

# BLCSREC: BRIDGING LANGUAGE AND COLLABORATIVE SEMANTICS FOR NEXT POI RECOMMENDATION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Next Point-of-Interest (POI) recommendation seeks to predict a user’s future location based on their past mobility patterns, a task essential for personalized location-based services. To enhance semantic modeling in this context, recent studies have applied large language models (LLMs), where early approaches relied on randomly assigned POI identifiers with limited representational capacity. More recent work has introduced semantic identifiers (SIDs) to better capture spatial and contextual correlations, leading to improved prediction accuracy. However, these approaches face several critical challenges: a semantic gap between LLM’s language semantics and POI-specific collaborative semantics, and the generation of invalid SIDs that do not correspond to real POIs. To tackle these challenges, we propose a novel framework called BLCSRec for bridging language and collaborative semantics in next POI recommendation. Specifically, we introduce an LLM-based POI profile generation method that summarizes user trajectories and integrates POI attributes with visitor information, and further employ an RQ-VAE to encode addresses, category, and these enriched textual profiles into semantic identifiers that capture both static attributes and collaborative context. To alleviate the gap between language and collaborative semantic, we incorporate explicit alignment tasks that map SIDs to/from textual descriptions and implicit alignment tasks that predict next POIs in asymmetric semantic formats. Furthermore, we employ GRPO reinforcement learning with a hierarchical reward structure to suppress invalid SID generation and enhance accuracy. Extensive experiments on three public datasets (NYC, TKY, and CA) demonstrate that our method consistently outperforms strong LLM-based baselines, achieving improvements of 7.3% in Acc@1 on NYC and 3.5% on CA, along with substantial reductions in invalid SID generation.

## 1 INTRODUCTION

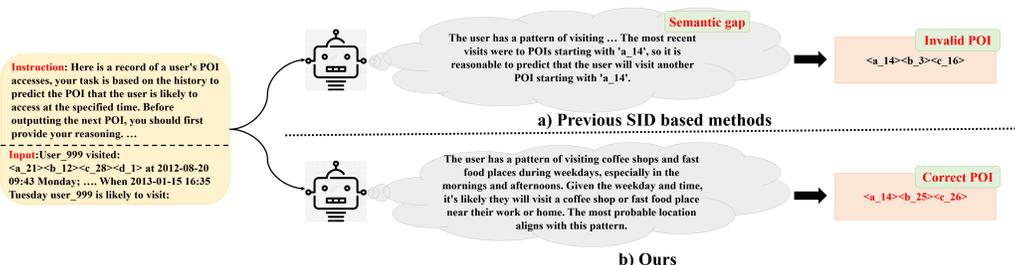


Figure 1: Challenge and Case Study

The proliferation of Location-Based Social Networks (LBSNs) has generated an unprecedented volume of user mobility data, typically captured as chronological sequences of check-ins at various Points-of-Interest (POIs). Capitalizing on these rich datasets, the task of next POI recommendation (Zhang et al., 2025a; Yu et al., 2024) has emerged as a pivotal research direction. This task is

054 fundamentally concerned with predicting a user’s subsequent destination by modeling the intricate  
055 patterns embedded within their historical movement trajectories.

056  
057 Early approaches(Feng et al., 2015; Graves, 2012; Wu et al., 2022; Luo et al., 2021; Yang et al.,  
058 2022; Yan et al., 2023; Luo et al., 2023; Feng et al., 2024) to next POI recommendation have  
059 predominantly utilized classic deep learning models, such as Long Short-Term Memory (LSTM)  
060 networks(Hochreiter & Schmidhuber, 1997) and Graph Neural Networks (GNNs)(Scarselli et al.,  
061 2009), to capture users’ temporal or spatial preferences. Following the demonstrated success of  
062 large language models (LLMs) in various recommendation tasks (Li et al., 2024b), recent studies  
063 have extended their use to next POI recommendation. While some approaches(Li et al., 2024a) con-  
064 struct prompts using randomly assigned POI identifiers, this setting limits the model’s ability to cap-  
065 ture semantic relations among locations. In contrast, prior work in other recommendation domains  
066 has demonstrated that using semantic identifiers (SIDs), where items with similar characteristics are  
067 assigned more similar SIDs, allows LLMs to better capture the relationships and similarities among  
068 items(Zhou et al., 2025; Zheng et al., 2024a; Rajput et al., 2023). Motivated by these findings, sev-  
069 eral recent efforts(Wang et al., 2025) have introduced SIDs into next POI recommendation, enabling  
LLMs to better exploit spatial and contextual correlations and thus improving predictive accuracy.

070 Despite their promising performance in next POI recommendation, approaches utilizing LLMs with  
071 SIDs are constrained by two primary limitations, as shown in Figure 1. a) **Semantic gap**: there exists  
072 a large gap between the language semantics modeled by LLMs and collaborative semantics implied  
073 by recommender systems. LLMs process natural language through token IDs, whereas recommen-  
074 dation systems utilize POI-specific identifiers (in this case, SIDs). This fundamental misalignment  
075 hinders the ability to fully harness the LLM’s representational power for recommendation tasks,  
076 as the models effectively learn within disparate semantic spaces. As shown in Figure 1a), when  
077 generating SIDs, the LLM fails to capture the semantics of index tokens and instead infers user  
078 preferences solely from their visit frequency. b) **Invalid POIs**: the inherent constraints of LLMs,  
079 specifically their finite context windows and imperfect instruction adherence, render it infeasible  
080 to inform the model of all valid SIDs through enumeration. Consequently, during unconstrained  
081 generation, the LLM is prone to producing spurious SIDs that do not correspond to any actual POI.

082 To address these challenges, we propose Bridging Language and Collaborative Semantics for Next  
083 POI Recommendation (BLCSRec). The framework first leverages an LLM to construct user profiles  
084 that capture individual preferences from check-in records, and POI profiles that summarize visit-  
085 ing patterns based on user preferences and temporal context. These profiles enrich the semantic  
086 space for SID generation, thereby enhancing the discriminability of POI representations and also  
087 provide LLM-comprehensible inputs and outputs, paving the way for an explicit alignment task.  
088 Each POI is then assigned a SID by RQ-VAE(Lee et al., 2022) using its address, category, and  
089 profile. The sequence of SIDs derived from a user’s visited POIs forms the training data for the  
090 LLM-based Sequential POI Prediction task. To mitigate the semantic gap, BLCSRec introduces two  
091 complementary strategies: (i) Explicit Alignment, which requires the LLM to map between SIDs  
092 and textual descriptions during training, grounding tokens in semantic meaning; and (ii) Implicit  
093 Alignment, which predicts the next POI’s SID or text from sequences of check-in records, encour-  
094 aging the model to exploit SID–semantic relations in sequential prediction. The three datasets are  
095 jointly used to train the LLM. Moreover, since reinforcement learning has been shown to increase  
096 the likelihood of generating preference-consistent samples while reducing the probability of non-  
097 preferred outputs(Yue et al., 2025; Xiong et al., 2025), we design a hierarchical reward structure that  
098 assigns higher rewards to correct SIDs, lower rewards to valid but incorrect SIDs, and the lowest  
099 rewards to invalid SIDs. GRPO(Shao et al., 2024) is then employed to suppress invalid generations  
and improve overall prediction accuracy. Our main contributions can be summarized as follows:

- 100 • We find that existing SID-based approaches employing LLMs for next POI recommenda-  
101 tion mainly suffer from two issues: a semantic gap between tokenized representations and  
102 POI-specific identifiers, and a non-negligible risk of generating invalid SIDs.
- 103 • We construct POI profiles to support SID generation and introduce both explicit and im-  
104 plicit alignment tasks, which alleviate the semantic gap by linking SIDs with their semantic  
105 representations. In addition, GRPO is incorporated to reduce the probability of producing  
106 invalid SIDs and to further improve prediction accuracy.
- 107 • We evaluate BLCSRec against strong baselines on three benchmark datasets (NYC, TKY,  
and CA), achieving consistent improvements in Acc@1 (e.g., +7.3% on NYC and +3.5%

on CA over the best baseline) while substantially reducing the generation of invalid SIDs, which demonstrates the effectiveness of our method.

