



Hyperbolic Multi-semantic Transition for Next POI Recommendation

Hongliang Qiao
The Hong Kong Polytechnic
University
Hong Kong SAR, China
hongliang.qiao@connect.polyu.hk

Shanshan Feng
Centre for Frontier AI Research,
A*STAR
Singapore, Singapore
victor_fengss@foxmail.com

Min Zhou
Huawei Cloud
Shenzhen, China
zhoumin27@huawei.com

WenTao Li
Bank of Communications
Wuhan, China
wentao_l@bankcomm.com

Fan Li*
The Hong Kong Polytechnic
University
Hong Kong SAR, China
fan-5.li@polyu.edu.hk

Abstract

The next Point-of-Interest (POI) recommendation has gained significant research interest, focusing on learning users' mobility patterns from sparse check-in data. Existing POI recommendation models face two main constraints. First, most models are based on Euclidean space and struggle with capturing the inherent hierarchical structures in historical check-ins. Second, various transition semantics in both one-hop and sequential transitions cannot be properly utilized to understand user movement trends. To overcome the above limitations, we introduce rotation operations in hyperbolic space, enabling the joint modeling of hierarchical structures and various transition semantics to effectively capture complex mobility patterns. Specifically, a novel hyperbolic rotation-based recommendation model HMST is developed for the next POI recommendation. To our knowledge, this is the first work to explore the hyperbolic rotations for the next POI recommendation tasks. Extensive experiments on three real-world datasets demonstrate the superiority of our proposed approach over the various state-of-the-art baselines.

CCS Concepts

• Computing methodologies → Artificial intelligence; • Information systems → Data mining.

Keywords

Next POI Recommendation; Hyperbolic Representation Learning

ACM Reference Format:

Hongliang Qiao, Shanshan Feng, Min Zhou, WenTao Li, and Fan Li. 2025. Hyperbolic Multi-semantic Transition for Next POI Recommendation. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3701716.3717802>

*The corresponding author.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.

WWW Companion '25, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1331-6/2025/04

<https://doi.org/10.1145/3701716.3717802>

1 Introduction

In recent years, location-based social networks (LBSNs) such as Foursquare have experienced rapid growth. In these social networks, users share their locations by checking in at Points-of-Interest (POIs). With the increasing availability of check-in data, POI recommendation [25] has been extensively investigated, helping users to explore their surroundings better and discover interesting locations based on their preferences. Among various POI recommendation tasks, next-POI recommendation [4, 21] is one of the most extensively studied recommendation tasks, aiming to suggest POIs for a user to visit subsequently given the recent check-in trajectory.

With the rapid advancement of machine learning, methods [7, 30, 34, 37, 50] that leverage various learning techniques have achieved huge improvements in POI recommendation tasks. Among these, methods [13, 34, 43] utilize graph representation learning to explore complex relations between entities in check-in data, such as users, POIs, and categories. These approaches construct node representations to model personalized movements and predict the next location a user might visit. While effective, most existing methods operate in Euclidean space, which limits their ability to capture complex patterns due to the inherent properties of Euclidean geometry [22]. In contrast, Feng et al. [9] demonstrated that the distributions of various entities have inherent hierarchical patterns and proposed HME to learn entity representations in hyperbolic space. Although hyperbolic representations offer greater expressive power than their Euclidean counterparts, HME treats all relationships equally and fails to capture the rich transition semantics in consecutive location visits, leading to limited prediction performance. Therefore, the potential of leveraging hyperbolic space for the next POI recommendation requires further exploration.

Multiple transition semantics occur simultaneously when users move between two POIs. In addition to the locations, user movements involve category and region transitions. Beyond user preferences, these transition semantics are crucial for understanding personalized movement patterns. The dense transition patterns in category and region semantics further provide additional information that mitigates the sparsity of POIs. Thus, integrating various transition semantics is essential for accurate recommendations. Furthermore, we identified two types of transitions in check-in data as is shown in Figure 1. The first is the one-hop transition between




Figure 1: Illustration of various transition semantics and two types of transitions in check-in data.

two POIs, capturing global correlations of POI transitions from historical check-in data. The second is the sequential transition that denotes multi-hop continuous movements within a user trajectory. This type of transition reflects the dynamic dependencies of user movements across multiple POIs and is vital for personalized recommendations. Leveraging both types of transitions can effectively enhance the performance of POI recommendation tasks.

