Task-Aware Meta-Learning on Heterogeneous Knowledge Graph for POI Recommendation

Jingyuan Wang[1], **Zhichun Wang***[1,2], **Tong Lu**[1],**Yiming Guan**[1]

[1] School of Artificial Intelligence, Beijing Normal University, Beijing, China
[2] Engineering Research Center of Intelligent Technology and Educational Application, Ministry of Education, Beijing, China
202321081040@mail.bnu.edu.cn, zcwang@bnu.edu.cn

Abstract

Point-of-Interest (POI) recommendation plays a pivotal role in location-based services by guiding users to discover new and relevant places. While graph-based methods have shown promising results, effectively modeling the diversity and dynamics of user preferences remains a key challenge. Addressing this requires richer representations of both POIs and user interests, as well as more adaptive learning strategies. In this work, we propose TMHKG, a **T**ask-aware **M**eta-learning framework with a **H**eterogeneous **K**nowledge **G**raph for POI recommendation. To enhance representation learning, TMHKG constructs a dual-view POI knowledge graph that integrates geographical proximity and user-aware category transitions, and models users' evolving interests from sequential visit histories. On top of enriched features, TMHKG adopts a task-aware meta-learning paradigm, treating each user's recommendation task as a separate meta-task. A generalizable recommendation policy is first learned from diverse training tasks and then quickly adapted to each user's unique behavior, enabling highly personalized predictions. Extensive experiments on two real-world datasets demonstrate that TMHKG consistently outperforms state-of-the-art baselines, highlighting its effectiveness in capturing complex user-POI interactions.

Code — https://github.com/kg-bnu/TMHKG

Introduction

With the widespread adoption of smart devices and the proliferation of location-based services, user interaction data has undergone a significant transformation, evolving from simple sequential records into intricate and complex graphs. This evolution underscores the increasing demand for advanced Point-of-Interest (POI) recommendation systems, which play a crucial role in helping individuals discover new places and personalize their experiences in a dynamic environment.

To meet this demand, Graph Neural Networks (GNNs) have emerged as a powerful paradigm, prized for their inherent ability to model complex structural dependencies within user-POI interaction graphs (Shu et al. 2022; Zhang and Ma 2024; Song et al. 2025). While these models have proven effective, their common reliance on uniform aggregation rules can pose challenges in capturing the full spectrum of individual user preferences. Such an approach, while powerful for general patterns, may not fully adapt to the unique ways different users weigh temporal and categorical factors in their decision-making. This observation suggests that the challenge may lie not only in the model architecture but also in the underlying training paradigm, which typically learns a single, static set of parameters for all users. This has led to a growing interest in meta-learning for its capacity to facilitate rapid adaptation. While its application has often been centered on addressing the cold-start problem (Mandal et al. 2022), we see a broader opportunity. By focusing primarily on new users, the full potential of meta-learning to personalize for the vast diversity across all user behaviors may not be fully realized (Li et al. 2025).

To explore this opportunity, we advocate for a paradigm shift and reframe personalized recommendation as a meta-learning problem where each user is treated as a unique task. To this end, we propose a novel framework, Task-aware Meta-learning on Heterogeneous Knowledge Graph (TMHKG), designed to learn a highly adaptive model that can quickly personalize to diverse and dynamic user preferences. Instead of learning a single static model, TMHKG learns a highly generalizable meta-model and rapidly adapts it to each user's specific behavioral patterns. Our main contributions are summarized as follows:

- We propose TMHKG, a novel framework that reframes personalized recommendation as an adaptive learning problem to address user diversity. To our knowledge, this is one of the first efforts to systematically apply a task-aware meta-learning paradigm to knowledge graph-enhanced recommendation.
- We design a dual-graph learning architecture that synergistically models user behavior. It captures dynamic preferences via an attribute-enriched interaction graph and enriches POI semantics using a knowledge graph enhanced with novel category-transition and spatial-proximity relations.
- We introduce a memory-aware task sampling strategy to effectively train the model on diverse user behaviors, and through extensive experiments on two real-world

*Corresponding author.

datasets, we demonstrate that TMHKG significantly outperforms state-of-the-art methods in personalized recommendation.

Related Work

Graph-based POI Recommendation

GNNs are central to modern POI recommendation for modeling structural dependencies. To better encode spatial and temporal contexts, many studies enhance interaction graphs with geographic and temporal attributes. STGN (Zhao et al. 2020) employs temporal gating mechanisms to filter sequential context during propagation. STP-UDGAT (Lim et al. 2020) further combines spatial and temporal information as attention inputs, allowing fine-grained modeling of user preference evolution. Substantial efforts have been made to improve GNN expressiveness and robustness. STGCN (Han et al. 2020) integrates spatio-temporal gating with GCNs for joint modeling of sequential trajectories, while SGL (Procopio, Tripodi, and Navigli 2021) introduces contrastive learning to strengthen representation robustness. Hybrid models that combine sequence modeling and graph aggregation have also shown effectiveness. For example, STAMP (Liu et al. 2018) incorporates short-term sequences into graph-based preferences, while GETNext (Yang, Liu, and Zhao 2022) combines graph contraction with global context learning via a Transformer to capture long-range dependencies. Despite their progress, these methods generally rely on uniform aggregation rules. This "one-size-fits-all" approach inherently limits their flexibility in modeling the nuanced preferences of individual users, which motivates our work to explore a more adaptive paradigm.