## 2 RELATED WORK

Early efforts in next POI recommendation typically formulated the problem as a sequential prediction task, often relying on handcrafted models such as Markov chains(He et al., 2016) or factorization techniques(Feng et al., 2015). With the advancement of deep learning, recurrent neural networks(RNNs) became the mainstream, where variants were designed to capture spatial-temporal dynamics and model both short- and long-term user preferences(Graves, 2012; Wu et al., 2022). More recently, attention-based architectures have shown superior performance by providing greater flexibility in modeling dependencies. Transformer-based models(Luo et al., 2021; Yang et al., 2022; Luo et al., 2023; Feng et al., 2024), such as GETNext(Yang et al., 2022), integrates a global trajectory flow graph with Transformer to jointly model collaborative, spatial-temporal, and semantic contexts. Parallel to this line of work, graph neural networks(GNNs) have been adopted to capture user-POI relationships beyond individual sequences(Zhang et al., 2021; Yan et al., 2023).

Following the demonstrated superior performance of LLMs in various recommendation tasks(Li et al., 2024b), recent research(Wang et al., 2023; Liu et al., 2024; Li et al., 2024a; Wang et al., 2025) has begun to explore their application to next POI recommendation. LLM4POI(Li et al., 2024a) involves constructing a prompt from a user’s historical sequence of visited POI identifiers and tasking the LLM with generating the ID for the next recommended POI. Recognizing that many high-performing LLM-based recommenders rely on semantic IDs (SIDs)(Zhou et al., 2025; Zheng et al., 2024a; Rajput et al., 2023), GNPR-SID(Wang et al., 2025) introduces SIDs by assigning a unique semantic identifier to each POI, ensuring that similar POIs possess more proximate SID representations. By constructing the input sequence with these SIDs, the LLM can more effectively capture the latent correlations between POIs, including spatial proximity and other contextual attributes, ultimately leading to improved recommendation accuracy.

Although existing LLM-based methods leveraging SIDs have achieved state-of-the-art performance in next POI recommendation, they overlook two critical issues. First, there remains a semantic gap between the language-level semantics captured by LLMs and the collaborative semantics inherent in recommender systems. Second, LLMs can occasionally generate invalid SIDs that do not correspond to any real POI, which may degrade recommendation reliability. To address these issues, we introduce a set of alignment tasks to mitigate the semantic gap and employ GRPO(Shao et al., 2024) to reduce the likelihood of generating invalid SIDs.

## 3 METHOD

To address the two primary challenges in existing SID-based next POI recommendation approaches, we propose tailored solutions. For the semantic gap, we first derive user preferences and temporal patterns by analyzing historical visitation records and timestamps for both users and POIs, then introduce a series of supplementary semantic alignment tasks during the LLM’s Supervised Fine-Tuning (SFT) process. These tasks encompass explicit alignment, which directs the LLM to generate corresponding SIDs/textual descriptions from POI textual descriptions/SIDs, and implicit alignment, which involves predicting subsequent POI SIDs or textual descriptions given input sequences with textual descriptions or SIDs. For invalid SID generation, we implement an additional Reinforcement Learning phase after SFT using the GRPO training methodology to encourage the model to output accurate and existing POI SIDs.

Our methodology comprises three primary phases: semantic ID construction, language-collaborative POI SID alignment, and reinforcement learning, as illustrated in Figure 2.

### 3.1 SEMANTIC ID CONSTRUCTION

Similar to approaches(Li et al., 2025; Deng et al., 2025; Zheng et al., 2024b; Yang et al., 2025b) in previous recommender systems, we eschew raw POI IDs in favor of vector quantization techniques, representing each POI with a compact set of discrete indices. While these indices are traditionally constructed using textual information, POIs in next-POI recommendation tasks typically have

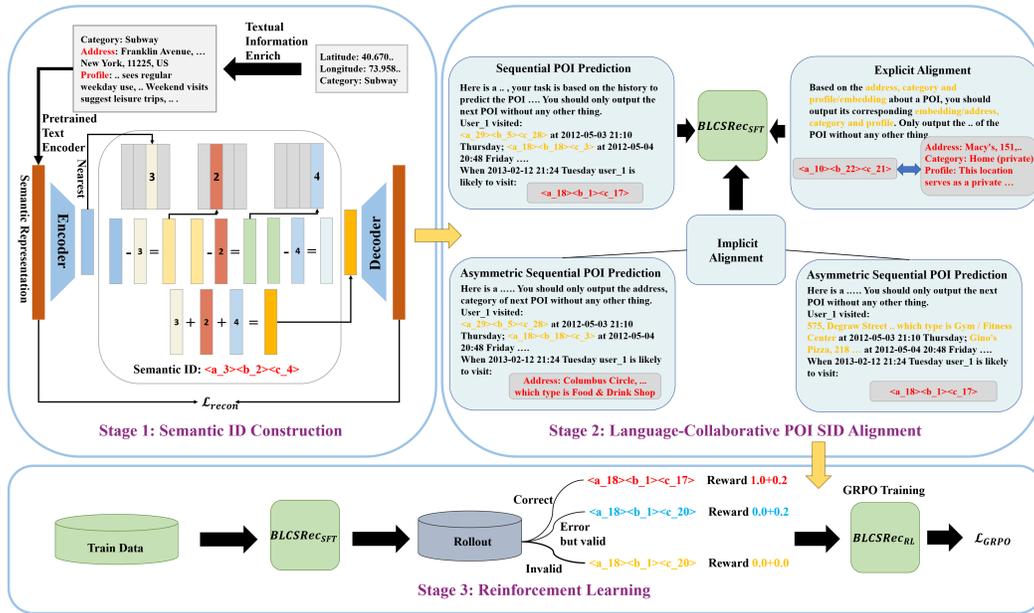


Figure 2: The overall architecture of our proposed framework.

limited text data, often consisting only of geographical coordinates and category information. To generate distinctive and semantically rich semantic IDs for each POI, we first construct semantic representation for each POI and subsequently employed RQ-VAE(Lee et al., 2022) for semantic ID construction.

### 3.1.1 SEMANTIC REPRESENTATION CONSTRUCTION

To construct semantically rich text representations, we enhance each POI’s textual description beyond its existing categorical information. Specifically, we enrich the spatial features and developed comprehensive POI profiles that capture both collaborative and temporal patterns.

**Spatial Features** For next-POI recommendation tasks, the conventional use of precise latitude-longitude coordinates poses challenges for LLMs due to their limitation in processing high-precision numerical inputs(Yang et al., 2025a). Although prior work(Wang et al., 2025) addressed this through Plus Code<sup>1</sup> discretization, this approach suffers from location ambiguity within the same code region. To solve these problems, We proposes converting coordinate pairs into U.S.-style address strings (number + street + city + state) via standard mapping services<sup>2</sup>. This representation maintains spatial granularity through shared address components among proximate locations while ensuring distinctiveness for POIs, thereby enabling LLMs to effectively capture spatial relationships without relying on precise numerical values.

**POI Profiles** Inspired by previous approaches(Liu et al., 2025) that generate item profiles from interaction patterns when basic item information is limited, we develop a two-stage LLM-based profiling process. First, we generate comprehensive user profiles capturing travel preferences by feeding each user’s historical POI interaction sequences (including locations, categories, and timestamps) into an LLM to summarize their categorical preferences, temporal patterns, and geographical trends. Subsequently, for each POI, we enrich its textual information by prompting the LLM to synthesize its collaborative and temporal patterns based on its location, category, and visitor information (including user profiles and visit timestamps). This approach enables the construction of semantically-rich textual representations that effectively capture both explicit and implicit POI characteristics.

<sup>1</sup><https://maps.google.com/pluscodes/>

<sup>2</sup><https://nominatim.org/>

Finally, we concatenate each POI’s categorical information, spatial features, and generated profile into a comprehensive textual description  $text$ , which serves as input to a pre-trained text encoder  $E_t$  to obtain the POI’s semantic representation  $p_s$ .