Motivated by this, we propose a new approach to model complex transition semantics in hyperbolic space. Specifically, our goal is to simultaneously preserve latent hierarchical structures and capture diverse transition semantics to enhance the accuracy of mobility predictions. However, achieving this is nontrivial due to several challenges. First, appropriate operations are required to model transition relations with diverse semantics. Previous methods often relied on Euclidean space operations to fuse different representations and incorporate various types of information. However, these operations cannot be directly applied to hyperbolic space, as they fail to effectively preserve hierarchical features. Second, various transition semantics in both one-hop and sequential transitions should be fully considered, effectively extracting semantics from both transition types to depict user movement trends. Moreover, all extracted transition semantics should be effectively utilized to make final recommendations.

To address these challenges, we explore rotation operations in hyperbolic space. That is, we first preserve latent hierarchical patterns to create expressive embeddings for users, POIs, categories, and regions. Then, we apply rotation transformations to these embeddings to capture various transition semantics. This approach offers several advantages for the POI tasks. Rotation operations do not change the norms of hyperbolic vectors, thus maintaining latent hierarchies. Diverse transition semantics can be effectively captured through different rotation angles. In this way, we encode transition semantics in distinct rotation matrices and utilize them to make recommendations.

In this work, we propose a new recommendation model HMST for the next POI recommendation. HMST consists of four main components. First, it captures user preferences by learning hyperbolic user and POI embeddings from user-POI interactions. Second, it encodes one-hop transition semantics, POI, category, and region transitions, into three learnable rotation matrices. Third, given a user trajectory, it constructs trajectory representations to extract

dynamic dependencies via hyperbolic attention operation and combines the learned rotation matrices to capture sequential transition semantics. Finally, it predicts the next POI, category, and region using recommendation scores calculated based on hyperbolic distances between trajectory representations and entity embeddings.

The main contributions are summarized as follows:

- We incorporate rotation operations in hyperbolic space to model multi-semantic transitions for the next POI recommendation. That is, we propose an effective approach to jointly preserve hierarchical structures and capture multi-semantic transitions in user mobility behaviors.
- We develop a hyperbolic rotation-based recommendation model HMST for the next POI recommendation, which contains heterogeneous graph learning and Lorentzian trajectory learning. To our knowledge, this is the first work to explore hyperbolic representations and rotation operations for the POI tasks.
- We conduct extensive experiments on three real-world check-in datasets, and empirical results demonstrate that HMST significantly outperforms various baselines, which indicates the advantages of the hyperbolic spaces and rotation techniques. Our code is available at <https://github.com/PlaymakerQ/HMST>.

2 Related Work

2.1 Next-POI Recommendation

The next POI recommendation task has attracted extensive research interest. Recent representative solutions include recurrent neural networks [19, 50], graph neural networks [13, 33, 34, 37, 43], and transformers [7, 41, 48]. For example, [37] introduces a hypergraph convolutional network to incorporate trajectory similarity. To capture collaborative signals from other users, [41] constructs a trajectory flow graph and designs a graph-enhanced transformer model for predicting a user's future mobility. [34] develops an adaptive graph representation network, which utilizes POI representations and an attention aggregation to capture the spatial-temporal dependencies. By using long and short-term preferences, [7] proposes a shared trajectory encoder for learning check-in sequences. However, these methods are based on Euclidean spaces where the hierarchical structures cannot be well preserved.

By providing auxiliary information for sparse POI-level data, hierarchical structures have been explored to enhance the next POI recommendation, including POI categories [44, 47, 49] and geographical regions [18–20, 36]. For instance, Nicholas et al. [19] develop a hierarchical multi-task graph recurrent network (HMT-GRN), using beam search to explore different user-region matrices. Different from existing approaches that use different levels of sequential transition separately, we introduce a novel strategy: utilizing hyperbolic space to capture the hierarchical structures and leveraging rotation operations to reflect the multi-semantic transitions simultaneously.

2.2 Hyperbolic Recommender Systems

To capture complex hierarchical patterns, hyperbolic learning techniques have been introduced for recommendation tasks [26, 29]. Recent developments include knowledge-aware recommendation [3,

6, 45], social recommendation [31, 42], session-based recommendation [11], and news recommendation [32]. For example, collaborative filtering techniques are combined with hyperbolic space [16, 39, 40], outperforming the conventional collaborative Filtering methods. Yang et al. [42] developed an HGSR framework to jointly consider the user-user and user-item relations. Notably, a hyperbolic metric embedding approach [9] is proposed for next-POI recommendation, which considers user, item, region, and category into the same Poincaréball model. Although can effectively model complex graph structures, existing studies usually focus on learning hyperbolic node representations, ignoring rich transition semantics in users' mobility behaviors.

