for Local-Life Recommendation

Hao Jiang[1] *, Guoquan Wang[1] *, Donglin Zhou[1†], Sheng Yu[1], Yang Zeng[1], Wencong Zeng[1], Kun Gai[2], Guorui Zhou[1] ‡

[1]Kuaishou Technology, Beijing, China
[2]Independent Researcher
{jianghao11, wangguoquan03, zhoudonglin, yusheng03, zhengchengyi, zengwencong, zhouguorui}@kuaishou.com,
gai.kun@qq.com

Abstract

Recent advances in Large Language Models (LLMs) have enhanced text-based recommendation by enriching traditional ID-based methods with semantic generalization capabilities. Text-based methods typically encode item textual information via prompt design and generate discrete semantic IDs through item tokenization. However, in domain-specific tasks such as local-life services, simply injecting location information into prompts fails to capture fine-grained spatial characteristics and real-world distance awareness among items. To address this, we propose **LGSID**, an LLM-Aligned Geographic Item Tokenization Framework for Local-life Recommendation. This framework consists of two key components: (1) **RL-based Geographic LLM Alignment**, and (2) **Hierarchical Geographic Item Tokenization**. In the RL-based alignment module, we initially train a list-wise reward model to capture real-world spatial relationships among items. We then introduce a novel G-DPO algorithm that uses pre-trained reward model to inject generalized spatial knowledge and collaborative signals into LLMs while preserving their semantic understanding. Furthermore, we propose a hierarchical geographic item tokenization strategy, where primary tokens are derived from discrete spatial and content attributes, and residual tokens are refined using the aligned LLM's geographic representation vectors. Extensive experiments on real-world Kuaishou industry datasets show that LGSID consistently outperforms state-of-the-art discriminative and generative recommendation models. Ablation studies, visualizations, and case studies further validate its effectiveness.

Code — https://github.com/JiangHaoPG11/LGSID

Introduction

With the rapid growth of local-life services, recommendation systems have become essential for meeting users' daily needs on major platforms such as Kuaishou and Meituan. However, traditional ID-based methods, which represent items using unique ID tokens and rely heavily on collaborative filtering signals (He et al. 2017), struggle to capture the real-world spatial characteristics and distance awareness of local-life recommendation. In such spatially constrained scenarios (Ma et al. 2024; Zhao et al. 2025), many items suffer from limited user interaction opportunities, leading to unfair exposure. Consequently, as the candidate item corpus grows, these methods encounter inherent performance bottlenecks (Zhang et al. 2024c; Wu et al. 2021).

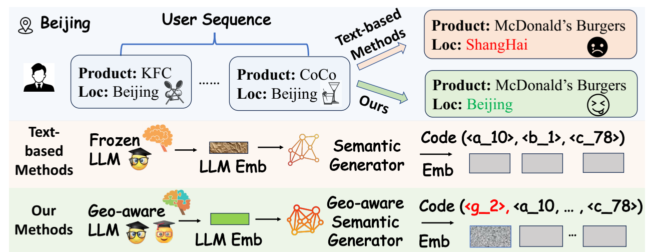

Figure 1: Illustration of challenges of text-based methods in local-life recommendation. Without geographic awareness, the system may recommend semantically relevant items that are inaccessible to users due to long distance.

Recent advances in Large Language Models (LLMs) offer an effective solution to the limited learning capacity of ID embeddings by leveraging their semantic understanding abilities (Achiam et al. 2023; Yang et al. 2025; Wang et al. 2024b). The prevailing paradigm in text-based methods, as illustrated in the middle part of Figure 1, involves designing prompts that incorporate domain features and item attributes for LLMs to generate semantic representations, which are subsequently quantized into discrete semantic IDs through a process known as item tokenization. For example, TIGER (Rajput et al. 2023) is one of the first to propose the RQ-VAE quantization model, which maps LLM representations to semantic IDs (SIDs). QARM (Luo et al. 2024) introduces the Res-Kmeans quantization model to tackle two key challenges in leveraging LLM representations: mismatch and unlearning. Furthermore, N-gram semantic IDs (Singh et al. 2024; Zheng et al. 2025b) and Cosine semantic IDs (Lin et al. 2025) have been developed to address the long-tail problem in existing quantization methods. To better align SIDs with recommendation, several methods further inject collaborative filtering signals during or after item tokenization, while underscoring the potential of item tokenization for recommendation (Wang et al. 2024a; Liu et al. 2025a).

*Equal contributions
†Work done during an internship at Kuaishou Technology
‡Corresponding author

Despite their successes, we argue that these methods still face two major challenges. (a) **Limited adaptation to downstream recommendation tasks with domain-specific knowledge.** They often treat LLMs merely as text encoders that transfer item textual descriptions into high-quality semantic representation, while focusing on designing novel quantization models to align with diverse downstream tasks. However, our key insight is that the quality and domain awareness of upstream LLMs fundamentally determine the upper bound of item tokenization performance. (b) **Weak integration of domain knowledge and semantic understanding.** They simply combine domain-specific signals with content attributes into prompts, assuming LLMs can naturally balance their importance. However, pre-trained LLMs often misinterpret such signals and tend to prioritize content relevance over geographic proximity. As shown in Figure 1, they may recommend a Shanghai restaurant to a user located in Beijing. This highlights the importance of item tokenization needs integrates item attributes and high-dimensional representations to better fuse semantic understanding with geographic awareness.

To overcome the limitations of existing text-based recommendation methods, we propose **LGSID**, an LLM-Aligned Geographic Item Tokenization framework for Local-life Recommendation. Unlike previous methods that focus on quantization models, LGSID adopts a post-training strategy that aligns LLM to optimize representations, thereby equipping them with real-world spatial awareness. To this end, we integrates a novel geographic item tokenization framework through a two-stage module: (1) **RL-based Geographic LLM Alignment** and (2) **Hierarchical Geographic Item Tokenization**. The RL-based alignment module first trains reward models using distance-aware list-wise sampling to capture and compress real-world spatial relationships. Building upon this, we introduce the G-DPO algorithm to incorporate geographic and collaborative filtering signals into the LLM. During LLM training, the G-DPO algorithm uses similarity regularization to dynamically balance semantic accuracy and geographic awareness, ensuring that domain-specific knowledge is effectively injected into the LLM while preserving its semantic understanding capabilities. Moreover, we propose a hierarchical geographic item tokenization method with a novel quantization strategy. This method first generates primary tokens based on spatial and content attributes, then refines residual tokens by leveraging the aligned LLM's geographic representations.