### 3.1.2 POI QUANTIZATION

After obtaining the POI semantic representations, we employ vector quantization to generate discrete SID for each POI. Specifically, we utilize a RQ-VAE(Lee et al., 2022) with the semantic representations as input to generate POI semantic IDs  $SID$ . RQ-VAE implements a multi-stage vector quantization process that recursively quantizes residual vectors from coarse to fine levels to produce a set of index tokens. For each semantic representation  $p_s$ , RQ-VAE first encodes it into a latent representation  $z$ . At each level  $l$ , we maintain a codebook  $\mathcal{C}^l = \{v_k^l\}_{k=1}^K$ , where each codebook vector  $v_k^l$  represents a learnable cluster center and  $K$  is the number of cluster centers. The residual quantization process can be formulated as:

$$c_i = \arg \min_k \|e_i - v_k^i\| \quad (1)$$

$$e_{i+1} = e_i - v_{c_i}^i \quad (2)$$

$$SID = [c_1, c_2, \dots, c_L] \quad (3)$$

where  $c_i$  is the  $i$ -th index token of the semantic ID,  $L$  denotes the number of quantization levels, and  $e_i$  is the residual vector in the  $i$ -th RQ level, and we set  $e_1 = z$ .

After quantizing the semantic representation into a SID, we obtain the corresponding latent representation  $\hat{z}$  by dequantizing each level’s index using the associated codebook vectors according to  $\hat{z} = \sum_{i=1}^L v_{c_i}^i$ . This latent representation is then passed through a decoder to reconstruct the semantic representation  $\hat{p}_s$ . The overall loss function of RQ-VAE is defined as follows.

$$\mathcal{L}_{recon} = \|p_s - \hat{p}_s\|_2^2 \quad (4)$$

$$\mathcal{L}_{RQ} = \sum_{i=1}^L \|sg[e_i] - v_{c_i}^i\|_2^2 + \beta \|e_i - sg[v_{c_i}^i]\|_2^2 \quad (5)$$

$$\mathcal{L}_{RQ-VAE} = \mathcal{L}_{recon} + \mathcal{L}_{RQ} \quad (6)$$

The loss function consists of two components:  $L_{recon}$ , which measures the reconstruction error, and  $L_{RQ}$ , which encourages the residual vectors to align closely with their assigned codebook entries. The stop-gradient operator  $sg[\cdot]$  is applied to prevent gradient flow through discrete variables, and the coefficient  $\beta$ , typically set to 0.25.

It is worth noting that the trained RQ-VAE may produce a non-negligible rate of conflicts, where multiple POIs are assigned the same SID. To ensure that each POI is associated with a distinct identifier, we append a unique suffix to the shared semantic ID. For instance, if two POIs are mapped to the same ID sequence  $\langle a.1 \rangle \langle b.1 \rangle \langle c.1 \rangle$ , they are further distinguished as  $\langle a.1 \rangle \langle b.1 \rangle \langle c.1 \rangle \langle d.0 \rangle$  and  $\langle a.1 \rangle \langle b.1 \rangle \langle c.1 \rangle \langle d.1 \rangle$ , respectively.

## 3.2 LANGUAGE-COLLABORATIVE POI SID ALIGNMENT

Although the constructed SIDs allow POIs with similar meanings to share similar prefix index tokens, which improves their accessibility to LLMs, the models themselves lack awareness of the semantic content encoded in index tokens such as  $\langle a.1 \rangle$ . To bridge this gap, we design a series of semantic alignment tasks to incorporate both language and collaborative semantics into the LLMs through fine-tuning. These tasks include: (1) sequential POI prediction, (2) explicit alignment, and (3) implicit alignment.

### 3.2.1 SEQUENTIAL POI PREDICTION

Our method leverages LLM for next-POI prediction as the primary training objective. Specifically, we construct personalized recommendation instruction based on users’ historical check-in patterns. These instruction, along with check-in histories, serve as input to the LLM to predict the POI a

270 target user will likely visit at a specified future time. The user’s check-in history is represented as a  
271 chronologically ordered sequence of POI semantic IDs.

272 Notably, the temporal aspects of historical check-in records are significant in two distinct dimen-  
273 sions: the timestamp information captures user-specific behavioral patterns at a granular level, while  
274 the day-of-week features encode cyclic behavioral variations in weekly routines (exemplified by stu-  
275 dents’ distinct patterns of visiting academic institutions on weekdays versus preferring residential or  
276 recreational venues on weekends). Consequently, we concatenate the chronologically ordered POI  
277 SIDs with their corresponding check-in timestamps and day-of-week indicators. This integrated  
278 approach enables more comprehensive modeling of personalized user information, effectively syn-  
279 thesizing collaborative semantics across both spatial and temporal dimensions.

280 However, merely tuning LLM through this approach fails to ensure adequate perception of feature  
281 information embedded within each index token of the SIDs. To address this limitation, we incor-  
282 porate both explicit and implicit alignment tasks to bridge the semantic gap between language and  
283 collaborative semantics in LLMs.

### 284 285 286 3.2.2 EXPLICIT ALIGNMENT 287

288 While our POI SIDs are constructed from POI addresses, categories, and profiles, they establish only  
289 weak correlations among POIs with similar linguistic semantics through shared prefix index tokens.  
290 This results in LLM’s limited ability to learn semantic information from these weak associations. To  
291 enhance LLM’s semantic understanding of each index token, we propose two explicit index token-  
292 language alignment tasks for LLM tuning.

293 Our alignment objectives are twofold: On the one hand, we aim to enable LLM to interpret a POI’s  
294 address, category, and profile when presented with its semantic ID. On the other hand, we seek  
295 to develop LLM’s capability to construct the appropriate semantic ID when given a POI’s address,  
296 category, and profile information. We implement this bidirectional alignment through two sequential  
297 tasks: first instructing LLM to generate the corresponding address, category, and profile from a POI’s  
298 semantic ID, then directing it to reconstruct the semantic ID based on these attributes.

### 300 301 3.2.3 IMPLICIT ALIGNMENT 302

303 While explicit alignment tasks enable LLM to perceive the mapping relationship between POI tex-  
304 tual information and semantic IDs, this understanding alone may be insufficient for practical ap-  
305 plication, analogous to how memorizing mathematical formulas doesn’t guarantee problem-solving  
306 ability(Shao et al., 2024). In the recommendation context, LLM might struggle to effectively utilize  
307 these learned mappings during actual recommendation tasks. Therefore, we designed two asymmet-  
308 ric POI prediction tasks as implicit alignment tasks specifically tailored to recommendation scenar-  
309 ios.

310 In sequential item prediction, both check-in histories and target POI are formatted using semantic  
311 ID representations, constituting a symmetric task where both conditions and targets are semantically  
312 based. To enhance semantic alignment further, we increase prediction complexity by varying the  
313 forms of conditions and targets, yielding different combinations of POI semantic representations.  
314 Specifically, we implement two asymmetric representations: First, replacing target POI semantic  
315 IDs with corresponding addresses and categories, directing LLM to generate POI attributes from  
316 semantic ID sequences; Second, substituting check-in history’s POI semantic IDs with addresses  
317 and categories, guiding LLM to generate expected POI semantic ID.

### 318 319 3.2.4 SUPERVISED FINE-TUNING 320

321 Following dataset construction for all three tasks, we integrated them for LLM’s SFT. Following con-  
322 ventional SFT practices for decoder-only architectures, we optimize the model parameters through  
323 token-level cross-entropy loss, maximizing the conditional probability of subsequent tokens given  
their preceding context, thereby obtaining our fine-tuned model  $\text{BLCSRec}_{\text{SFT}}$ .

Table 1: Comparison of different models on three datasets: NYC, TKY, and CA. We present the model inputs including the POI representation approach (POI Repr.) and whether visit timestamps are utilized.