2.3 Rotation Techniques

The rotation techniques were originally designed to learn various relations patterns in the knowledge graph learning problems [5, 24, 27]. To further incorporate underlying hierarchical structures in the knowledge graph, hyperbolic rotations have been developed [1] for knowledge graph embedding. Motivated by the rotated Lorentzian graph embedding model [8], [23] utilizes rotation transformation to model asymmetric diffusion patterns in information cascades. Different from these studies, we attempt to explore the rotation operations to capture the rich multiple transition semantics in both one-hop transition and sequential transition in check-in data.

3 Preliminary

3.1 Notations and Problem Definition

Let $\mathcal{U}$, $\mathcal{L}$, $\mathcal{C}$, and $\mathcal{R}$ be the set of users, POIs, categories, and regions, respectively. A POI l is a location typically associated with geographical coordinates and a category $c \in \mathcal{C}$ to denote its function. Regions in $\mathcal{R}$ are constructed by dividing the entire area based on collected geographical coordinates to identify the region r a POI l belongs to. A check-in $q = (u, l, t, c, r)$ records that a user u visits a specific POI l at timestamp t , c and r denote the category and region information of l , respectively. A check-in trajectory s_i^u records all check-ins generated by a user u in a certain period¹ which can be viewed as a sequence $s_i^u = \{q_1, q_2, \dots, q_m\}$.

From historical check-ins, we could further extract different relations that have been shown effective in the next POI recommendation. In this work, we use edges to represent these entities and relations. (1) User-POI edge $e_{u,l} := (u, l)$: If a user u has visited a POI l , there is a User-POI edge (u, l) to indicate the interaction between the user and POI. User-POI interactions often indicate user preferences. We use $\mathcal{E}_{UL}$ to denote the set of all user-POI edges. (2) POI-POI edge $e_{l_1,l_2} := (l_1, l_2)$: Two locations l_1 and l_2 visited by a user within 6 hours are assumed to exhibit distinct sequential patterns [9]. In this way, we can extract all one-hop transition relations in the edge set $\mathcal{E}_{LL}$.

Next POI Recommendation Problem. Given a set of historical check-ins and a user's recent trajectory $s_i^u = \{q_1, q_2, \dots, q_m\}$, the purpose of the next-POI recommendation is to recommend a potential POI $l_{m+1} \in \mathcal{L}$ that the user u would visit next.

¹Similar to [41], we partition all check-ins of a specified user into check-in trajectories by 24-hour intervals.

3.2 Preliminary Knowledge

Hyperbolic Geometry. In this work, we adopt the Lorentz model for describing hyperbolic spaces. A n -dimensional Lorentz model is defined as the Riemannian manifold $\mathbb{L}_\beta^n = (\mathcal{H}_\beta^n, g_\beta)$ with a negative curvature $-1/\beta$, where $\beta > 0$ is the curvature parameter and $g_\beta = \text{diag}([-1, 1, 1, \dots, 1])$ is the metric tensor. When $\beta = 1$, the Lorentz space equals the unit hyperboloid model. The point set in the Lorentz model is defined as $\mathbb{L}_\beta^n = \{\mathbf{x} \in \mathbb{R}^{n+1} | \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = -\beta\}$, where $\mathbf{x} = (x_0, x_1, \dots, x_n)$ is $n+1$ dimensional vector with $x_0 = \sqrt{\beta + \sum_{i=1}^n x_i^2} > 0$. $\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}} = -x_0 \cdot y_0 + \sum_{i=1}^n x_i \cdot y_i$ denotes the Lorentzian scalar product. Particularly, the squared Lorentzian distance between $\mathbf{x}, \mathbf{y} \in \mathbb{L}_\beta^n$ is defined as:

$$d_{\mathcal{L}}^2(\mathbf{x}, \mathbf{y}) = -2\beta - 2\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}. \quad (1)$$

Recent work [2, 15] proved that the squared Lorentzian distance could effectively depict hierarchical data.