In summary, our main contributions are as follows.

- We identify the limitations of existing LLM-driven item tokenization methods and highlight the importance of aligning LLMs with domain-specific knowledge.
- We propose LGSID, a two-stage item tokenization framework tailored for spatially constrained scenarios such as local-life recommendation, with RL-based LLM alignment and hierarchical geographic item tokenization.
- We conduct comprehensive experiments on real-world industry datasets, demonstrating that LGSID significantly improves performance in both discriminative and generative recommendation models.

Related Work

Item Tokenization

Recent advances in item tokenization (e.g., RQ-VAE and VQ-VAE) have driven the development of text-based recommendation (Rajput et al. 2023). A key challenge in these methods is aligning item tokenization with recommendation. Existing item tokenization methods can be broadly categorized into two main types: two-stage methods and end-to-end methods. In the two-stage methods, early efforts such as LC-Rec (Zheng et al. 2024) introduces semantic alignment to integrate recommendation signals, while LETTER (Wang et al. 2024a) jointly aligns semantic and collaborative spaces. Building on these, QARM (Luo et al. 2024) learns semantic IDs using Res-Kmeans guided by user interaction distributions. EAGER (Wang et al. 2024c) adopts two-stream codebooks to model both semantic and collaborative signals, UTGRec (Zheng et al. 2025a) integrates multimodal semantics with co-occurrence patterns into universal codes. Furthermore, SC (Li et al. 2025) regularize semantic encoders with ID embeddings from pre-trained CF models. In end-to-end methods, tokenization and recommendation are jointly alignmented and optimized. UnifiedSID (Lin et al. 2025) employs RQ-VAE trained with cosine and euclidean distances to integrate semantic and ID tokens, while ETEGRec (Liu et al. 2025a) introduces sequence-item and preference-semantic alignment objectives with generative models training. However, these methods largely overlook domain-specific constraints. For example, in local-life recommendation, users interact only with items within a limited geographic radius. SIDs without geographic awareness often recommend items that match user interests but are outside deliverable distances, negatively affects system efficiency.

LLM Alignment for Recommendation

A key challenge in applying LLMs to recommendation is incorporating task-specific awareness. Existing methods can be categorized into two groups. The first focuses on designing LLMs tasks to stimulate real-world knowledge and reasoning. For example, LGHRec (Luo et al. 2025) uses chain-of-thought reasoning to distill item descriptions into semantic IDs fused with vanilla IDs for GNN. GNPR-SID (Wang et al. 2025) incorporates domain attributes like location information into prompts to capture downstream signals. SIIT (Chen et al. 2024) iteratively refines tokenization via self-improvement. The other group focuses on fine-tuning LLMs. NoteLLM (Zhang et al. 2024a) and NoteLLM-2 (Zhang et al. 2024b) compress items via prompts and integrate collaborative signals through supervised fine-tuning. AlignRec (Liu et al. 2024) introduces alignment objectives for multimodal and user–item consistency. LLMEmb (Liu et al. 2025b) uses supervised contrastive fine-tuning to align LLM embeddings with collaborative data. LARM (Liu et al. 2025c) distills open-source LLM knowledge into smaller models, while Lu et al. (Lu et al. 2024) strengthen LLMs' alignment with recommendation instructions. These methods emphasize semantic alignment. However, LLMs need balance domain-specific constraints while accurately capturing users' true preferences based on item content.