Model	Inputs		$Acc@1$			$rate_i$		
	POI Repr.	Seq	NYC	TKY	CA	NYC	TKY	CA
PRME	RID	-	0.1159	0.1052	0.0521	-	-	-
LSTM	RID	-	0.1305	0.1335	0.0665	-	-	-
PLSPL	RID	-	0.1917	0.1889	0.1072	-	-	-
STAN	RID	-	0.2231	0.1963	0.1104	-	-	-
GETNext	RID	-	0.2435	0.1829	0.1357	-	-	-
STHGCN	RID	-	0.2734	0.2950	0.1730	-	-	-
TPG	RID	-	0.2555	0.1420	0.1749	-	-	-
ROTAN	RID	-	0.3106	0.2458	0.2199	-	-	-
LLM4POI	RID	300	0.3372	<u>0.3035</u>	0.2065	-	-	-
GNPR-SID	SID	50	0.3310	0.2801	0.2119	1.42%	1.31%	1.63%
<b>BLCSR<sub>rec</sub><sup>SFT</sup></b>	SID	50	<u>0.3589</u>	0.2979	<u>0.2239</u>	<u>0.72%</u>	<u>0.40%</u>	<u>1.40%</u>
<b>BLCSR<sub>rec</sub><sup>RL</sup></b>	SID	50	<b>0.3618</b>	<b>0.3049</b>	<b>0.2277</b>	<b>0.00%</b>	<b>0.05%</b>	<b>1.12%</b>

### 3.3 REINFORCEMENT LEARNING

As GRPO has been demonstrated to steer model outputs toward higher-reward outcomes (Yue et al., 2025), in the reinforcement learning stage, we utilize the GRPO algorithm to enhance answer accuracy and reduce the likelihood of the model generating invalid SIDs.

In our RL stage, we incorporate two reward signals: the validity reward and the accuracy reward. The validity reward encourages the model to generate semantic IDs corresponding to real-world POIs; a score of 0.2 is assigned if the predicted semantic ID maps to an existing POI, and 0 otherwise. The accuracy reward promotes precise prediction by assigning a score of 1 when the predicted semantic ID  $\hat{y}_i$  matches the ground truth  $y_i$ , and 0 otherwise.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETTINGS

**Datasets** We evaluate our approach on three real-world datasets: Foursquare NYC (Yang et al., 2015), Foursquare TKY (Yang et al., 2015), and Gowalla CA (Cho et al., 2011). Following previous works (Yan et al., 2023), we preprocess each dataset by filtering out POIs and users with fewer than 10 check-ins. The data is chronologically sorted, with an 8-1-1 split for training, validation, and testing. To maintain consistency between training and evaluation, we exclude users and POIs from the test set that do not appear in the training data. Check-in records are then grouped by users and arranged chronologically. During evaluation, each user’s last visited POI serves as ground truth, while previous visits constitute the input sequence. Notably, during training, users’ historical records are treated as sequential input and concatenated with the test set to a specified length before inference.

**Evaluation Metrics** Consistent with previous LLM-based methods (Li et al., 2024a; Wang et al., 2025), our model generates a single POI recommendation per inference. Therefore, we adopt the accuracy metric for top-1 recommendation  $acc@1$ . Additionally, for LLM-based methods utilizing semantic IDs as input and output, we introduce an invalid SID rate metric  $rate_i$  to assess the probability of generating invalid SIDs.

### 4.2 MAIN RESULTS

We evaluate our approach against current state-of-the-art methods across the NYC, TKY, and CA datasets, with results presented in Table 1. Our method **BLCSR<sub>rec</sub><sup>RL</sup>** demonstrates significant

improvements in  $Acc@1$  metrics, achieving gains of 7.3%, 0.5%, and 3.5% respectively over the best existing approaches. Notably, even our SFT model **BLCSRECSFT** surpasses current state-of-the-art performance on NYC and CA datasets. On the TKY dataset, our method achieves comparable performance to LLM4POI(Li et al., 2024a) using only 50 check-in records, versus LLM4POI’s 300 records. Compared to method GNPR-SID(Wang et al., 2025), which also employs semantic IDs, both our SFT and RL models exhibit superior performance in both accuracy and invalid rate metrics, demonstrating enhanced effectiveness and robustness. We attribute these improvements to three key factors:

- **Enhanced POI Representation:** Our approach generates comprehensive POI profiles incorporating visit preferences and temporal characteristics, enriching textual information and enabling more distinctive semantic IDs with reduced conflicts.
- **Semantic Alignment:** The introduction of alignment tasks during SFT bridges the semantic gap between natural language and semantic IDs, enabling the LLM to better capture the semantic content of each index token, simultaneously improving accuracy while reducing invalid SIDs occurrences.
- **RL Optimization:** The post-SFT reinforcement learning phase effectively increases the probability of generating correct semantic IDs while reducing invalid SIDs predictions.

### 4.3 ABLATION STUDY

The ablation study was conducted on the NYC and CA datasets to examine the effect of each component in our framework. Evaluation was based on  $Acc@1$  and  $rate_i$ . Beginning with the pretrained Qwen2.5-7B-Instruct(Yang et al., 2024) model without POI profiles, we incrementally added POI profile incorporation(profile), sequential POI prediction fine-tuning(SP), explicit(EA) and implicit alignment(IA) objectives, and finally GRPO training(GRPO).

The results in Table 2 show steady gains in accuracy and reductions in invalid SIDs generation. On NYC,  $Acc@1$  improves from 0.2601 (baseline) to 0.3618 with GRPO, while  $rate_i$  drops from 4.35% to 0.00%. A similar trend is observed on CA, where  $Acc@1$  rises from 0.1695 to 0.2277 and  $rate_i$  decreases from 1.85% to 1.12%. These findings indicate that each component contributes to accuracy improvement, with explicit and implicit alignment yielding notable gains, and GRPO providing the most effective balance between accuracy and invalid SIDs reduction.

Method	NYC		CA	
	$Acc@1$	$rate_i$	$Acc@1$	$rate_i$
Qwen2.5-7B	0.2601	4.35%	0.1695	1.85%
+ profile	0.2901	4.02%	0.1718	1.72%
+ SP	0.3357	1.30%	0.2151	1.62%
+ EA	0.3546	1.12%	0.2179	1.59%
+ IA	0.3589	0.72%	0.2239	1.41%
+ GRPO	0.3618	0.00%	0.2277	1.12%

Table 2: Ablation Study.

### 4.4 SEMANTIC ANALYSIS

To evaluate whether our alignment task effectively language and collaborative semantics in LLM, we conducted a semantic analysis on the NYC dataset from two perspectives: LLM token embedding visualization analysis and LLM performance on semantically similar negative items.

#### 4.4.1 EMBEDDING VISUALIZATION ANALYSIS

We investigate the relationship between POI index tokens and the original semantic space of the LLM. Following prior studies(Zheng et al., 2024b), we apply Principal Component Analysis (PCA)(Yang et al., 2004) to visualize embeddings of different token types. In Figure 3a), “Tokens w/ Alignment” refers to all SID index tokens produced by the LLM when the alignment task is incorporated, whereas “Tokens w/o Alignment” denotes all SID index tokens derived from the LLM trained without the alignment task during SFT, and “POI Texts” refers to tokens derived from POI addresses, categories and profiles. The 2D visualization indicates that, after incorporating the alignment task, the distribution of SID index tokens in the semantic space becomes closer to the

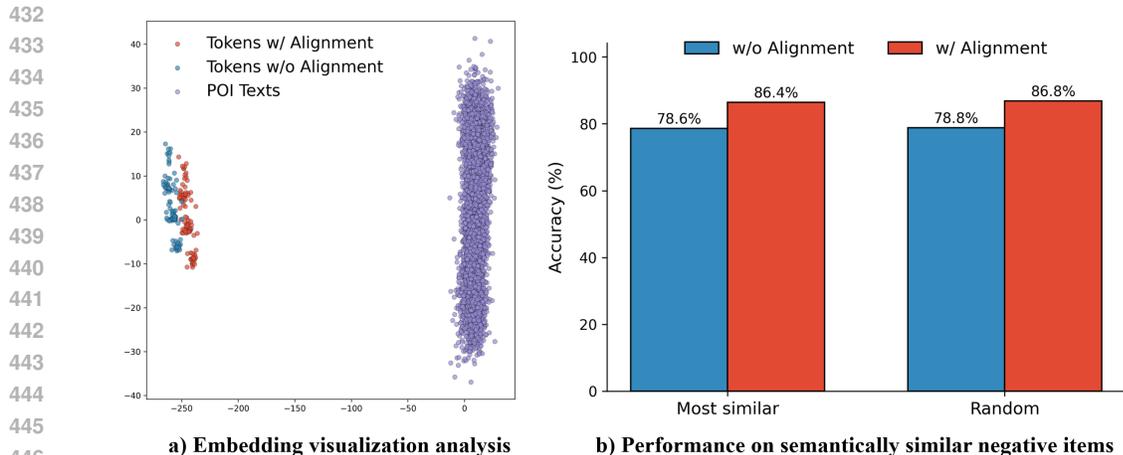


Figure 3: Semantic Analysis

LLM semantic space of POI textual information compared to the setting without alignment. This observation indicates that the alignment task can effectively reduce the semantic gap between LLM language and collaborative semantics.