For any $\mathbf{x} \in \mathbb{L}_\beta^n$, there exists a n -dimensional vector space $\mathcal{T}_{\mathbf{x}}\mathbb{L}_\beta^n$, referred to as the tangent space centered at $\mathbf{x}$. The exponential map $\exp_{\mathbf{x}}^\beta : \mathbb{L}_\beta^n \rightarrow \mathcal{T}_{\mathbf{x}}\mathbb{L}_\beta^n$ and the logarithmic map $\log_{\mathbf{x}}^\beta : \mathcal{T}_{\mathbf{x}}\mathbb{L}_\beta^n \rightarrow \mathbb{L}_\beta^n$ are used to perform mappings between the tangent space and hyperbolic space [22].

In this paper, we omit β and fix its value at $\beta = 1$. We use the origin $\mathbf{o}$ as a reference point to perform tangent space transformations $\exp_{\mathbf{o}}(\cdot)$ and $\log_{\mathbf{o}}(\cdot)$ between $\mathbb{L}^n$ and $\mathcal{T}_{\mathbf{o}}\mathbb{L}^n$. For more details about hyperbolic space, please refer to the survey paper [22].

Rotation. The rotation was originally proposed for knowledge graph embedding [27]. Since rotations are hyperbolic isometries [1], they can be applied to hyperbolic spaces. The rotation operation can be described by a block-diagonal matrix, which is defined as:

$$\text{Rot}_\Theta = \text{diag}\left(\mathbf{R}(\theta_{r,1}), \mathbf{R}(\theta_{r,2}), \dots, \mathbf{R}(\theta_{r,n/2})\right), \quad (2)$$

where $\Theta = \{\theta_{r,1}, \dots, \theta_{r,n/2}\}$ are the rotation parameters. The 2×2 block $\mathbf{R}(\theta_{r,i})$ is calculated as:

$$\mathbf{R}(\theta_{r,i}) = \begin{bmatrix} \cos(\theta_{r,i}) & -\sin(\theta_{r,i}) \\ \sin(\theta_{r,i}) & \cos(\theta_{r,i}) \end{bmatrix}. \quad (3)$$

In this work, we use $\text{Rot}(\mathbf{x})$ to denote applying rotation operation on a given hyperbolic vector $\mathbf{x} \in \mathbb{L}^n$, and all parameters of rotation matrices are learnable.

4 The Proposed Method

4.1 Overview of Framework

Figure 2 shows the framework of the proposed model HMST, which comprises four components. Firstly, we exploit user-POI interactions to extract user preferences in hyperbolic embeddings. Second, we model three kinds of one-hop transition semantics simultaneously with three learnable rotation matrices to capture global transition relations between two POIs. Third, we apply the Lorentzian attention mechanism to model dynamic dependencies in the check-in trajectory and combine its output with learned transition rotation matrices to generate trajectory representations with corresponding semantics. In this manner, both the hierarchical structures and




Figure 2: Illustration of the framework of HMST. HMST has four components. First, user-POI interaction captures user preference information. Second, multi-semantic one-hop transition models one-hop user movements with three semantics. Third, multi-semantic sequential transition captures dynamic dependencies and models continuous multi-hop movements with three semantics. Finally, multi-task prediction recommends the next possible POI, category, and region, respectively.

diverse transition semantics can be captured in two types of transitions, enabling comprehensive multi-semantic transitions to enhance recommendation accuracy. Lastly, we perform the next POI, category, and region recommendations by calculating the scores based on hyperbolic distances.

4.2 User-POI Interaction

In our work, each entity (user, POI, region, and category) is assigned a hyperbolic embedding vector in the Lorentz model $\mathbb{L}^n$.

We first capture user preference information from all observed historical user-POI interactions, i.e., all edges $\{(u, l) \mid (u, l) \in \mathcal{E}_{UL}\}$. Notably, user-POI interactions often exhibit power-law distributions, indicating implicit hierarchical structures that can be effectively captured through hyperbolic embeddings. Inspired by the classic network embedding methods [10, 28], we learn user-POI interactions in an unsupervised manner using a specially designed score function with rotation operations. For each user $u \in \mathcal{U}$, a preference rotation matrix Rot_u is assigned to more effectively capture the user's preference information. The key idea is that when a pair (u, l) is directly connected, their embeddings, after a rotation transformation, should be closer in the embedding space.