Meta-learning based POI Recommendation

Meta-learning offers a powerful paradigm for building adaptive recommendation models. In the context of POI recommendation, its application has frequently centered on addressing the cold-start problem. For example, MFNP (Sun et al. 2021) incorporates region-specific user preferences through meta-learning-based adaptive fusion, and MetaKG (Du et al. 2022) introduces meta-learners to alleviate the cold-start issue. More recently, HyperMAN (Wang et al. 2025) combines hypergraph modeling with meta-learning to handle cold-start situations. While these methods have shown promise, their primary focus on cold-start scenarios means the broader challenge of modeling the substantial behavioral diversity across all users remains less explored. This focus can limit their ability to effectively capture the unique preferences of individual users, regardless of their activity level.

Methodology

We now present the proposed TMHKG framework, which integrates task-aware meta-learning with heterogeneous interaction and knowledge graph modeling to improve personalized POI recommendation, particularly under diverse and dynamic user preferences. As shown in Figure 1, the model consists of four tightly coupled components: (1)**Interaction Graph Learning**. We construct an *attribute-enriched heterogeneous interaction graph* to capture user–POI interactions, where edges encode temporal, spatial, and frequency attributes. We then learn user and POI representations over this graph using an *edge-aware propagation* mechanism; (2)**Knowledge Graph Learning**. We enhance the POI knowledge graph with *category-transition edges* and *spatial-proximity edges* to model semantic and spatial relations among POIs, and apply a *relation-aware propagation* scheme for representation learning; (3)**Meta-Learning Optimization**. We adopt a meta-learning optimization paradigm that leverages the sampled tasks to train the model in a way that enhances its adaptability and generalization, particularly for cold-start users and unseen scenarios; (4)**Memory-Aware Task Sampling**. To improve the diversity and effectiveness of meta-learning tasks, we design a *memory-aware task selection* strategy that considers both structural heterogeneity and user behavior entropy, dynamically sampling informative and representative tasks. We now detail each component.

Interaction Graph Learning

To capture users' structural preferences under rich contextual information, we construct a lightweight heterogeneous interaction graph based on attribute-aware edges and complement it with a trajectory encoder that models sequential user behavior. The final user representation is generated via a residual fusion mechanism.

Graph Construction. We define the interaction graph as $G_{\text{int}} = (V, E), V = \mathcal{U} \cup \mathcal{P}$, where $\mathcal{U}$ and $\mathcal{P}$ denote the sets of users and POIs, respectively. An edge $(u, p) \in \mathcal{E}$ exists if user u has visited POI p, and is associated with the following behavioral attributes: temporal attribute τ_{up} (discretized time slot), spatial attribute d_{up} (bucketed geospatial distance), and frequency attribute f_{up} (normalized check-in frequency). These attributes are encoded into embedding vectors and concatenated into an edge representation:

$$\mathbf{a}_{up} = [\mathbf{e}_{\tau_{up}} \,\|\, \mathbf{e}_{d_{up}} \,\|\, f_{up}] \tag{1}$$

Edge-aware GNN Propagation. We adopt an edge-aware attention-based GNN for user and POI representation learning. At each GNN layer l, the message aggregation and node representation update are performed as follows:

$$\mathbf{h}_i^{(l+1)} = \text{ReLU}\left(\mathbf{W}_0 \mathbf{h}_i^{(l)} + \sum_{j \in \mathcal{N}(i)} g(\mathbf{a}_{ij}) \cdot \alpha_{ij} \cdot \mathbf{W}\mathbf{h}_j^{(l)}\right) \tag{2}$$

where $\mathbf{h}_i^{(l)}$ is the representation of node i at layer l, $g(\mathbf{a}_{ij}) = sigmoid(\mathbf{w}_g^\top \mathbf{a}_{ij})$ is the edge-level gate that modulates the intensity of propagation based on edge attributes, and α_{ij} is the attention coefficient computed as:

$$\alpha_{ij} = \frac{\exp\left(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}\mathbf{h}_i^{(l)} \,\|\, \mathbf{W}\mathbf{h}_j^{(l)}])\right)}{\sum_{k \in \mathcal{N}(i)} \exp\left(\text{LeakyReLU}(\mathbf{a}^\top [\mathbf{W}\mathbf{h}_i^{(l)} \,\|\, \mathbf{W}\mathbf{h}_k^{(l)}])\right)} \tag{3}$$

where a and $\mathbf{W}$ denote the trainable attention parameters and transformation matrix, respectively. $g(\mathbf{a}_{ij})$ controls the