#### 4.4.2 PERFORMANCE ON SEMANTICALLY SIMILAR NEGATIVE ITEMS

To examine why combining language and collaborative semantics benefits LLMs in next POI recommendation, we design an auxiliary ranking task with negative samples. Specifically, two types of negatives are constructed: (1) most similar negatives, identified by cosine similarity between POI text embeddings, and (2) random negatives, sampled uniformly from the valid POI set. The evaluation reuses the same test set as the sequential POI prediction task, where the model must distinguish the ground-truth item from a negative, the performance is reported in terms of accuracy.

The results showed in Figure 3b) indicate that, after incorporating the alignment tasks, the LLM outperforms its non-aligned counterpart in both distinguishing POIs with similar language semantics and identifying random negative POIs. This demonstrates that our proposed explicit and implicit alignment tasks effectively reduce the semantic gap between recommendation and natural language tasks, enabling the model to capture semantic information within index tokens and thereby improve recommendation performance.

## 5 CONCLUSION

In this work, we address two critical issues in LLM-based next POI recommendation systems related to SID integration: the semantic gap between language semantics modeled by LLMs and collaborative semantics in POI identifiers, and the generation of invalid SIDs that undermine prediction reliability. To tackle these challenges, we propose BLCSRec, which leverages LLMs to construct user and POI profiles capturing preferences, temporal patterns, and collaborative signals, thereby enriching SID generation via RQ-VAE for more discriminative representations. Furthermore, we introduce explicit alignment tasks to ground SIDs in textual semantics through bidirectional mapping, and implicit alignment tasks to encourage the model to exploit SID-semantic relations in asymmetric sequential prediction. Additionally, we employ GRPO reinforcement learning with a hierarchical reward mechanism that prioritizes correct SIDs, penalizes invalid ones. Experimental results on three public datasets (NYC, TKY, and CA) show that our approach consistently outperforms strong baselines, with improvements of 7.3% in Acc@1 on NYC and 3.5% on CA, while substantially reducing invalid SID generation. These findings underscore the importance of semantic alignment and reinforcement strategies for enhancing the robustness and accuracy of LLM-based recommender systems, offering insights for broader applications in sequential recommendation tasks involving spatial-temporal data.

## REFERENCES

- 486  
487  
488 Eunjoon Cho, Seth A. Myers, and Jure Leskovec. Friendship and mobility: user movement in  
489 location-based social networks. In Chid Apté, Joydeep Ghosh, and Padhraic Smyth (eds.), *Pro-*  
490 *ceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and*  
491 *Data Mining, San Diego, CA, USA, August 21-24, 2011*, pp. 1082–1090. ACM, 2011. doi:  
492 10.1145/2020408.2020579. URL <https://doi.org/10.1145/2020408.2020579>.
- 493 DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu,  
494 Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu,  
495 Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao  
496 Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan,  
497 Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao,  
498 Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding,  
499 Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang  
500 Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai  
501 Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang,  
502 Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang,  
503 Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang,  
504 Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang,  
505 R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye,  
506 Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, and S. S. Li. Deepseek-r1: Incentivizing  
507 reasoning capability in llms via reinforcement learning. *CoRR*, abs/2501.12948, 2025. doi: 10.  
508 48550/ARXIV.2501.12948. URL <https://doi.org/10.48550/arXiv.2501.12948>.
- 509 Jiaxin Deng, Shiyao Wang, Kuo Cai, Lejian Ren, Qigen Hu, Weifeng Ding, Qiang Luo, and Guorui  
510 Zhou. Onerec: Unifying retrieve and rank with generative recommender and iterative preference  
511 alignment. *CoRR*, abs/2502.18965, 2025. doi: 10.48550/ARXIV.2502.18965. URL <https://doi.org/10.48550/arXiv.2502.18965>.
- 512 Shanshan Feng, Xutao Li, Yifeng Zeng, Gao Cong, Yeow Meng Chee, and Quan Yuan. Personalized  
513 ranking metric embedding for next new POI recommendation. In Qiang Yang and Michael J.  
514 Wooldridge (eds.), *Proceedings of the Twenty-Fourth International Joint Conference on Artificial*  
515 *Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pp. 2069–2075. AAAI Press,  
516 2015. URL <http://ijcai.org/Abstract/15/293>.
- 517 Shanshan Feng, Feiyu Meng, Lisi Chen, Shuo Shang, and Yew Soon Ong. ROTAN: A rotation-  
518 based temporal attention network for time-specific next POI recommendation. In Ricardo Baeza-  
519 Yates and Francesco Bonchi (eds.), *Proceedings of the 30th ACM SIGKDD Conference on Knowl-*  
520 *edge Discovery and Data Mining, KDD 2024, Barcelona, Spain, August 25-29, 2024*, pp. 759–  
521 770. ACM, 2024. doi: 10.1145/3637528.3671809. URL <https://doi.org/10.1145/3637528.3671809>.
- 522 Alex Graves. *Supervised Sequence Labelling with Recurrent Neural Networks*, volume 385 of *Stud-*  
523 *ies in Computational Intelligence*. Springer, 2012. ISBN 978-3-642-24796-5. doi: 10.1007/  
524 978-3-642-24797-2. URL <https://doi.org/10.1007/978-3-642-24797-2>.
- 525 Jing He, Xin Li, Lejian Liao, Dandan Song, and William K. Cheung. Inferring a personalized next  
526 point-of-interest recommendation model with latent behavior patterns. In Dale Schuurmans and  
527 Michael P. Wellman (eds.), *Proceedings of the Thirtieth AAAI Conference on Artificial Intelli-*  
528 *gence, February 12-17, 2016, Phoenix, Arizona, USA*, pp. 137–143. AAAI Press, 2016. doi:  
529 10.1609/AAAI.V30I1.9994. URL <https://doi.org/10.1609/aaai.v30i1.9994>.
- 530 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):  
531 1735–1780, 1997. doi: 10.1162/NECO.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- 532 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
533 and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth Inter-*  
534 *national Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.  
535 OpenReview.net, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.