Specifically, for an observed edge $e_{u,l} = (u, l) \in \mathcal{E}_{UL}$, the user node has a hyperbolic embedding $\mathbf{u} \in \mathbb{L}^n$, and the POI node has an embedding $\mathbf{l} \in \mathbb{L}^n$. A score function $S(\cdot)$ is used to measure the

preference score between u and l as:

$$S(u, l) = -d_{\mathcal{L}}^2(\mathbf{u}^R, \mathbf{l}) + b_u + b_l, \quad (4)$$

where $\mathbf{u}^R = \text{Rot}_u(\mathbf{u})$ is the user embedding after rotation operation of Rot_u . The probability $\Pr(e_{u,l})$ of observing the edge (u, l) can be calculated using the score function $S(u, l)$ as follows:

$$\Pr(e_{u,l}) = \frac{\exp(S(u, l))}{\sum_{l_k \in \mathcal{L}} \exp(S(u, l_k))}. \quad (5)$$

Note that directly calculating Eq. (6) is computationally expensive, as all nodes in $\mathcal{L}$ must be enumerated. To address this, we employ the widely used negative sampling technique. This allows us to estimate Eq. (6) in logarithmic form as follows:

$$\log \Pr(e_{u,l}) \approx \log \sigma(S(u, l)) + \sum_{l_k \in \mathcal{N}} \log \sigma(-S(u, l_k)), \quad (6)$$

where $\sigma(x) = 1/(1 + e^{-x})$ is the Sigmoid function and $\mathcal{N}$ is the set of random sampled negative nodes.

The goal is to maximize the joint probability of all user-POI interactions to learn the user and POI embeddings, as well as the preference rotation parameters. Assuming that each edge $e_{u,l} \in \mathcal{E}_{UL}$ is independent, the log probability is calculated as follows:

$$O_{ul} = \sum_{e_{u,l} \in \mathcal{E}_{UL}} \log \Pr(e_{u,l}). \quad (7)$$

That is, we optimize O_{ul} to capture the preferences of all users from the user-POI interactions.

4.3 Multi-semantic One-hop Transition

We then model one-hop multi-semantic transitions using the edge set $\{(l_a, l_b) | (l_a, l_b) \in \mathcal{E}_{LL}\}$ to capture global user movements. Multi-semantic considers not only the transition between locations l_a and l_b , but also the transitions in their category (c_{l_a}, c_{l_b}) and region (r_{l_a}, r_{l_b}) information, helping to mitigate POI data sparsity by leveraging the less sparse category and region information.

We propose a novel approach using various rotation operations to capture location, category, and region transitions for an observed user movement $l_a \rightarrow l_b$ simultaneously. Here, rotation operations are adopted since they can flexibly represent various transition relationships while preserving latent hierarchical patterns.

Similar to the way for learning user-POI interactions, we further incorporate category and region transition information when maximizing the joint probability of observing all location transitions $\{(l_a, l_b) | (l_a, l_b) \in \mathcal{E}_{LL}\}$. We first model each semantic transition and then combine them for one-hop multi-semantic transition.

We use three trainable rotation matrices Rot_{ll} , Rot_{cc} , and Rot_{rr} to capture location, category, and region transition, respectively. Analogous to Eq. (4) and (6), the probability of POI-POI sequential transition in the log form can be calculated as:

$$\log \Pr(e_{l_a, l_b}) \approx \log \sigma(S(l_a, l_b)) + \sum_{l_k \in \mathcal{N}_{l_a}} \log \sigma(-S(l_a, l_k)), \quad (8)$$

$$S(l_a, l_b) = -d_{\mathcal{L}}^2(\mathbf{l}_a^R, \mathbf{l}_b) + b_{l_a} + b_{l_b}, \quad (9)$$

where $\mathbf{l}_a^R = \text{Rot}_{ll}(\mathbf{l}_a)$, b_{l_a} , b_{l_b} are POI specific biases. Similarly, the logarithmic probability of the category transition is calculated as:

$$\log \Pr(e_{c_a, c_b}) \approx \log \sigma(S(c_a, c_b)) + \sum_{c_k \in \mathcal{N}_{c_a}} \log \sigma(-S(c_a, c_k)), \quad (10)$$

$$S(c_a, c_b) = -d_{\mathcal{L}}^2(\mathbf{c}_a^R, \mathbf{c}_b) + b_{c_a} + b_{c_b}, \quad (11)$$

where, where $\mathbf{c}_a^R = \text{Rot}_{cc}(\mathbf{c}_a)$, b_{c_a} , b_{c_b} are category specific biases. The logarithmic probability of the region transition is