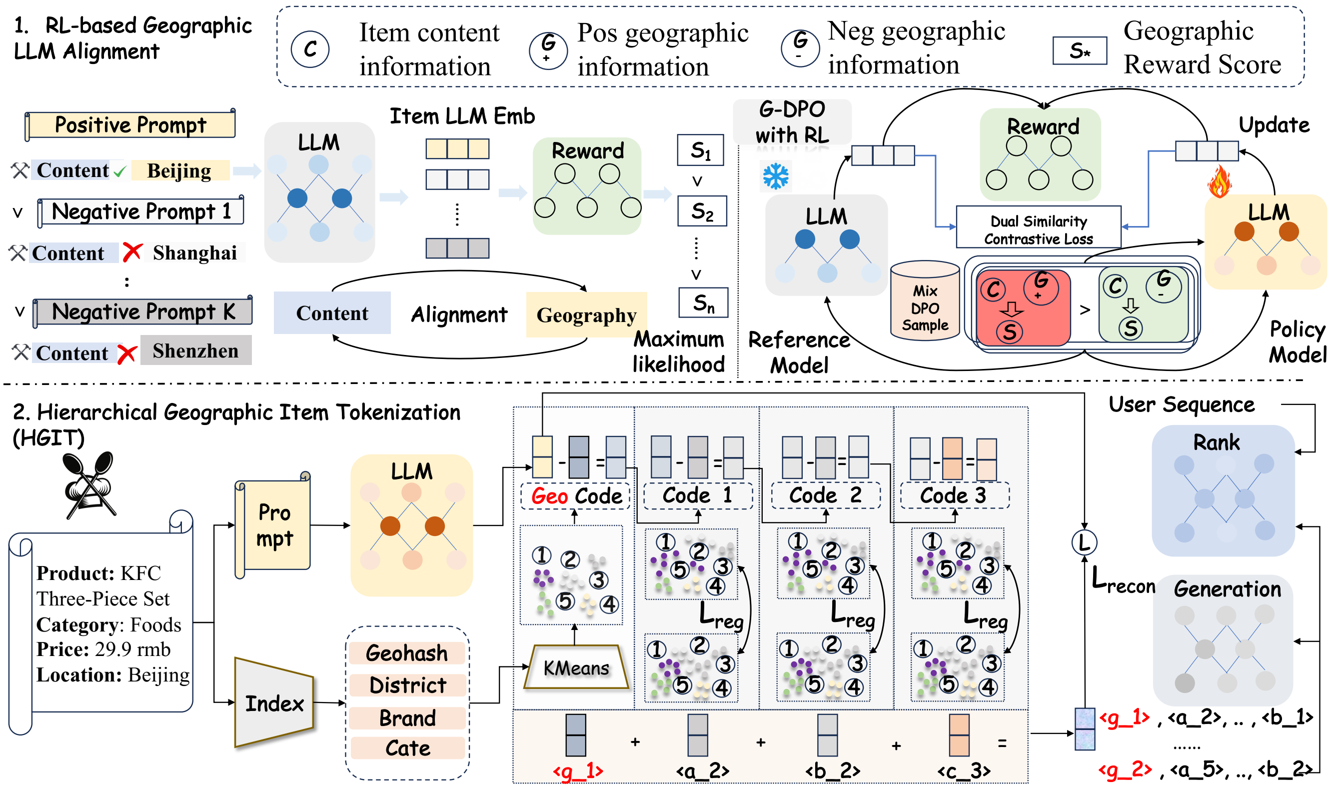

Figure 2: Model structure of LGSID. The upper part illustrates the pipeline of RL-based Geographic LLM Alignment, while the lower part depicts the pipeline of Hierarchical Geographic Item Tokenization.

Preliminaries

Geography Prompt Design

We design prompts for LLMs to derive item semantic representations by combining their textual descriptions $\mathbf{T}_i$ (e.g., name, brand, category, price) with geographical attributes (e.g., province, city, town), as detailed in Appendix: A. Specially, items are first encoded into semantic representations $\mathbf{E}_i$, which are then quantized into discrete semantic IDs S_i.

We apply our LGSID to both two mainstream recommendation paradigms with following pipelines. (1) discriminative recommendation: Users $\mathcal{U} = \{u_1, \dots, u_M\}$ and items $\mathcal{I} = \{i_1, \dots, i_N\}$ are encoded into their unique ID embeddings, where each item is represented by its textual description $\mathbf{T}_i$ and semantic IDs S_i. A user encoder $f_u(\cdot)$ learns $\mathbf{e}_u = f_u(u)$, and Top-k items $\{i_1, \dots, i_k\}$ are generated by matching $\mathbf{e}_u$ with $(ID_i, \mathbf{T}_i, S_i)$. (2) generative recommendation: Items are first mapped to semantic IDs $S_i = (s_i^1, \dots, s_i^L)$ via quantized tokenization. Given a user's SID history $\{S_{i_1}, \dots, S_{i_t}\}$, the model predicts the next SID $S_{i_{t+1}}$, corresponding to the most likely next item.

Methodology

Overview

In this paper, we propose a novel geographic item tokenization framework for local-life recommendation, comprising two key modules: RL-based Geographic LLM Alignment and Hierarchical Geographic Item Tokenization. The upper part of Figure 2 illustrates the RL-based alignment module, which aligns LLMs with real-world spatial knowledge while preserving their semantic understanding ability. The lower part of Figure 2 illustrates the HGIT pipeline, which fuses spatial and discrete content features with rich semantic representations derived from LLMs to better balance item semantics and domain-specific characteristics.

RL-based Geographic LLM Alignment

Existing methods often incorporate domain-specific information into prompts to enhance LLMs awareness (Wang et al. 2025). However, LLMs mainly rely on semantic similarity, capturing only coarse spatial relations from textual relevance and struggling with fine-grained distinctions between textually similar locations, e.g., confusing "Suzhou, Anhui, City" with "Suzhou, Jiangsu, City". To address this, we propose an RL-based post-training strategy to align LLMs with real-world geographic knowledge.

Geography-aware Reward Model Training. In real-world recommendation systems with large and dynamic candidate pools, manually labeling item pairs for reinforcement learning is impractical. To overcome this challenge, we first train a list-wise reward model $\mathcal{R}(i)$, which predicts the geographic relevance score between item content and its corresponding location based on their LLM representations. This allows the model to internalize generalizable spatial knowledge via a neural network.

Specifically, we first compute the pairwise geodesic distance between items using their latitude and longitude, and adopt a density-aware hard negative sampling strategy that selects K negative samples based on spatial distance. We construct prompt sequences by fixing the item content and replacing the location with that of each negative sample using a prompt mismatching strategy, as defined in Eq. 1.

$$P_i = \left[P_{\text{content}}, P^{\text{i}}_{\text{location}}\right], \quad P_i^{j^-} = \left[P_{\text{content}}, P^{\text{j}^-}_{\text{location}}\right], \quad (1)$$

where P_{content} is item text attributes, P^{i}_{location} is the true location, and $P^{j^-}_{\text{location}}$ is a negative location from sampled items. This generates a prompt sequences $P_i = [P_i, P_i^1, \ldots, P_i^k]$.

The prompts are encoded by the LLM into embeddings $\mathbf{E} = [\mathbf{E}^i, \mathbf{E}^i_1, \ldots, \mathbf{E}^i_k]$, where $\mathbf{E}^i_j \in \mathbb{R}^d$ and d denotes the embedding dimension. Meanwhile, we design a list-wise architecture to facilitate the learning of spatial proximity between items. Specifically, for each prompt-mismatched sampled item, the model feeds the LLM representation $\mathbf{E}^i_j$ into a multi-layer perceptron (MLP) to predict a reward score, which quantifies the relationship between the target item content and location j, as defined in Eq. 2.

$$r_{i,j} = \text{MLP}(\mathbf{E}^i_j), \quad (2)$$

where $r_{i,j}$ denotes the reward score between the i-th target item and j-th sampled item. Next, we assign soft labels $p_{i,j}$ based on the distances between items. Specifically, given a candidate list of K prompt sequence sorted by distance from near to far, we define the soft labels as shown in Eq. 3.

$$p_{i,j} = K - Rank(\text{dis}_{i,j}) + 1, \quad (3)$$

where $\text{dis}_{i,j}$ denotes the haversine distance between items i and j in the list. This distance-based labeling strategy encourages the model to prioritize items that are geographically closer to the target. The reward model is trained using a weighted binary cross-entropy loss, as defined in Eq. 4.

$$\mathcal{L}_{\text{RM}} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j} p_{i,j} \log \sigma(r_{i,j}), \quad (4)$$

where $\sigma(\cdot)$ is the sigmoid activation function, N is the batch size, and $\mathcal{L}_{\text{RM}}$ denotes the reward model's training loss.