- 540 Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image  
541 generation using residual quantization. In *IEEE/CVF Conference on Computer Vision and  
542 Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 11513–11522.  
543 IEEE, 2022. doi: 10.1109/CVPR52688.2022.01123. URL [https://doi.org/10.1109/  
544 CVPR52688.2022.01123](https://doi.org/10.1109/CVPR52688.2022.01123).
- 545 Guanghan Li, Xun Zhang, Yufei Zhang, Yifan Yin, Guojun Yin, and Wei Lin. Semantic conver-  
546 gence: Harmonizing recommender systems via two-stage alignment and behavioral semantic  
547 tokenization. In Toby Walsh, Julie Shah, and Zico Kolter (eds.), *AAAI-25, Sponsored by the  
548 Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadel-  
549 phia, PA, USA*, pp. 12040–12048. AAAI Press, 2025. doi: 10.1609/AAAI.V39I11.33311. URL  
550 <https://doi.org/10.1609/aaai.v39i11.33311>.
- 551 Peibo Li, Maarten de Rijke, Hao Xue, Shuang Ao, Yang Song, and Flora D. Salim. Large language  
552 models for next point-of-interest recommendation. In Grace Hui Yang, Hongning Wang, Sam  
553 Han, Claudia Hauff, Guido Zuccon, and Yi Zhang (eds.), *Proceedings of the 47th International  
554 ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024,  
555 Washington DC, USA, July 14-18, 2024*, pp. 1463–1472. ACM, 2024a. doi: 10.1145/3626772.  
556 3657840. URL <https://doi.org/10.1145/3626772.3657840>.
- 557 Yongqi Li, Xinyu Lin, Wenjie Wang, Fuli Feng, Liang Pang, Wenjie Li, Liqiang Nie, Xiangnan  
558 He, and Tat-Seng Chua. A survey of generative search and recommendation in the era of large  
559 language models. *CoRR*, abs/2404.16924, 2024b. doi: 10.48550/ARXIV.2404.16924. URL  
560 <https://doi.org/10.48550/arXiv.2404.16924>.
- 561 Shuai Liu, Ning Cao, Yile Chen, Yue Jiang, and Gao Cong. nextlocllm: next location prediction  
562 using llms. *CoRR*, abs/2410.09129, 2024. doi: 10.48550/ARXIV.2410.09129. URL <https://doi.org/10.48550/arXiv.2410.09129>.
- 563 Yibin Liu, Ang Li, and Shijian Li. Hierarchical interaction summarization and contrastive prompting  
564 for explainable recommendations, 2025. URL <https://arxiv.org/abs/2507.06044>.
- 565 Yan Luo, Haoyi Duan, Ye Liu, and Fu-Lai Chung. Timestamps as prompts for geography-aware  
566 location recommendation. In Ingo Frommholz, Frank Hopfgartner, Mark Lee, Michael Oakes,  
567 Mounia Lalmas, Min Zhang, and Rodrygo L. T. Santos (eds.), *Proceedings of the 32nd ACM In-  
568 ternational Conference on Information and Knowledge Management, CIKM 2023, Birmingham,  
569 United Kingdom, October 21-25, 2023*, pp. 1697–1706. ACM, 2023. doi: 10.1145/3583780.  
570 3615083. URL <https://doi.org/10.1145/3583780.3615083>.
- 571 Yingtao Luo, Qiang Liu, and Zhaocheng Liu. STAN: spatio-temporal attention network for next  
572 location recommendation. In Jure Leskovec, Marko Grobelnik, Marc Najork, Jie Tang, and Leila  
573 Zia (eds.), *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-  
574 23, 2021*, pp. 2177–2185. ACM / IW3C2, 2021. doi: 10.1145/3442381.3449998. URL <https://doi.org/10.1145/3442381.3449998>.
- 575 Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz  
576 Heldt, Lichan Hong, Yi Tay, Vinh Q. Tran, Jonah Samost, Maciej Kula, Ed H. Chi, and Ma-  
577 hesh Sathiamoorthy. Recommender systems with generative retrieval. In Alice Oh, Tris-  
578 tan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Ad-  
579 vances in Neural Information Processing Systems 36: Annual Conference on Neural Infor-  
580 mation Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16,  
581 2023*, 2023. URL [http://papers.nips.cc/paper\\_files/paper/2023/hash/  
582 20dcab0f14046a5c6b02b61da9f13229-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2023/hash/20dcab0f14046a5c6b02b61da9f13229-Abstract-Conference.html).
- 583 Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini.  
584 The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.  
585 doi: 10.1109/TNN.2008.2005605.
- 586 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, Y. K. Li,  
587 Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open  
588 language models. *CoRR*, abs/2402.03300, 2024. doi: 10.48550/ARXIV.2402.03300. URL  
589 <https://doi.org/10.48550/arXiv.2402.03300>.
- 590
- 591
- 592
- 593

- 594 Dongsheng Wang, Yuxi Huang, Shen Gao, Yifan Wang, Chengrui Huang, and Shuo Shang. Generative next POI recommendation with semantic ID. *CoRR*, abs/2506.01375, 2025. doi: 10.48550/ARXIV.2506.01375. URL <https://doi.org/10.48550/arXiv.2506.01375>.
- 595  
596  
597
- 598 Xinglei Wang, Meng Fang, Zichao Zeng, and Tao Cheng. Where would I go next? large language models as human mobility predictors. *CoRR*, abs/2308.15197, 2023. doi: 10.48550/ARXIV.2308.15197. URL <https://doi.org/10.48550/arXiv.2308.15197>.
- 599  
600
- 601 Yuxia Wu, Ke Li, Guoshuai Zhao, and Xueming Qian. Personalized long- and short-term preference learning for next POI recommendation. *IEEE Trans. Knowl. Data Eng.*, 34(4):1944–1957, 2022. doi: 10.1109/TKDE.2020.3002531. URL <https://doi.org/10.1109/TKDE.2020.3002531>.
- 602  
603  
604
- 605
- 606 Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong Zhang, Caiming Xiong, and Hanze Dong. A minimalist approach to LLM reasoning: from rejection sampling to reinforce. *CoRR*, abs/2504.11343, 2025. doi: 10.48550/ARXIV.2504.11343. URL <https://doi.org/10.48550/arXiv.2504.11343>.
- 607  
608  
609
- 610 Xiaodong Yan, Tengwei Song, Yifeng Jiao, Jianshan He, Jiaotuan Wang, Ruopeng Li, and Wei Chu. Spatio-temporal hypergraph learning for next POI recommendation. In Hsin-Hsi Chen, Wei-Jou (Edward) Duh, Hen-Hsen Huang, Makoto P. Kato, Josiane Mothe, and Barbara Poblete (eds.), *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*, pp. 403–412. ACM, 2023. doi: 10.1145/3539618.3591770. URL <https://doi.org/10.1145/3539618.3591770>.
- 611  
612  
613  
614  
615  
616
- 617 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report. *CoRR*, abs/2412.15115, 2024. doi: 10.48550/ARXIV.2412.15115. URL <https://doi.org/10.48550/arXiv.2412.15115>.
- 618  
619  
620  
621  
622  
623  
624
- 625 Dingqi Yang, Daqing Zhang, Vincent W. Zheng, and Zhiyong Yu. Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns. *IEEE Trans. Syst. Man Cybern. Syst.*, 45(1):129–142, 2015. doi: 10.1109/TSMC.2014.2327053. URL <https://doi.org/10.1109/TSMC.2014.2327053>.
- 626  
627  
628  
629
- 630 Haotong Yang, Yi Hu, Shijia Kang, Zhouchen Lin, and Muhan Zhang. Number cookbook: Number understanding of language models and how to improve it. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025a. URL <https://openreview.net/forum?id=BWS5gVjgeY>.
- 631  
632  
633
- 634 Jian Yang, David Zhang, Alejandro F. Frangi, and Jing-Yu Yang. Two-dimensional PCA: A new approach to appearance-based face representation and recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(1):131–137, 2004. doi: 10.1109/TPAMI.2004.10004. URL <http://doi.ieeeecomputersociety.org/10.1109/TPAMI.2004.10004>.
- 635  
636  
637  
638
- 639 Song Yang, Jiamou Liu, and Kaiqi Zhao. Getnext: Trajectory flow map enhanced transformer for next POI recommendation. In Enrique Amigó, Pablo Castells, Julio Gonzalo, Ben Carterette, J. Shane Culpepper, and Gabriella Kazai (eds.), *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, pp. 1144–1153. ACM, 2022. doi: 10.1145/3477495.3531983. URL <https://doi.org/10.1145/3477495.3531983>.
- 640  
641  
642  
643  
644
- 645 Yuhao Yang, Zhi Ji, Zhaopeng Li, Yi Li, Zhonglin Mo, Yue Ding, Kai Chen, Zijian Zhang, Jie Li, Shuanglong Li, and Lin Liu. Sparse meets dense: Unified generative recommendations with cascaded sparse-dense representations. *CoRR*, abs/2503.02453, 2025b. doi: 10.48550/ARXIV.2503.02453. URL <https://doi.org/10.48550/arXiv.2503.02453>.
- 646  
647