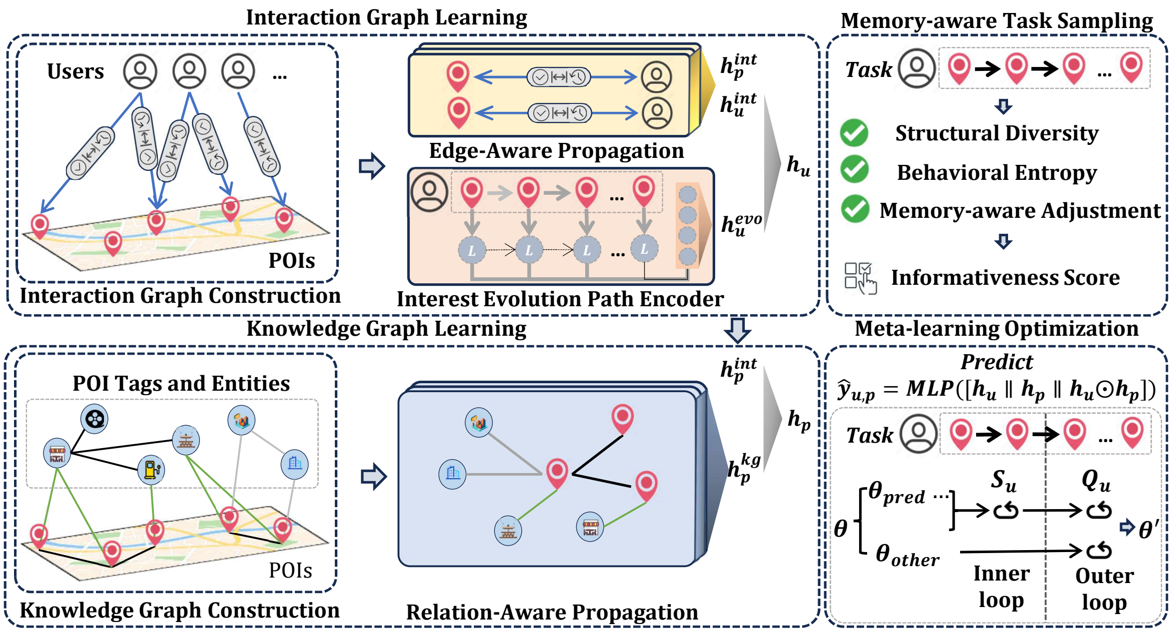

Figure 1: The overall framework of our proposed TMHKG model.

contribution of each edge based on the edge attributes, while α_{ij} captures the relative importance of neighboring nodes through the attention mechanism. After L layers, we fuse all layers to obtain the final structural representation:

$$\mathbf{h}_i^{\text{int}} = \frac{1}{L}\sum_{l=1}^{L}\mathbf{h}_i^{(l)} \tag{4}$$

Interest Evolution Path Encoder. Given a user's check-in sequence $S_u = [p_1, p_2, \ldots, p_n]$, we first obtain POI embeddings from the interaction graph and use a single-layer LSTM encodes the sequence:

$$[\mathbf{z}_1, \ldots, \mathbf{z}_n] = \text{LSTM}([\mathbf{h}_{p_1}^{\text{int}}, \ldots, \mathbf{h}_{p_n}^{\text{int}}]), \quad \forall p_i \in S_u \tag{5}$$

We apply attention pooling to extract trajectory-level semantics:

$$\alpha_i = \frac{\exp(\mathbf{q}^\top \tanh(\mathbf{W}_a \mathbf{z}_i))}{\sum_{j=1}^{n}\exp(\mathbf{q}^\top \tanh(\mathbf{W}_a \mathbf{z}_j))}, \quad \mathbf{h}_u^{\text{evo}} = \sum_{i=1}^{n}\alpha_i \cdot \mathbf{z}_i \tag{6}$$

where $\mathbf{q}$ is a trainable query vector, and $\mathbf{W}_a$ is a trainable projection matrix. The attention pooling allows the model to focus on the most important parts of the user's trajectory.

User Representation Fusion. To adaptively integrate structural and sequential signals, we introduce a residual fusion mechanism:

$$\lambda_u = sigmoid\left(\text{MLP}([\mathbf{h}_u^{\text{int}} \parallel \mathbf{h}_u^{\text{evo}}])\right) \tag{7}$$

where $\lambda_u \in (0, 1)$ is the trajectory importance score predicted by a shallow MLP. The final user representation is:

$$\mathbf{h}_u = \mathbf{h}_u^{\text{int}} + \lambda_u \cdot \mathbf{h}_u^{\text{evo}} \tag{8}$$

Knowledge Graph Learning

To enhance POI semantics with structured side information, we construct a knowledge graph that connects POI nodes with various semantic entities and relations derived from metadata and textual tags. We treat it as an independent, complementary source of structural knowledge and learn POI representations through a relation-aware message propagation mechanism.

Relation-aware Graph Propagation. Inspired by the relation projection idea in TransR (Lin et al. 2015) and its graph-based adaptations, we design a lightweight propagation scheme that distinguishes different edge semantics by projecting node embeddings into relation-specific spaces during message passing. Formally, given a knowledge graph $\mathcal{G}_{\text{KG}} = (\mathcal{V}, \mathcal{E})$ with relation types $r \in \mathcal{R}$, and a relation-specific projection matrix $\mathbf{M}_r$ for each relation, we introduce a virtual self-loop relation r_{self} to preserve node identity features during propagation. The augmented relation set is defined as $\mathcal{R}' = \mathcal{R} \cup \{r_{\text{self}}\}$, and each node i is associated with a self-loop edge (i, r_{self}, i).

The hidden representation of node i at layer $l+1$ is computed as:

$$\mathbf{h}_i^{(l+1)} = \text{ReLU}\left(\sum_{r\in\mathcal{R}'}\sum_{j\in\mathcal{N}_r(i)}\frac{1}{c_{i,r}} \cdot \mathbf{W}_r \cdot \mathbf{M}_r \mathbf{h}_j^{(l)}\right) \tag{9}$$

where $\mathcal{N}_r(i)$ denotes the set of neighbors of node i under relation r (including i itself when $r = r_{\text{self}}$), $\mathbf{W}_r \in \mathbb{R}^{d\times d}$ is a learnable transformation matrix, and $\mathbf{M}_r$ projects the input embedding into the semantic space of relation r. The normalization constant $c_{i,r}$ is set to the cardinality of the neighbor set, i.e., $|\mathcal{N}_r(i)|$.