G-DPO Algorithm with RM. Building upon pre-trained reward models, we propose G-DPO, an enhanced algorithm inspired by Direct Preference Optimization (DPO) (Rafailov et al. 2023), as illustrated in the upper right part of Figure 2. In G-DPO, we introduce a domain-mixed sampling strategy for preference learning, denoted by $\mathcal{D}_{\text{mix}} = \mathcal{D}_{\text{dc}} \cup \mathcal{D}_{\text{gc}}$. This database combines two data types: domain collaborative pairs and geography constrained pairs.

For domain collaborative pairs $\mathcal{D}_{\text{dc}}$, we utilize user historical behaviors to enhance collaborative signal awareness. Items frequently co-occurrence by users tend to exhibit both semantic and geographic similarity (Zhang et al. 2024a). Specifically, the co-occurrence score is defined in Eq. 5.

$$s_{i_a,i_b} = \sum_{u=1}^{U} \mathbb{I}\left[i_a \in H_u \wedge i_b \in H_u\right], \quad (5)$$

where $\mathbb{I}[\cdot]$ is the indicator function, and H_u represents user's click history. We retain pairs (i_a, i_b) with $s_{i_a,i_b} > s_{\text{th}}$ as domain collaborative sample pairs, where s_{th} is the threshold.

For geographically constrained pairs $\mathcal{D}_{\text{gc}}$, we randomly sample items outside the target item's location to form pairs (i_a, i_r), ensuring diversity and efficiency within million-scale candidate pools. Furthermore, G-DPO aligns the policy model π_θ with a reference model π_{ref} using domain-mixed sample pairs $(i^+, i^-) \in \mathcal{D}_{\text{mix}}$. Each item pair is scored by the reward model $\mathcal{R}$, which takes embeddings from either the policy $\mathbf{E}_{\pi_\theta}(i)$ or the reference model $\mathbf{E}_{\text{ref}}(i)$ and outputs a distance-based score. The alignment loss in G-DPO is given in Eq. 6.

$$\begin{aligned}\mathcal{L}_{\text{align}} = -\mathbb{E}_{(i^+,i^-)} \log \sigma \Big(\beta \Big(&\mathcal{R}(\mathbf{E}_{\pi_\theta}(i^+)) - \mathcal{R}(E_{\pi_\theta}(i^-)) \\ &-\mathcal{R}(\mathbf{E}_{\pi_{\text{ref}}}(i^+)) + \mathcal{R}(\mathbf{E}_{\pi_{\text{ref}}}(i^-))\Big)\Big),\end{aligned} \quad (6)$$

where β controls alignment sharpness, $\sigma(\cdot)$ is the sigmoid function, and $\mathcal{R}(\cdot)$ is the pre-trained reward model. To preserve LLMs' semantic representations during G-DPO updates, we introduce an in-batch contrastive loss as a similarity regularizer, as shown in Eq. 7.

$$\begin{aligned}\mathcal{L}_{\text{sim}} = \mathbb{E}_{i\in\mathcal{B}}\Big[&\|E_{\pi_\theta}(i) - E_{\pi_{\text{ref}}}(i)\|_2^2 \\ &- \frac{1}{|\mathcal{B}|-1} \sum_{j\neq i} \|E_{\pi_\theta}(i) - E_{\pi_{\text{ref}}}(j)\|_2^2\Big].\end{aligned} \quad (7)$$

For each instance i in a batch $\mathcal{B}$, we pull its representation closer to the reference model output while pushing it away from other reference embeddings in the same batch. We combine the contrastive and DPO loss with weight λ, as shown in Eq. 8.

$$\mathcal{L}_{\text{G-DPO}} = \mathcal{L}_{\text{align}} + \lambda \mathcal{L}_{\text{sim}}. \quad (8)$$

Hierarchical Geographic Item Tokenization

We introduce a Hierarchical Geographic Item Tokenization module on top of aligned LLM semantic representations.

For the first layer, we construct a multi-dimensional feature that integrates different types of attributes for geography-aware token initialization. To mitigate the inefficiency of high-dimensional one-hot encoding, we adopt composite embeddings for discrete categorical features. Specifically, geography-aware codes f_{geo} are formed by latitude and longitude. Meanwhile, administrative codes f_{admin} are fixed scaling factors used for normalization, determined by the province ID, city ID, and district ID. Similarly, category codes f_{cat} are fixed scaling factors based on the primary category and the secondary category. Brand codes f_{brand} are fixed scaling factors determined by the brand ID. The final clustering feature vector $\mathbf{F}$ is constructed via weighted concatenation of all components, as shown in Eq. 9.

$$\mathbf{F} = [w_{\text{admin}} \cdot f_{\text{admin}},\ w_{\text{geo}} \cdot f_{\text{geo}},\ w_{\text{cat}} \cdot f_{\text{cat}},\ w_{\text{brand}} \cdot f_{\text{brand}}], \quad (9)$$

where w_* are empirically chosen coefficients reflecting the relative importance of each feature type. Using clustering feature vector $\mathbf{F}$, we apply MiniBatch K-Means to generate a vocabulary of first-layer geographic tokens. The cluster representations are then computed as the mean of LLM embeddings within each token group, producing the first-layer cluster centers $\boldsymbol{\mu}^{(1)}$.