$$\log \Pr(e_{r_a, r_b}) \approx \log \sigma(S(r_a, r_b)) + \sum_{r_k \in \mathcal{N}_{r_a}} \log \sigma(-S(r_a, r_k)), \quad (12)$$

$$S(r_a, r_b) = -d_{\mathcal{L}}^2(\mathbf{r}_a^R, \mathbf{r}_b) + b_{r_a} + b_{r_b}, \quad (13)$$

where, where $\mathbf{r}_a^R = \text{Rot}_{rr}(\mathbf{r}_a)$, b_{r_a} , b_{r_b} are region specific biases. $\mathcal{N}_{ll}$, $\mathcal{N}_{cc}$, and $\mathcal{N}_{rr}$ are related sampled negative node sets. Hence, the overall one-hop multi-semantic transition for $l_a \rightarrow l_b$ can be calculated by combining probabilities across all semantics:

$$\log \Pr_{mst}(l_a \rightarrow l_b) = \log \Pr(e_{l_a, l_b}) + \log \Pr(e_{c_a, c_b}) + \log \Pr(e_{r_a, r_b}). \quad (14)$$

After that, the objective function to maximize the joint probability of observing all one-hop multi-semantic transitions is:

$$O_{mst} = \sum_{(l_a, l_b) \in \mathcal{E}_{LL}} \log \Pr_{mst}(l_a \rightarrow l_b). \quad (15)$$

We maximize O_{mst} to learn the embeddings and parameters of three rotation matrices. In this way, we effectively encode global one-hop transition relations with different semantics.

4.4 Multi-semantic Sequential Transition

One-hop transition is insufficient for making personalized next POI recommendations, as it fails to capture the complex dynamic dependencies within a user's check-in trajectory. To effectively predict users' future mobility, it is essential to explore sequential transition patterns across multiple consecutive locations. These dynamic dependencies help us better understand user movement patterns over time. Hence, we incorporate the sequential transitions to make final recommendations.

Unlike previous methods [41] that directly fuse all information, which lacks interpretability, our method leverages learned embeddings and various global transition semantics to refine the trajectory modeling process. Since trajectory transitions also involve location, category, and region information, it is natural to combine the mentioned semantics to enhance the final recommendations.

Given trajectory $s = \{q_1, q_2, \dots, q_m\}$ of a user u , we first fuse the check-in representations and then generate the trajectory embedding $\mathbf{s}$ which will be used for predict the next POI.

The check-in representation $\mathbf{q}_i \in \mathbb{L}^n$ of a check-in record $q_i = (u, l_i, c_i, r_i)$ can be represented as:

$$\mathbf{q}_i = \exp_{\mathbf{o}}(\log_{\mathbf{o}}(\mathbf{u}) + \log_{\mathbf{o}}(\mathbf{l}_i) + \log_{\mathbf{o}}(\mathbf{c}_i) + \log_{\mathbf{o}}(\mathbf{r}_i)), \quad (16)$$

where $\mathbf{u}$, $\mathbf{l}_i$, $\mathbf{c}_i$, and $\mathbf{r}_i$ represent the user, POI, category, and region embeddings, respectively, $\log(\cdot)$ and $\exp(\cdot)$ are mappings between the tangent space at space origin $\mathcal{T}_{\mathbf{o}}\mathbb{L}^n$ and the Lorentz model $\mathbb{L}^n$.

We extract and summarize the dependencies of historical check-ins in the trajectory by a Lorentzian attention module. Using the last check-in as the query, we integrate all the historical check-ins to calculate the trajectory embedding $\mathbf{s}$:

$$\mathbf{s} = \sqrt{\beta} \cdot \frac{\sum_{i=1}^m \alpha_i \mathbf{q}_i^v}{\|\sum_{i=1}^m \alpha_i \mathbf{q}_i^v\|_{\mathcal{L}}}, \quad (17)$$

$$\alpha_i = \frac{\exp\left(\frac{-d_{\mathcal{L}}^2(\mathbf{q}_m^q, \mathbf{q}_i^k)}{\sqrt{n}}\right)}{\sum_{k=1}^m \exp\left(\frac{-d_{\mathcal{L}}^2(\mathbf{q}_m^q, \mathbf{q}_k^k)}{\sqrt{n}}\right)}, \quad (18)$$

where $\mathbf{q}^q = \text{HL}^q(\mathbf{q})$, $\mathbf{q}^k = \text{HL}^k(\mathbf{q})$ and $\mathbf{q}^v = \text{HL}^v(\mathbf{q})$. Note that $\text{HL}(\cdot)$ denotes the Lorentzian linear layer [2].