- 648 Jian Yu, Lucas Guo, Jiayu Zhang, and Guiling Wang. A survey on graph neural network-  
649 based next POI recommendation for smart cities. *J. Reliab. Intell. Environ.*, 10(3):299–  
650 318, 2024. doi: 10.1007/S40860-024-00233-Z. URL [https://doi.org/10.1007/  
651 s40860-024-00233-z](https://doi.org/10.1007/s40860-024-00233-z).
- 652 Yang Yue, Zhiqi Chen, Rui Lu, Andrew Zhao, Zhaokai Wang, Yang Yue, Shiji Song, and Gao  
653 Huang. Does reinforcement learning really incentivize reasoning capacity in llms beyond the  
654 base model? *CoRR*, abs/2504.13837, 2025. doi: 10.48550/ARXIV.2504.13837. URL <https://doi.org/10.48550/arXiv.2504.13837>.
- 655  
656  
657 Jiyong Zhang, Xin Liu, Xiaofei Zhou, and Xiaowen Chu. Leveraging graph neural networks for  
658 point-of-interest recommendations. *Neurocomputing*, 462:1–13, 2021. doi: 10.1016/J.NEUCOM.  
659 2021.07.063. URL <https://doi.org/10.1016/j.neucom.2021.07.063>.
- 660  
661 Qianru Zhang, Peng Yang, Junliang Yu, Haixin Wang, Xingwei He, Siu-Ming Yiu, and Hongzhi  
662 Yin. A survey on point-of-interest recommendation: Models, architectures, and security. *IEEE*  
663 *Trans. Knowl. Data Eng.*, 37(6):3153–3172, 2025a. doi: 10.1109/TKDE.2025.3551292. URL  
664 <https://doi.org/10.1109/TKDE.2025.3551292>.
- 665  
666  
667 Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie,  
668 An Yang, Dayiheng Liu, Junyang Lin, Fei Huang, and Jingren Zhou. Qwen3 embedding: Advanc-  
669 ing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*,  
670 2025b.
- 671  
672 Yuze Zhao, Jintao Huang, Jinghan Hu, Xingjun Wang, Yunlin Mao, Daoze Zhang, Zeyinzi Jiang,  
673 Zhikai Wu, Baole Ai, Ang Wang, Wenmeng Zhou, and Yingda Chen. Swift:a scalable lightweight  
674 infrastructure for fine-tuning, 2024. URL <https://arxiv.org/abs/2408.05517>.
- 675  
676 Bowen Zheng, Yupeng Hou, Hongyu Lu, Yu Chen, Wayne Xin Zhao, Ming Chen, and Ji-Rong Wen.  
677 Adapting large language models by integrating collaborative semantics for recommendation. In  
678 *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, pp. 1435–1448, 2024a.  
679 doi: 10.1109/ICDE60146.2024.00118.
- 680  
681 Bowen Zheng, Yupeng Hou, Hongyu Lu, Yu Chen, Wayne Xin Zhao, Ming Chen, and Ji-Rong Wen.  
682 Adapting large language models by integrating collaborative semantics for recommendation. In  
683 *40th IEEE International Conference on Data Engineering, ICDE 2024, Utrecht, The Netherlands,*  
684 *May 13-16, 2024*, pp. 1435–1448. IEEE, 2024b. doi: 10.1109/ICDE60146.2024.00118. URL  
685 <https://doi.org/10.1109/ICDE60146.2024.00118>.
- 686  
687 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, Zhangchi Feng, and  
688 Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In *Pro-*  
689 *ceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume*  
690 *3: System Demonstrations)*, Bangkok, Thailand, 2024c. Association for Computational Linguis-  
691 tics. URL <http://arxiv.org/abs/2403.13372>.
- 692  
693 Guorui Zhou, Jiabin Deng, Jinghao Zhang, Kuo Cai, Lejian Ren, Qiang Luo, Qianqian Wang, Qigen  
694 Hu, Rui Huang, Shiyao Wang, Weifeng Ding, Wuchao Li, Xinchun Luo, Xingmei Wang, Zexuan  
695 Cheng, Zixing Zhang, Bin Zhang, Boxuan Wang, Chaoyi Ma, Chengru Song, Chenhui Wang,  
696 Di Wang, Dongxue Meng, Fan Yang, Fangyu Zhang, Feng Jiang, Fuxing Zhang, Gang Wang,  
697 Guowang Zhang, Han Li, Hengrui Hu, Hezheng Lin, Hongtao Cheng, Hongyang Cao, Huanjie  
698 Wang, Jiaming Huang, Jiapeng Chen, Jiaqiang Liu, Jinghui Jia, Kun Gai, Lantao Hu, Liang Zeng,  
699 Liao Yu, Qiang Wang, Qidong Zhou, Shengzhe Wang, Shihui He, Shuang Yang, Shujie Yang, Sui  
700 Huang, Tao Wu, Tiantian He, Tingting Gao, Wei Yuan, Xiao Liang, Xiaoxiao Xu, Xugang Liu,  
701 Yan Wang, Yi Wang, Yiwu Liu, Yue Song, Yufei Zhang, Yunfan Wu, Yunfeng Zhao, and Zhanyu  
Liu. Onerec technical report. *CoRR*, abs/2506.13695, 2025. doi: 10.48550/ARXIV.2506.13695.  
URL <https://doi.org/10.48550/arXiv.2506.13695>.

## A APPENDIX

### A.1 GENAI USAGE DISCLOSURE

LLMs play a central role in our work. First, we employ LLMs to generate profiles for users and POIs. Second, since our approach follows a generative recommendation paradigm, the underlying generative model we adopt is also an LLM. In addition, we utilize LLMs to refine the writing of this paper. Beyond these uses, no generative models are involved in any other part of our work, including the codebase, figures, tables, or experimental data.

### A.2 BASELINES

Our experimental evaluation encompasses diverse baseline approaches across five categories: Traditional Sequential Models including PRME(Feng et al., 2015), a metric embedding approach for personalized ranking, LSTM(Graves, 2012), a recurrent neural architecture for sequence modeling, and PLSPL(Wu et al., 2022), a personalized sequential recommendation framework; Transformer-based Architectures comprising STAN(Luo et al., 2021), which incorporates spatio-temporal attention mechanisms, and GETNext(Yang et al., 2022), leveraging transformer encoders for POI prediction; Temporal-aware Frameworks including TPG(Luo et al., 2023), employing timestamp-guided prediction strategies, and ROTAN(Feng et al., 2024), integrating temporal awareness in POI recommendations; Graph-based Approaches represented by STHGCN(Yan et al., 2023), utilizing hierarchical graph convolutions for spatio-temporal modeling; and LLM-based Methods: LLM4POI(Li et al., 2024a), the current state-of-the-art approach that utilizes random numeric identifiers and comprehensive user check-in histories; GNPR-SID(Wang et al., 2025), which leverages semantic IDs constructed from discretized POI features and processes user check-in sequences through these semantic representations.

### A.3 IMPLEMENTATION DETAILS

During semantic ID construction, we employ GPT-41-mini<sup>3</sup> for generating user and POI profiles, followed by Qwen3-Embedding-0.6B(Zhang et al., 2025b) for text representation generation. The RQ-VAE(Lee et al., 2022) module incorporates 3 codebook layers, each comprising 32 index tokens for NYC and 64 for TKY and CA datasets, with a vector dimensionality of 64. For the base model, we select Qwen2.5-7B-Instruct(Yang et al., 2024), maintaining comparable parameter scale with previous works(Li et al., 2024a; Wang et al., 2025). During the SFT stage, we apply LoRA(Hu et al., 2022) fine-tuning to all linear layers, with dataset-specific rank settings of 128 for NYC and 1024 for both TKY and CA. Training is conducted using the LLaMA-Factory(Zheng et al., 2024c) framework on 4 RTX 4090 GPUs, with a constant learning rate of 1e-5, a batch size of 1 per GPU, gradient accumulation steps set to 4, and a sequence length of 4096 tokens. In the RL stage, we fine-tune the linear layers within the LLM’s Attention modules using LoRA, employing a learning rate of 1e-6 and lora rank of 128. The training data used for the sequential POI prediction task is reused as training data for GRPO. We initialize GRPO training from model **BLCSR**<sub>recSFT</sub>. This stage is implemented using the ms-swift(Zhao et al., 2024) framework across 8 RTX 4090 GPUs, with a batch size of 1 per GPU, gradient accumulation steps set to 4,  $G$  set to 16 and a sequence length of 4096 tokens. To address potential convergence issues arising from the underrepresentation of our primary sequential POI prediction task, we augmented its presence by triplicating the corresponding data. To facilitate the comprehension of LLMs, we first translate the Japanese addresses in the TKY dataset, obtained from mapping services, into English using translation software<sup>4</sup> before further utilization.