Behavior-augmented Relation Extensions. To further infuse the KG with user-driven semantics, we introduce two novel relation types: (1) *category transition edges*, built from co-occurrence patterns in user trajectories, capturing probabilistic transitions between semantic types of POIs, and (2) *spatial proximity edges*, connecting POIs within a fixed geographic radius, weighted by an exponential distance-decay function. These new relations are seamlessly incorporated into the propagation process by treating them as additional relations in the knowledge graph. Each relation type has its own learnable projection matrix $\mathbf{M}_r$ and transformation matrix $\mathbf{W}_r$, allowing the model to adapt and enhance POI representations with user-centric semantics. These extensions provide additional behavioral and spatial inductive biases to better capture complex user preferences in the POI space.

POI Representation Fusion. After independently encoding POI representations from the interaction graph and the knowledge graph, we fuse the two views to obtain a unified POI representation that captures both behavioral and semantic information with a gated fusion mechanism.

Task Formulation and Meta-learning Optimization

To capture diverse and evolving user preferences, we formulate POI recommendation as a meta-learning problem, inspired by the Model-Agnostic Meta-Learning (MAML) (Finn, Abbeel, and Levine 2017) framework, where each user is treated as a distinct task. The goal is to learn a generalizable model that can be efficiently adapted to user-specific behavior via limited interaction data, following MAML's principle of optimizing for fast adaptation.

Task Definition. For each user u, we define a meta-task $\mathcal{T}_u = (\mathcal{S}_u, \mathcal{Q}_u)$, where $\mathcal{S}_u$ denotes the support set (few-shot historical check-ins) and $\mathcal{Q}_u$ is the query set used for meta-evaluation. The model is trained to maximize ranking performance on $\mathcal{Q}_u$ after adapting on $\mathcal{S}_u$.

Scoring Function and Training Objectives To estimate the affinity between users and POIs, we employ a learnable matching network that captures high-order interactions between heterogeneous embeddings. Given a user representation $\mathbf{h}_u \in \mathbb{R}^d$ and a POI representation $\mathbf{h}_p \in \mathbb{R}^d$, both obtained via multi-source fusion modules, we define the prediction score as:

$$\hat{y}_{u,p} = \text{MLP}([\mathbf{h}_u \,\|\, \mathbf{h}_p \,\|\, \mathbf{h}_u \odot \mathbf{h}_p]) \tag{10}$$

Here, $[\cdot\|\cdot\|\cdot]$ denotes vector concatenation, and $\odot$ is element-wise multiplication.

To optimize the model, we adopt a multi-task learning objective that integrates recommendation accuracy and representation consistency.

For each user task $\mathcal{T}_u = (\mathcal{S}_u, \mathcal{Q}_u)$, the main goal is to rank observed POIs higher than unobserved ones. We adopt a Bayesian Personalized Ranking (BPR) loss over the query set:

$$\mathcal{L}_{\text{rank}}^{(u)} = -\sum_{(p^+,p^-)\in\mathcal{Q}_u} \log \sigma\left(\hat{y}_{u,p^+} - \hat{y}_{u,p^-}\right) \tag{11}$$

where $\sigma(\cdot)$ is the sigmoid function.

To encourage semantic alignment between trajectory-based and structure-based user representations, we minimize their cosine distance:

$$\mathcal{L}_{\text{align}}^{(u)} = 1 - \frac{\mathbf{h}_u^{\text{int}} \cdot \mathbf{h}_u^{\text{evo}}}{\|\mathbf{h}_u^{\text{int}}\|_2 \cdot \|\mathbf{h}_u^{\text{evo}}\|_2} \tag{12}$$

To promote representational coherence across the interaction and knowledge views of POIs, we apply a cosine similarity-based alignment loss:

$$\mathcal{L}_{\text{poi-align}} = \sum_{p\in\mathcal{P}_{\text{batch}}} \left[1 - \frac{\mathbf{h}_p^{\text{int}} \cdot \mathbf{h}_p^{\text{kg}}}{\|\mathbf{h}_p^{\text{int}}\|_2 \cdot \|\mathbf{h}_p^{\text{kg}}\|_2}\right] \tag{13}$$

The overall loss for each task $\mathcal{T}_u$ is formulated as a weighted sum:

$$\mathcal{L}^{(u)} = \mathcal{L}_{\text{rank}}^{(u)} + \beta(\alpha \cdot \mathcal{L}_{\text{align}}^{(u)} + (1-\alpha) \cdot \mathcal{L}_{\text{poi-align}}) \tag{14}$$

where α and β are two hyperparameters to control the consistency loss.

Parameter Grouping. We partition all trainable parameters θ into the following groups:

$$\theta = \{\theta_{\text{int-graph}}, \theta_{\text{kg-graph}}, \theta_{\text{iepe}}, \theta_{\text{user-fuse}}, \theta_{\text{fuse}}, \theta_{\text{pred}}\} \tag{15}$$

where $\theta_{\text{int-graph}}$ and $\theta_{\text{kg-graph}}$ include all learnable parameters from the interaction graph and the knowledge graph encoders, respectively (e.g., relation projection matrices, attention weights, and edge gates). θ_{iepe} corresponds to the parameters of the Interest Evolution Path Encoder, $\theta_{\text{user-fuse}}$ controls the residual-style user representation fusion, θ_{fuse} handles the POI-level fusion between interaction- and KG-based embeddings, and θ_{pred} is the user–POI scoring module.