For residual layers ($l \geq 2$), we employ learnable cluster centers with euclidean distance-based assignment. The process of the cluster centers $\boldsymbol{\mu}^{(l)}$ with Eq. 10, while the quantized representation and residual update with Eq. 11 and 12.

$$\mathbf{z}^{(l)} = \arg\min_{k} \|\mathbf{R}^{(l-1)} - \boldsymbol{\mu}^{(l)}\|_2^2, \tag{10}$$

$$\mathbf{Q}^{(l)} = \boldsymbol{\mu}^{(l)}[\mathbf{z}^{(l)}], \tag{11}$$

$$\mathbf{R}^{(l)} = \mathbf{R}^{(l-1)} - \mathbf{Q}^{(l)}. \tag{12}$$

The primary objective is to minimize the reconstruction loss between the original embeddings and their quantized representations. For input embeddings $\mathbf{X}$, we apply a combination of absolute reconstruction losses with Eq. 13.

$$\mathcal{L}_{\text{recon}} = \|\mathbf{X} - \sum_{l=1}^{L} \mathbf{Q}^{(l)}\|_2^2. \tag{13}$$

To promote balanced utilization of the learned clusters and prevent cluster collapse, we introduce an entropy-based regularization term. For each layer l, the cluster usage distribution is computed as shown in Eq. 14.

$$p_k^{(l)} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}[\mathbf{z}_i^{(l)} = k], \tag{14}$$

where $\mathbb{I}[\cdot]$ is the indicator function. The regularization loss encourages uniform cluster usage through KL divergence:

$$\mathcal{L}_{\text{reg}}^{(l)} = \text{KL}\left(p^{(l)} \| \mathbf{u}\right) = \sum_{k=1}^{K_l} p_k^{(l)} \log \frac{p_k^{(l)}}{1/K_l}, \tag{15}$$

where $\mathbf{u} = [1/K_l, \ldots, 1/K_l]$ represents the uniform distribution over K_l clusters. The complete training objective combines reconstruction with cluster usage regularization:

$$\mathcal{L}_{\text{HGIT}} = \mathcal{L}_{\text{recon}} + \lambda_{\text{reg}} \sum_{l=2}^{L} \mathcal{L}_{\text{reg}}^{(l)}, \tag{16}$$

where λ_{reg} controls the strength of the regularization. Note that regularization is only applied to learnable layers ($l \geq 2$) since the first layer uses pre-computed geographic clusters.

Experiment

In this section, we conduct extensive experiments on real-world datasets to address the following research questions:

- **RQ1**: How does LGSID improve upon existing tokenization for SOTA discriminative and generative models?
- **RQ2**: How does fine-tuning affect the LLM's ability on enhancing geographic awareness while preserving semantic understanding?
- **RQ3**: How accurately do item representations and quantized IDs reflect real-world geographic proximity?
- **RQ4**: How does LGSID exhibit superior geographic awareness compared to other SID methods?

Experiment Settings

The datsets and evaluation, finetune settings, parameters settings and baseline models are detailed in Appendix: B.

Overall Performance (RQ1)

The results in Discriminative Recommendation. Table 1 reports offline AUC on the Kuaishou local-life dataset when DIN (Zhou et al. 2018), DIEN (Zhou et al. 2019), ETA (Chen et al. 2021), SIM (Pi et al. 2020), and TWIN (Si et al. 2024) are augmented with different item tokenization schemes. LGSID delivers the largest absolute gains. These improvements stem from the two components of LGSID. For attention-based DIN, DIEN and SIM, flat item IDs severely limit the granularity at which spatial proximity can influence attention scores. After the G-DPO phase of LGSID injects aligned LLM's spatial knowledge, each ID becomes a geography-aware embedding that encodes real-world spatial distance and neighborhood co-visit patterns. For TWIN and ETA, which operate under tight latency constraints, compact yet informative codes are critical. LGSID's hierarchical quantization first compresses geo-textual attributes into coarse primary tokens, then progressively refines residuals with geographic context. This yields richer representations without expanding the embedding table, lifting TWIN by 3.65% and ETA by 3.71%.

The results in Generative Recommendation. Table 2 shows the results of generative recommendation. We mainly compare two generative recommendation model, including TIGER (Rajput et al. 2023) and OneRec (Deng et al. 2025) with different quantization methods. We find that Lin et al. (Lin et al. 2025) obtains the worst performance, possibly because it uses different distance functions across various levels of codewords, leading to convergence difficulties and challenges in model optimization. RQ-VAE (Rajput et al. 2023) and RQ-VAE-ngram (Zheng et al. 2025b) achieve similar performance on TIGER and OneRec. However, these quantization methods did not consider geographical constraints, resulting in suboptimal performance in local-life recommendation scenarios. Our LGSID introduces RL-based Alignment to generate geographically aware representations and transfers these representations into semantic IDs through hierarchical geographic item tokenization.

RL-based LLM Alignment Analysis (RQ2)

This experiment validates the effectiveness of our RL-based LLM Alignment by measuring model improvements before and after fine-tuning. We design two metrics based on retrieving the top-k items by similarity to a target item embedding: (1) **semantic similarity**, measured by the semantic similarity of the retrieved items, and (2) **geographic awareness**, measured by the coverage of retrieved items sharing the same province (state), city, and town as the target item. Based on Table 3, we have following conclusion. (1) Pure semantic understanding is insufficient for geographic awareness, as text similarity captures terms but not real distances. (2) The reward model effectively compresses and transfers geographic knowledge into the LLM, where list-wise modeling improves T@5 from 0.1601 to 0.5584. (3) Incorpo-

Method	DIN	DIEN	SIM	TWIN	ETA
Origin	0.5859	0.6255	0.5884	0.5898	0.5903
+ Res-KMeans (Luo et al. 2024)	0.6100$_{\uparrow+0.0241}$	0.6369$_{\uparrow+0.0114}$	0.6063$_{\uparrow+0.0179}$	0.6087$_{\uparrow+0.0189}$	0.6077$_{\uparrow+0.0174}$
+ RQ-VAE (Rajput et al. 2023)	0.6185$_{\uparrow+0.0326}$	0.6364$_{\uparrow+0.0109}$	0.6111$_{\uparrow+0.0227}$	0.6153$_{\uparrow+0.0255}$	0.6153$_{\uparrow+0.0250}$
+ Lin et al. (Lin et al. 2025)	0.6161$_{\uparrow+0.0302}$	0.6368$_{\uparrow+0.0113}$	0.6107$_{\uparrow+0.0223}$	0.6148$_{\uparrow+0.0250}$	0.6148$_{\uparrow+0.0245}$
+ RQ-VAE-ngram (Zheng et al. 2025b)	0.6163$_{\uparrow+0.0304}$	0.6354$_{\uparrow+0.0099}$	0.6116$_{\uparrow+0.0232}$	0.6129$_{\uparrow+0.0231}$	0.6145$_{\uparrow+0.0242}$
+ LGSID (Ours)	**0.6276**$_{\uparrow+0.0417}$	**0.6484**$_{\uparrow+0.0229}$	**0.6224**$_{\uparrow+0.0340}$	**0.6263**$_{\uparrow+0.0365}$	**0.6274**$_{\uparrow+0.0371}$

Table 1: Performance comparison with AUC improvements over the original discriminative recommendation models.