### A.4 TRAINING DATA EXAMPLES

This section illustrates the training samples for the three tasks in our Language-Collaborative POI SID Alignment stage. An example for the sequential POI prediction task is presented as follows:

<sup>3</sup><https://openai.com/index/gpt-4-1/>

<sup>4</sup><https://translate.google.com>

**Instruction:** Here is a record of a user’s POI accesses, your task is based on the history to predict the POI that the user is likely to access at the specified time. You should only output the next POI without any other thing.

**Input:** User\_3 visited: <a\_6><b\_7><c\_14> at 2012-04-07 23:09 Saturday; <a\_24><b\_16><c\_31><d\_2> at 2012-04-08 02:13 Sunday; .... When 2012-09-15 16:58 Saturday user\_3 is likely to visit:

**Output:** <a\_5><b\_0><c\_20>

The training data for explicit alignment task is demonstrated as follows:

**Instruction:** Based on the embedding about a POI, you should output its corresponding address, category and profile. Only output the address, category and profile of the POI without any other thing.

**Input:** <a\_5><b\_0><c\_20>

**Output:** Address: John F. Kennedy International Airport, ...

Category: Airport

Profile: ...

**Instruction:** Based on the address, category and profile about a POI, you should output its corresponding embedding. Only output the embedding of the POI without any other thing.

**Input:** Address: John F. Kennedy International Airport, ...

Category: Airport

Profile: ...

**Output:** <a\_5><b\_0><c\_20>

The training data for implicit alignment task is demonstrated as follows:

**Instruction:** Here is a record of a user’s POI accesses, your task is based on the history to predict the POI that the user is likely to access at the specified time. You should only output the address, category of next POI without any other thing

**Input:** User\_3 visited: <a\_6><b\_7><c\_14> at 2012-04-07 23:09 Saturday; <a\_24><b\_16><c\_31><d\_2> at 2012-04-08 02:13 Sunday; .... When 2012-09-15 16:58 Saturday user\_3 is likely to visit:

**Output:** John F. Kennedy International Airport, JFK Access Road, Queens, Queens County, City of New York, New York, 11430, United States which type is Airport

**Instruction:** Here is a record of a user’s POI accesses, your task is based on the history to predict the POI that the user is likely to access at the specified time. You should only output the next POI without any other thing.

**Input:** User\_3 visited: 115, Court Street, ... which type is Movie Theater at 2012-04-07 23:09 Wednesday; Shake Shack, 409, Fulton Mall, ... which type is Burger Joint at 2012-04-08 02:13 Wednesday; .... When 2012-09-15 16:58 Sunday user\_3 is likely to visit:

**Output:** <a\_3><b\_7><c\_44>

## A.5 GRPO

During the training of GRPO, for each training iteration, given the input  $q$ , we sample  $G$  candidate outputs from the previous policy  $\pi_{\text{old}}$ . Each candidate receives a reward  $r$ , and we compute the group-relative advantage  $\hat{A}_i$  for candidate  $i$  as

$$\hat{A}_i = \frac{r_i - \mu}{\sigma} \quad (7)$$

where  $\mu$  and  $\sigma$  denote the mean and standard deviation of rewards within the candidate group. Outputs with above-average rewards thus gain higher advantages, guiding prioritized optimization. The policy update then maximizes the following clipped surrogate objective:

$$\mathcal{L}_{GRPO}(\theta) = \mathbb{E}_{o_i \sim \pi_\theta} \left[ \frac{1}{G} \sum_{i=1}^G \min(\rho_i \hat{A}_i, \text{clip}(\rho_i, 1-\epsilon, 1+\epsilon) \hat{A}_i) \right] - \beta \text{KL}[\pi_\theta \parallel \pi_{\text{old}}], \quad (8)$$

where  $\pi_\theta$  is the new policy,  $\rho_i = \frac{\pi_\theta(o_i|q)}{\pi_{\text{old}}(o_i|q)}$  is the importance-sampling ratio,  $\epsilon$  is the clipping hyperparameter, and the KL divergence term with weight  $\beta$  constrains policy deviation for stability which is calculated as  $\text{KL}[\pi_\theta \parallel \pi_{\text{old}}] = \frac{\pi_{\text{old}}}{\pi_\theta} - \log \frac{\pi_{\text{old}}}{\pi_\theta} - 1$ .

#### A.6 SUPERVISED FINE-TUNING LOSS

The SFT loss is as follows:

$$\mathcal{L}_{sft} = \sum_{t=1}^T p_\theta(x_t | x_1, x_2, \dots, x_{t-1}) \quad (9)$$

Where  $T$  denotes the length of the output sequence,  $x_t$  represents the token being predicted,  $x_1, x_2, \dots, x_{t-1}$  refers to the preceding tokens in the sequence,  $\theta$  denotes the trainable model parameters, and  $p_\theta(x_t | x_1, x_2, \dots, x_{t-1})$  represents the predicted probability distribution of the model.

#### A.7 EVALUATION METRICS DEFINITION

$acc@1$  is defined as:

$$acc@1 = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i = y_i) \quad (10)$$

where  $N$  denotes the number of test samples,  $\mathbb{I}(\cdot)$  represents the indicator function that yields 1 when the condition is satisfied and 0 otherwise, and  $\hat{y}_i, y_i$  denote the model’s prediction and ground truth, respectively. And  $rate_i$  is formulated as:

$$rate_i = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i \notin V) \quad (11)$$

where  $V$  denotes the set of valid semantic IDs.

#### A.8 CASE STUDY

In Figure 1b), we further illustrate the reasoning process of our GRPO-enhanced model prior to predicting the next POI. The example shows that the model can effectively associate information from a sequence of SIDs with their corresponding categories, addresses, and related attributes, enabling more accurate next-POI predictions. Compared with previous SID-based approaches, our method allows the model to capture the semantic content embedded in index tokens, thereby improving recommendation accuracy.

#### A.9 IMPACT OF HYPERPARAMETERS

Table 3: Impact of validity reward when the predicted semantic ID is valid

reward value	$Acc@1$		$rate_i$	
	NYC	CA	NYC	CA
0.1	0.3604	0.2267	0.58%	1.26%
0.2	0.3618	0.2277	0.00%	1.12%
0.3	0.3589	0.2258	0.00%	1.07%
0.4	0.3575	0.2247	0.00%	0.98%
0.5	0.3502	0.2239	0.00%	0.98%

Inspired by the reward design in DeepSeek-R1(DeepSeek-AI et al., 2025), we assign an accuracy reward of 1.0 when the predicted SID is correct and 0.0 when it is incorrect. Beyond this, we

864 investigate the effect of introducing a validity reward when the predicted SID is valid. The results,  
865 shown in Table 3, indicate that small validity rewards (0.1–0.2) slightly improve Acc@1 in both  
866 NYC and CA datasets compared to higher reward values. For instance, setting the reward to 0.2  
867 yields the best performance (0.3618 in NYC and 0.2277 in CA). However, larger rewards (larger  
868 than 0.3) lead to a gradual decline in accuracy, suggesting that excessive emphasis on validity may  
869 interfere with accuracy optimization. Meanwhile, the invalid prediction rate decreases as the reward  
870 increases, with values above 0.3 driving the rate to nearly 0% in NYC and below 1% in CA. This  
871 demonstrates a trade-off between improving prediction validity and maintaining overall accuracy.

872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917