Inner-loop Adaptation. For each task $\mathcal{T}_u$, we perform task-specific adaptation on a subset of parameters using the support set $\mathcal{S}_u$. The adapted parameter group is defined as:

$$\theta_{\text{adapt}} = \{\theta_{\text{int-graph}}, \theta_{\text{user-fuse}}, \theta_{\text{pred}}\} \tag{16}$$

These parameters are updated via one or more gradient steps:

$$\theta'_u = \theta_{\text{adapt}} - r_{inner} \nabla_{\theta_{\text{adapt}}} \mathcal{L}_{\mathcal{S}_u}(\theta) \tag{17}$$

where r_{inner} is the inner-loop learning rate and $\mathcal{L}_{\mathcal{S}_u}$ is the loss computed on the support set.

Outer-loop Optimization. After adaptation, the updated parameters θ'_u are used to evaluate the loss on the query set $\mathcal{Q}_u$, and the meta-objective is aggregated over a batch of tasks $\mathcal{B}$:

$$\mathcal{L}_{\text{meta}} = \sum_{u\in\mathcal{B}} \mathcal{L}_{\mathcal{Q}_u}(\theta'_u) \tag{18}$$

All parameters in θ are then updated using the outer-loop learning rate r_{outer}:

$$\theta \leftarrow \theta - r_{outer} \nabla_\theta \mathcal{L}_{\text{meta}} \tag{19}$$

This meta-learning framework enables the model to acquire globally transferable knowledge while retaining the flexibility to quickly personalize representations for diverse user behaviors.

Memory-aware Task Sampling

To enhance the robustness and diversity of meta-training, we propose a memory-aware task sampling strategy that prioritizes structurally distinct and behaviorally diverse user tasks. Unlike uniform sampling, our method computes an informativeness score for each candidate task and samples from a soft distribution derived from these scores.

Task Scoring. Each user task $\mathcal{T}_u$ is assigned a composite score s_u that integrates two signals: structural distinctiveness and behavioral entropy.

For structural diversity, we measure how different a user's interaction-graph-based embedding $\mathbf{z}_u^{\text{graph}}$ is from the mean embedding of previously selected tasks. Let $\mathcal{H}$ denote the set of already selected tasks in the current episode. The structural score is defined as:

$$s_u^{\text{struct}} = \left\|\mathbf{h}_u^{\text{graph}} - \bar{\mathbf{h}}\right\|_2, \quad \bar{\mathbf{h}} = \frac{1}{|\mathcal{H}|}\sum_{v\in\mathcal{H}} \mathbf{h}_v^{\text{graph}} \tag{20}$$

This formulation reduces computational overhead by avoiding pairwise comparisons and emphasizing deviation from the overall task distribution.

For behavioral entropy, we compute the uncertainty of a user's check-in distribution over POI categories. Let p_c^u denote the proportion of user u's check-ins in category c, then:

$$s_u^{\text{entropy}} = -\sum_c p_c^u \log p_c^u \tag{21}$$

The final score is computed as a combination of the two min-max-normalized metrics:

$$s_u = s_u^{\text{struct}-norm} + s_u^{\text{entropy}-norm} \tag{22}$$

Memory-aware Adjustment. To reduce selection redundancy, we apply a penalty to frequently chosen tasks. Let $\text{freq}(u)$ denote the number of times task u has been selected so far. The adjusted score is computed as:

$$\tilde{s}_u = \frac{s_u}{1+\log(1+\text{freq}(u))} \tag{23}$$

We select meta-training tasks from the candidate pool by sampling according to the adjusted scores. This sampling strategy encourages task diversity while allowing moderately informative tasks to be selected, improving generalization during meta-learning.

Experiment

To comprehensively evaluate the effectiveness of our proposed TMHKG framework, we conduct extensive experiments on two real-world POI recommendation datasets. Our goal is to investigate whether TMHKG can effectively leverage heterogeneous interaction and knowledge graphs, along with a task-aware meta-learning strategy, to improve the accuracy and adaptability of personalized POI recommendation. To this end, we formulate the following research questions:

- **RQ1**: How does TMHKG perform compared to state-of-the-art POI recommendation models?
- **RQ2**: How do the main components of TMHKG contribute to the overall performance?
- **RQ3**: How do different hyperparameter settings impact the performance of TMHKG?
- **RQ4**: Can TMHKG effectively address the cold-start problem?

Experiment Settings

Datasets. We use two public datasets to evaluate the effectiveness of TMHKG: Yelp and NYCRestaurant. The two datasets are publicly accessible and vary in size and sparsity, making our experiments more convincing. **Yelp**: This dataset provides real-world data related to businesses including reviews, photos, check-ins, and attributes like hours, parking availability, and ambience. **NYCRestaurant**: This dataset includes check-in, tip, and tag data of restaurant venues in NYC collected from Foursquare from 24 October 2011 to 20 February 2012. In our experiment, we preprocess the three datasets by filtering out the inactive users and unpopular POIs. Specifically, we filter out those users who have fewer than 10 check-in POIs and those POIs that are visited by fewer than 10 users. Then, 3:1:1 of each processed dataset is selected as the training, validation, and testing datasets. Furthermore, following the knowledge graph construction methodology employed in the knowledge graph learning module, we construct a knowledge graph based on POI metadata and Freebase. The statistics of two datasets are described in Table 1.