Method	TIGER				OneRec			
	Hit@5	Hit@10	NDCG@5	NDCG@10	Hit@5	Hit@10	NDCG@5	NDCG@10
RQ-VAE (Rajput et al. 2023)	0.3087	0.3880	0.2255	0.2512	0.3739	0.4534	0.2798	0.3056
Lin et al. (Lin et al. 2025)	0.1767	0.2067	0.1335	0.1432	0.2950	0.3346	0.2272	0.2401
RQ-VAE-ngram (Zheng et al. 2025b)	0.2991	0.3769	0.2158	0.2411	0.3626	0.4358	0.2720	0.2957
LGSID (Ours)	**0.3921**	**0.5077**	**0.2817**	**0.3191**	**0.4435**	**0.5537**	**0.3304**	**0.3661**
IMP	**27.01%**	**30.83%**	**24.94%**	**27.05%**	**18.63%**	**22.13%**	**18.09%**	**19.79%**

Table 2: Performance comparison with different quantization methods in generative recommendation models. The **IMP** metric indicates the relative improvement of LGSID over the best-performing baseline (excluding LGSID).

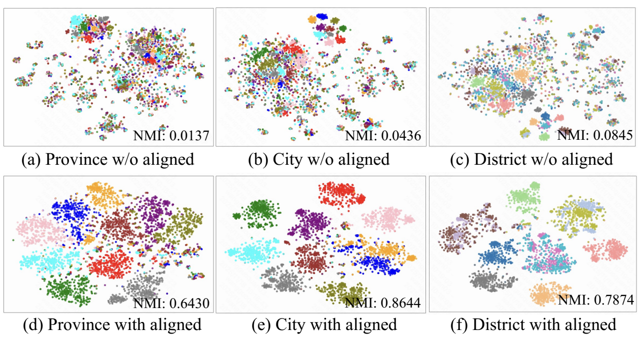

Figure 3: T-SNE visualization of items around cluster centroids across tokenization methods.

rating density-aware list-wise modeling further boosts T@5 from 0.1601 to 0.6114, enhancing near-distance sensitivity. However, this comes at the expense of semantic comprehension; applying a mixed-sample strategy mitigates this trade-off by improving sample discrimination and integrating collaborative signals. (4) Over emphasizing geographic awareness alone does not ensure better downstream performance, so we introduce textual similarity regularization to maintain semantics while achieving optimal results.

Visualization Analysis (RQ3)

The T-SNE visualization in Figure 3 shows that after RL-based Geographic LLM Alignment, the cluster centers of Province, City, and District are significantly closer. The Normalized Mutual Information (NMI) quantitatively measures the agreement between the clustering partition discovered by the model and the ground-truth geographic labels. It jumps from 0.0137–0.0845 to 0.6430–0.8644. The key to this improvement lies in: G-DPO first uses distance-aware list-wise rewards to inject real-world spatial relationships into the

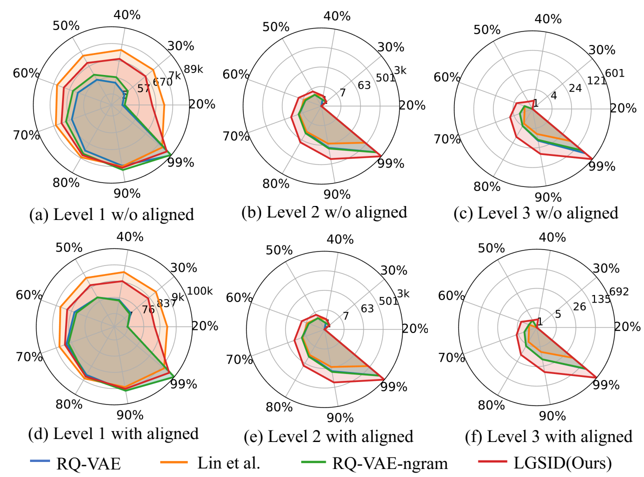

Figure 4: Token quantile percentiles across hierarchical levels for local-life items.

LLM, allowing its token embeddings to carry inherent geographic priors.

The radar chart in Figure 4 shows the quantile values at different percentile levels, where larger areas indicate superior coverage performance across the distribution spectrum. Higher quantile values indicate that tokens can represent more instances. At Level-1, our LGSID method demonstrates remarkable consistency between aligned and non-aligned settings, with identical coverage patterns observed in both scenarios. LGSID maintains 11k coverage in 90% quantile, while RQ-VAE decay to 8k. As we progress to finer granularities (Level-2 and Level-3), LGSID's advantages become more pronounced. Its area in the radar chart is the largest compared to other methods.

Method	Similarity			Province Coverage (P@K)			City Coverage (C@K)			Town Coverage (T@K)		
	Top@5	Top@10	Top@100	P@5	P@10	P@100	C@5	C@10	C@100	T@5	T@10	T@100
Origin	0.9204	0.9133	0.8833	0.8716	0.8410	0.6681	0.7342	0.6827	0.4372	0.1601	0.1328	0.0552
DPO-PR	0.8771	0.8679	0.8286	0.9001	0.8752	0.7477	0.7478	0.7013	0.5064	0.1452	0.1167	0.0445
DPO-LR	0.7595	0.7478	0.7088	0.8995	0.8648	0.6560	0.8681	0.8254	0.5783	0.5584	0.4966	0.2480
DPO-LRD	0.7411	0.7288	0.6876	0.8715	0.8302	0.6012	0.8277	0.7755	0.5043	0.6114	0.5435	0.2620
DPO-LRDM	0.8107	0.7954	0.7401	0.9047	0.8773	0.7261	0.7812	0.7329	0.5218	0.1816	0.1481	0.0625
DPO-LRDMS	0.8856	0.8754	0.8283	0.9960	0.9936	0.9662	0.9548	0.9352	0.8130	0.4030	0.3525	0.2260
G-DPO (Ours)	0.8977	0.8892	0.8504	0.9905	0.9852	0.9307	0.9173	0.8858	0.7065	0.294	0.2432	0.1290
IMP	**-2.47%**	**-2.64%**	**-3.72%**	**+13.64%**	**+17.15%**	**+39.31%**	**+24.94%**	**+29.75%**	**+61.60%**	**+83.64%**	**+83.13%**	**+133.70%**