		Yelp	Foursquare
G_{int}	#Users	30,887	3,112
	#POIs	18,995	3,298
	#Check-ins	860,888	27,149
G_{kg}	#Entities	47,487	5,629
	#Relations	32	5
	#Triplets	95,864	10,331

Table 1: Dataset statistics.

Baselines. To demonstrate the effectiveness of our proposed TMHKG framework, we compare it with three categories of representative POI recommendation methods: knowledge graph-based methods (KGIN, KGAT), behavior evolution-based methods (GETNext, STGCN), and meta-learning-based methods (MetaKG, HyperMAN). The detailed descriptions are as follows: *KGIN* (Wang et al. 2021) is a knowledge graph-enhanced recommendation method that integrates multi-hop semantic relations among entities into user and item representations via high-order graph convolution and cross-neighbor propagation; *KGAT* (Wang et al. 2019) introduces a graph attention mechanism to model the varying importance of different knowledge relations, enabling more interpretable and discriminative interest representations; *GETNext* (Yang, Liu, and Zhao 2022) constructs a spatio-temporal-semantic heterogeneous graph based on user mobility trajectories and learns preference evolution paths by modeling multi-type relations; *STGCN* (Han et al. 2020) captures spatio-temporal dependencies in user mobility trajectories to model sequential POI visiting behaviors; *MetaKG* (Du et al. 2022) employs collaborative-aware and knowledge-aware meta-learners to capture user preferences and knowledge associations, effectively addressing cold-start issues; *HyperMAN* (Wang et al. 2025) combines meta-learning with high-order structural modeling, enhancing the model's transferability and adaptability across user-specific tasks through multi-task training and structural contrastive signals, and constructs a task-aware recommendation model on multi-relational hypergraphs.

Evaluation metrics. We choose *Recall@K* and *NDCG@K* as the evaluation metrics. *Recall@K* measures the proportion of relevant POIs that are present in the top k recommended items for a given user. The Normalized Discounted Cumulative Gain (NDCG) penalizes irrelevant items ranked higher and rewards relevant items placed at top positions. Both metrics range from 0 to 1, where a value closer to 1 indicates that the recommendation system is effectively ranking relevant POIs higher in the list.

Parameter Settings and Implementation. We implement our TMHKG and all baselines in Pytorch and carefully tune the key parameters. For a fair comparison, we fix the embedding size to 64 for all models, and the embedding parameters are initialized with the Xavier method (Glorot and Bengio 2010). We optimize our method with Adam (Adam et al. 2014) and set the batch size to 4096. A grid search is conducted to confirm the optimal settings. We tune the learning rate η among $\{0.0001, 0.0003, 0.001, 0.003\}$. The best settings for hyperparameters in all comparison methods are researched by either empirical study or by following the original papers.

Performance Comparison (RQ1)

We report the empirical results of all methods in Table 2. The improvements and statistical significance test are performed between TMHKG and the strongest baselines highlighted with an underline. Analyzing such performance comparison, we have the following observations: (1)**Superior Performance Across All Metrics.** The proposed method achieves the best performance across all evaluation metrics, demonstrating the overall effectiveness of the framework. On the two real-world datasets, Yelp and NYCRestaurant, our method consistently outperforms representative baselines in terms of Recall@k and NDCG@k. Notably, on the NDCG@20 metric, our method improves by 5.77% and 4.77% compared to the best-performing baselines on Yelp and NYCRestaurant, respectively, highlighting its advantage in personalized ranking quality; (2)**Effective Modeling of Dynamic User Interests.** Our method exhibits significant advantages in capturing the dynamic evolution of user interests, outperforming behavior-evolution modeling approaches. While models such as GETNext and STGCN leverage user trajectories to model interest evolution, most existing methods fail to consider the synergy between knowledge structures and preference dynamics. In contrast, our method integrates interest-path modeling, knowledge graph propagation, and task-level adaptive mechanisms to effectively capture latent evolutionary trends in users' spatiotemporal behavior sequences, thereby achieving stronger preference representation capabilities; (3)**Superiority Over Knowledge-Enhanced and Meta-Learning-Based Methods.** The incorporation of knowledge graphs and task modeling mechanisms yields substantial improvements over existing knowledge-enhanced and meta-learning-based recommendation methods. For instance, while KGIN and KGAT can model knowledge graph structures, they lack personalized task modeling, resulting in limited performance. HyperMAN incorporates meta-learning but does not explicitly model interest evolution paths, leading to constrained interpretability and stability in recommendations. In comparison, our method leverages fine-grained knowledge structure modeling and task-driven representation optimization, achieving consistently superior results across multiple evaluation metrics.

Ablation Study(RQ2)

To validate the contribution of each component, we conduct an ablation study with five variants:(1)**w/o IEPE** removes the Interest Evolution Path Encoder to assess dynamic preference modeling;(2)**w/o KG** omits the knowledge graph entirely;(3)**w/o KG+** uses the original KG without our proposed relational enhancements;(4)**w/o Sampling** replaces our memory-aware sampling with a random strategy;(5)**w/o RCL** removes the representation consistency loss.

The results of our ablation study, presented in Table 3, demonstrate the integral contribution of each component to TMHKG's performance. The full model consistently outperforms all variants. The most significant performance degradation is observed when removing the knowledge graph entirely, which causes Recall@10 on Yelp to plummet from 0.0901 to 0.0664. This underscores the critical role of external semantic knowledge. Similarly, removing the dynamic interest modeling or the adaptive task sampling also leads to substantial drops in performance, highlighting the necessity of modeling user evolution and employing an effective meta-training strategy. Furthermore, using the base knowledge graph without our structural enhancements or removing the consistency loss results in a noticeable, albeit smaller, decline in performance. This confirms the value of our behavior-aware edge enhancements and the regularization effect of the alignment loss. In conclusion, each component of the proposed model plays a crucial role in performance improvement, collectively supporting the recommendation system's accuracy, personalization, and generalization capabilities.

Parameter Analysis (RQ3)

Here we investigate the effect of different settings of the L, L', α, and β parameters in our framework. Specifically, we conduct experiments on the Yelp dataset, keeping all other parameters fixed at their optimal values while varying one parameter at a time. As shown in Figure 2, the performance initially improves as the parameters increases, reaches a peak, and then declines.

For the number of GNN layers, L(interaction) and L'(knowledge), performance peaks at a moderate depth. Too few layers result in insufficient feature propagation, while too many lead to overfitting. The hyperparameters α and β control the consistency loss. α balances the user and POI alignment objectives, while β balances the overall consistency loss against the main ranking task. For both, extreme values hurt performance by either under-regularizing or overpowering the primary ranking objective. Our results

Dataset	Yelp						NYCRestraunt					
Model	**R@5**	**R@10**	**R@20**	**N@5**	**N@10**	**N@20**	**R@5**	**R@10**	**R@20**	**N@5**	**N@10**	**N@20**
KGIN	0.0459	0.0756	0.1269	0.0506	0.0617	0.0847	0.0432	0.0817	0.0915	0.0568	0.0604	0.0792
KGAT	0.0485	0.0798	0.1291	0.0549	0.0642	0.0891	0.0457	0.0821	0.1092	0.0515	0.0595	0.0799
STGN	0.0518	0.0844	0.1328	0.0564	0.0836	0.0956	0.0492	0.0820	0.1108	0.0564	0.0789	0.0900
GETN	0.0524	0.0859	0.1304	0.0584	0.0821	0.0970	0.0510	0.0816	0.1169	0.0565	0.0807	0.0902
MetaK	0.0517	0.0839	0.1326	0.0566	0.0825	0.0945	0.0507	0.0818	0.1141	0.0543	0.0785	0.0899
Hyper	0.0519	0.0841	0.1348	0.0584	0.0827	0.0961	0.0511	0.0819	0.1154	0.0550	0.0794	0.0901
Ours	**0.0564**	**0.0901**	**0.1416**	**0.0614**	**0.0879**	**0.1026**	**0.0537**	**0.0863**	**0.1225**	**0.0599**	**0.0852**	**0.0945**
Imprv.	+7.63%	+4.89%	+5.05%	+5.14%	+5.15%	+5.77%	+5.09%	+5.12%	+4.79%	+5.46%	+5.57%	+4.77%

Table 2: Performance comparison in Recall@k and NDCG@k on two datasets.

Model	Yelp		NYCRestaurant	
	R@10	N@10	R@10	N@10
w/o IEPE	0.0815	0.0839	0.0849	0.0782
w/o KG	0.0664	0.0681	0.0645	0.0632
w/o KG+	0.0881	0.0837	0.0842	0.0828
w/o Sampling	0.0783	0.0768	0.0736	0.0697
w/o RCL	0.0884	0.0872	0.0855	0.0831
Full Model	**0.0901**	**0.0879**	**0.0863**	**0.0852**

Table 3: Effect of ablation study.

show that optimal performance is achieved when these components are properly balanced.

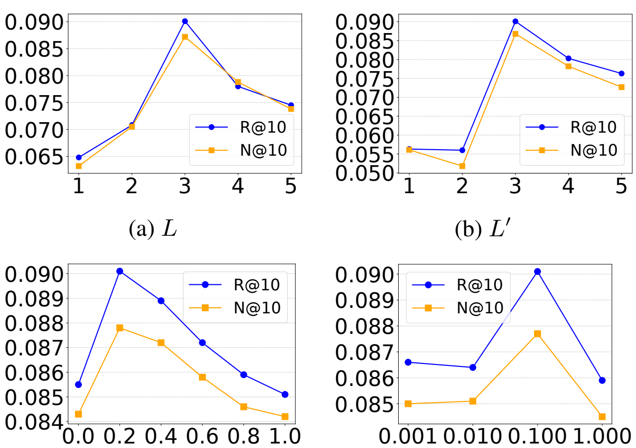

(a) L (b) L'

(c) α (d) β

Figure 2: Effect of L, L', α and β.

Cold-start Scenarios Analysis (RQ4)

To investigate the ability of TMHKG to handle cold-start scenarios, we design a dedicated experiment focusing on the user cold-start (UC) setting. Specifically, we conduct experiments on both the Yelp and NYCRestaurant datasets, where we simulate the UC condition by selecting users with very few historical interactions as the cold-start users. For each dataset, we sort all users by their number of interactions and designate the bottom 20% of users as cold-start users. These users are entirely excluded from the training set and only used for evaluation, while the remaining users are used for training and validation.We compare TMHKG with two representative baselines: GETNext and HyperMAN. Both baselines are trained and evaluated under the same cold-start data splits.