Table 3: Evaluation of G-DPO variants: Origin (baseline), DPO-PR (point-wise reward), DPO-LR (list-wise reward), DPO-LRD (list-wise reward + density-aware sampling), DPO-LRDM (list-wise reward + density-aware sampling + domain mixed preference pairs), and DPO-LRDMS (list-wise reward + density-aware sampling + domain mixed preference pairs + similarity regularization). Metrics include similarity, province (P@K), city (C@K), and town (T@K) coverage for $K = \{5, 10, 100\}$.

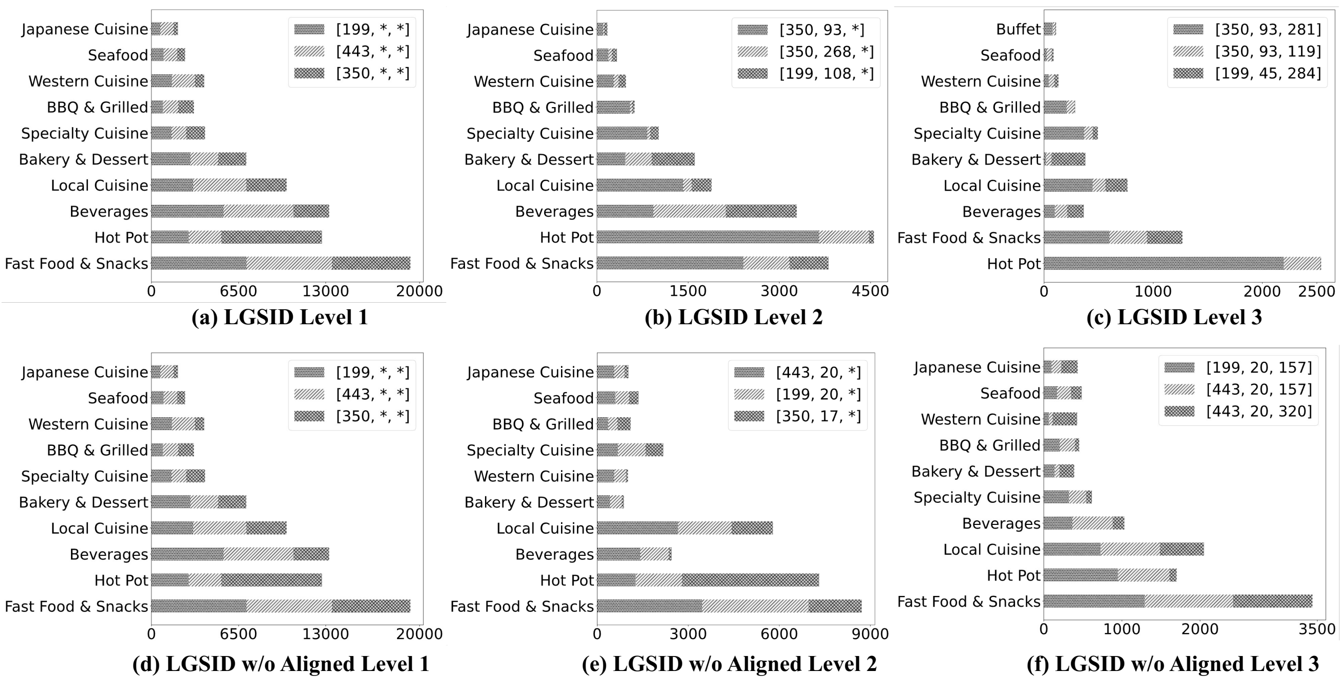

(a) LGSID Level 1 **(b) LGSID Level 2** **(c) LGSID Level 3**

(d) LGSID w/o Aligned Level 1 **(e) LGSID w/o Aligned Level 2** **(f) LGSID w/o Aligned Level 3**

Figure 5: Hierarchical category frequency distribution of LGSID for different SID prefixes (Aligned vs Unaligned).

Case Study (RQ4)

Figures 5 illustrate the allocation of three-layer discrete tokens produced by the hierarchical quantizer in our LGSID, with and without RL-based G-DPO alignment, respectively. Since the first layer relies on pre-computed geographic feature clusters, the category distributions in Figure 5(a) (alignment) and Figure 5(d) (misalignment) are similar. Figure 5(b) shows that, after RL-based G-DPO alignment, the first-layer token ([350, 93, *]) cleanly groups the entire BBQ & Grilled branch into a single coarse identifier. In contrast, Figure 5(e) depicts the same layer without alignment: the same restaurants are scattered across ([199, 20, *]), ([443, 20, *]), and ([350, 17, *]), because the LLM embeddings have not been post-trained with G-DPO to respect distance-aware rewards. As a result, the unaligned Level-1 tokens lose category cohesion; the hierarchical quantizer can no longer rely on a shared root to refine subcategories at later levels, highlighting the importance of upstream LLMs and the quality of their embeddings.

Conclusion

In this paper, LGSID is designed to equip semantic IDs with real-world spatial awareness for local-life recommendation. Specifically, LGSID trains a list-wise reward model with density-aware negative sampling to capture relative spatial distances, then injects geographic knowledge using a novel G-DPO algorithm. Moreover, a hierarchical geographic tokenization strategy generates a sequence of spatial-aware discrete tokens, enabling efficient compression and reconstruction. Extensive results on both discriminative and generative recommendation models demonstrate improved geographic awareness in scenarios of different granularity.

References

Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Chen, Q.; Pei, C.; Lv, S.; Li, C.; Ge, J.; and Ou, W. 2021. End-to-end user behavior retrieval in click-through rateprediction model. *arXiv preprint arXiv:2108.04468*.

Chen, R.; Ju, M.; Bui, N.; Antypas, D.; Cai, S.; Wu, X.; Neves, L.; Wang, Z.; Shah, N.; and Zhao, T. 2024. Enhancing item tokenization for generative recommendation through self-improvement. *arXiv preprint arXiv:2412.17171*.

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*.

He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, 173–182.

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.

Lin, G.; Hua, Z.; Feng, T.; Yang, S.; Long, B.; and You, J. 2025. Unified semantic and ID representation learning for deep recommenders. *arXiv preprint arXiv:2502.16474*.

Liu, E.; Zheng, B.; Ling, C.; Hu, L.; Li, H.; and Zhao, W. X. 2025a. Generative recommender with end-to-end learnable item tokenization. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 729–739.

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, Y.; Cao, J.; Wang, S.; Wen, S.; Chen, X.; Wu, X.; Yang, S.; Liu, Z.; Gai, K.; and Zhou, G. 2025c. LLM-Alignment Live-Streaming Recommendation. *arXiv preprint arXiv:2504.05217*.

Liu, Y.; Zhang, K.; Ren, X.; Huang, Y.; Jin, J.; Qin, Y.; Su, R.; Xu, R.; Yu, Y.; and Zhang, W. 2024. Alignrec: Aligning and training in multimodal recommendations. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 1503–1512.

Lu, W.; Lian, J.; Zhang, W.; Li, G.; Zhou, M.; Liao, H.; and Xie, X. 2024. Aligning large language models for controllable recommendations. *arXiv preprint arXiv:2403.05063*.

Luo, H.; Wu, B.; Jia, H.; Zhu, Q.; and Shan, L. 2025. Llm-cot enhanced graph neural recommendation with harmonized group policy optimization. *arXiv preprint arXiv:2505.12396*.

Luo, X.; Cao, J.; Sun, T.; Yu, J.; Huang, R.; Yuan, W.; Lin, H.; Zheng, Y.; Wang, S.; Hu, Q.; et al. 2024. Qarm: Quantitative alignment multi-modal recommendation at kuaishou. *arXiv preprint arXiv:2411.11739*.

Ma, G.; Wang, H.; Zhao, J.; Yan, R.; and Tang, H. 2024. Successive POI recommendation via brain-inspired spatiotemporal aware representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 574–582.

Pi, Q.; Zhou, G.; Zhang, Y.; Wang, Z.; Ren, L.; Fan, Y.; Zhu, X.; and Gai, K. 2020. Search-based user interest modeling with lifelong sequential behavior data for click-through rate prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, 2685–2692.

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.; Guan, L.; Sun, Z.; Zang, X.; Lu, J.; Hui, Y.; Cao, X.; Yang, Z.; Zheng, Y.; Leng, D.; et al. 2024. Twin v2: Scaling ultra-long user behavior sequence modeling for enhanced ctr prediction at kuaishou. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 4890–4897.

Singh, A.; Vu, T.; Mehta, N.; Keshavan, R.; Sathiamoorthy, M.; Zheng, Y.; Hong, L.; Heldt, L.; Wei, L.; Tandon, D.; et al. 2024. Better generalization with semantic ids: A case study in ranking for recommendations. In *Proceedings of the 18th ACM Conference on Recommender Systems*, 1039–1044.

Wang, D.; Huang, Y.; Gao, S.; Wang, Y.; Huang, C.; and Shang, S. 2025. Generative Next POI Recommendation with Semantic ID. *arXiv preprint arXiv:2506.01375*.

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, Y.; Chu, Z.; Ouyang, X.; Wang, S.; Hao, H.; Shen, Y.; Gu, J.; Xue, S.; Zhang, J.; Cui, Q.; et al. 2024b. Llmrg: Improving recommendations through large language model reasoning graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 19189–19196.

Wang, Y.; Xun, J.; Hong, M.; Zhu, J.; Jin, T.; Lin, W.; Li, H.; Li, L.; Xia, Y.; Zhao, Z.; et al. 2024c. 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.

Wu, J.; Wang, X.; Feng, F.; He, X.; Chen, L.; Lian, J.; and Xie, X. 2021. Self-supervised graph learning for recommendation. In *Proceedings of the 44th international ACM SIGIR* *conference on research and development in information retrieval*, 726–735.

Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Zhang, C.; Wu, S.; Zhang, H.; Xu, T.; Gao, Y.; Hu, Y.; and Chen, E. 2024a. Notellm: A retrievable large language model for note recommendation. In *Companion Proceedings of the ACM Web Conference 2024*, 170–179.

Zhang, C.; Zhang, H.; Wu, S.; Wu, D.; Xu, T.; Zhao, X.; Gao, Y.; Hu, Y.; and Chen, E. 2024b. Notellm-2: Multimodal large representation models for recommendation. *arXiv preprint arXiv:2405.16789*.

Zhang, S.; Chen, L.; Wang, C.; Li, S.; and Xiong, H. 2024c. Temporal graph contrastive learning for sequential recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 9359–9367.

Zhao, S.; Chen, W.; Shi, B.; Zhou, L.; Lin, S.; and Wan, H. 2025. Spatial-Temporal Knowledge Distillation for Takeaway Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 13365–13373.

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. 2025a. Universal Item Tokenization for Transferable Generative Recommendation. *arXiv preprint arXiv:2504.04405*.

Zheng, C.; Huang, M.; Pedchenko, D.; Rangadurai, K.; Wang, S.; Nahum, G.; Lei, J.; Yang, Y.; Liu, T.; Luo, Z.; et al. 2025b. Enhancing Embedding Representation Stability in Recommendation Systems with Semantic ID. *arXiv preprint arXiv:2504.02137*.

Zhou, G.; Mou, N.; Fan, Y.; Pi, Q.; Bian, W.; Zhou, C.; Zhu, X.; and Gai, K. 2019. Deep interest evolution network for click-through rate prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 5941–5948.

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.