As shown in Table 4, TMHKG consistently achieves superior performance compared to the baselines in the UC setting on both datasets. This result demonstrates the effectiveness of our task-aware meta-learning strategy and memory-aware task sampling mechanism in improving model adaptability and generalization to users with sparse historical data.

Model	Yelp		NYCRestaurant	
	R@10	N@10	R@10	N@10
GETNext	0.0411	0.0402	0.0362	0.0324
HyperMAN	0.0566	0.0532	0.0489	0.0488
Ours	**0.0820**	**0.0812**	**0.0790**	**0.0707**

Table 4: Performance under user cold-start setting.

Conclusion

In this paper, we proposed TMHKG, a task-aware meta-learning framework for personalized POI recommendation. Our framework learns to adapt to individual users by synergistically modeling dynamic preferences from interaction graphs and rich semantics from an enhanced knowledge graph, guided by a memory-aware task sampling strategy. Extensive experiments demonstrate that TMHKG significantly outperforms state-of-the-art methods by effectively addressing the challenge of user preference diversity. Future work will focus on exploring additional context signals and further optimizing the task adaptation process.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (No. 62276026).

References

Adam, K. D. B. J.; et al. 2014. A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 1412(6).

Du, Y.; Zhu, X.; Chen, L.; Fang, Z.; and Gao, Y. 2022. Metakg: Meta-learning on knowledge graph for cold-start recommendation. *IEEE Transactions on knowledge and data engineering*, 35(10): 9850–9863.

Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.

Glorot, X.; and Bengio, Y. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 249–256. JMLR Workshop and Conference Proceedings.

Han, H.; Zhang, M.; Hou, M.; Zhang, F.; Wang, Z.; Chen, E.; Wang, H.; Ma, J.; and Liu, Q. 2020. STGCN: a spatial-temporal aware graph learning method for POI recommendation. In *2020 IEEE International Conference on Data Mining (ICDM)*, 1052–1057. IEEE.

Li, H.; Wan, H.; Chen, Y.; Ye, D.; Gel, Y.; and Jiang, H. 2025. TMetaNet: Topological Meta-Learning Framework for Dynamic Link Prediction. *arXiv preprint arXiv:2506.00453*.

Lim, N.; Hooi, B.; Ng, S.-K.; Wang, X.; Goh, Y. L.; Weng, R.; and Varadarajan, J. 2020. STP-UDGAT: Spatial-temporal-preference user dimensional graph attention network for next POI recommendation. In *Proceedings of the 29th ACM international conference on information & knowledge management*, 845–854.

Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; and Zhu, X. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29.

Liu, Q.; Zeng, Y.; Mokhosi, R.; and Zhang, H. 2018. STAMP: short-term attention/memory priority model for session-based recommendation. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 1831–1839.

Mandal, D.; Medya, S.; Uzzi, B.; and Aggarwal, C. 2022. Metalearning with graph neural networks: Methods and applications. *ACM SIGKDD Explorations Newsletter*, 23(2): 13–22.

Procopio, L.; Tripodi, R.; and Navigli, R. 2021. SGL: Speaking the graph languages of semantic parsing via multilingual translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 325–337.

Shu, T.; Shi, L.; Zhu, C.; and Liu, X. 2022. A graph neural network framework based on preference-aware graph diffusion for recommendation. *Frontiers in Psychiatry*, 13: 1012980.

Song, X.; Liu, Z.; Meng, L.; Chu, D.; Yu, J.; and Sheng, Q. Z. 2025. Accurate POI recommendation for random groups with improved graph neural networks and a multi-negotiation model. *Scientific Reports*, 15(1): 7531.

Sun, H.; Xu, J.; Zheng, K.; Zhao, P.; Chao, P.; and Zhou, X. 2021. MFNP: A meta-optimized model for few-shot next POI recommendation. In *IJCAI*, volume 2021, 3017–3023.

Wang, J.; Zhang, T.; Zhang, L.; Bai, Y.; Li, X.; and Jin, J. 2025. HyperMAN: Hypergraph-enhanced Meta-learning Adaptive Network for Next POI Recommendation. *arXiv preprint arXiv:2503.22049*.

Wang, X.; He, X.; Cao, Y.; Liu, M.; and Chua, T.-S. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 950–958.

Wang, X.; Huang, T.; Wang, D.; Yuan, Y.; Liu, Z.; He, X.; and Chua, T.-S. 2021. Learning intents behind interactions with knowledge graph for recommendation. In *Proceedings of the web conference 2021*, 878–887.

Yang, S.; Liu, J.; and Zhao, K. 2022. GETNext: Trajectory flow map enhanced transformer for next POI recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on research and development in information retrieval*, 1144–1153.

Zhang, J.; and Ma, W. 2024. Hybrid structural graph attention network for POI recommendation. *Expert Systems with Applications*, 248: 123436.

Zhao, P.; Luo, A.; Liu, Y.; Xu, J.; Li, Z.; Zhuang, F.; Sheng, V. S.; and Zhou, X. 2020. Where to go next: A spatio-temporal gated network for next poi recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 34(5): 2512–2524.