Moreover, we combine $\mathbf{s}$ with various transition rotation matrices to obtain trajectory representations with different transition semantics. Specifically, we compute $\mathbf{s}_{ll} = \text{Rot}_{ll}(\mathbf{s})$, $\mathbf{s}_{cc} = \text{Rot}_{cc}(\mathbf{s})$, and $\mathbf{s}_{rr} = \text{Rot}_{rr}(\mathbf{s})$. Finally, using these trajectory representations, we integrate global one-hop transitions and sequential transition information with latent hierarchical features to make the final next POI recommendations.

4.5 Multi-task Prediction

After capturing transition information from both one-hop and sequential aspects, we predict the next location a user might visit. Unlike methods that focus solely on predicting the next POI, our approach also predicts potential transitions between categories and regions. Using the learned trajectory representations with specific semantics, we simultaneously predict the next POI, category, and region. In other words, we integrate the results of multiple tasks to make the final recommendation.

Table 1: Dataset statistics.

	#User	#POI	#Cat	#Check-in	#Trajectory
NYC	1,048	4,981	318	103,941	14,130
TKY	2,282	7,833	290	405,000	65,499
CA	3,957	9,690	296	238,369	45,123

For a given trajectory, we calculate the next POI recommendation score based on the hyperbolic distance between the trajectory representation and all POI embeddings. The next POI recommendation score for POI l_i is calculated as:

$$y_{l_i} = \exp \left(-d_{\mathcal{L}}^2(s_{ll}, l_i) \right) + b_{l_i}, \quad (19)$$

where b_{l_i} is the recommendation bias for l_i . Similarly, we can compute the recommendation score for the next category c_j and region r_k respectively as:

$$y_{c_j} = \exp \left(-d_{\mathcal{L}}^2(s_{cc}, c_j) \right) + b_{c_j}, \quad (20)$$

$$y_{r_k} = \exp \left(-d_{\mathcal{L}}^2(s_{rr}, r_k) \right) + b_{r_k}. \quad (21)$$

In practice, given a trajectory s , the recommendation scores for POI, category, and region are combined into vectors: $y_{ll} \in \mathbb{R}^{1 \times |\mathcal{L}|}$, $y_{cc} \in \mathbb{R}^{1 \times |\mathcal{C}|}$, and $y_{rr} \in \mathbb{R}^{1 \times |\mathcal{R}|}$, respectively. Accordingly, we calculate three cross-entropy losses O_l , O_c , and O_r to train our model's ability to predict the next POI, category, and region.

The overall loss function $\mathcal{O}$ is the combination of the user-POI interaction learning objective O_{ul} , one-hop multi-semantic transition objective O_{mst} and all recommendation objectives O_l , O_c , and O_r , which is calculated as:

$$\mathcal{O} = -(O_{ul} + O_{mst}) + O_l + O_c + O_r. \quad (22)$$

We minimize $\mathcal{O}$ to train our model, optimizing it using the efficient Riemannian Adam algorithm [14].

5 Experiments

5.1 Experimental Settings

5.1.1 Datasets. We conduct experiments on three commonly used real-word datasets **FourSquare-NYC** [38], **FourSquare-TKY** [38] and **Gowalla-CA** [46]. Each dataset is divided into training, validation, and test sets in chronological order with a ratio of 80% for training, 10% for validation, and 10% for testing. The statistics of datasets are shown in Table 1.

5.1.2 Experimental Setups. We use similar evaluation protocol as [41]. The curvature parameter β is set to 1. The default batch size is 128. The number of dimensions n is set to 64. The number of negative samples is set to 10. We use Top-k Accuracy Rates (Acc@k) and Mean Reciprocal Rank (MRR) as evaluation metrics. k is set to 5 and 10. Our method is implemented in Python on a Ubuntu server equipped with an Intel(R) i9-10900X 3.70GHz CPU and two NVIDIA GTX 4090 GPUs.

5.1.3 Baselines. We compare HMST with the following 9 methods to evaluate its effectiveness. The baselines can be categorized into two main classes:

(1) Euclidean methods: **LSTM** [12], **PLSPL** [35], **GETNext** [41], **AGRAN** [34], **CLSPRec** [7], **MCN4Rec** [17]. GETNext, AGRAN, CLSPRec, and MCN4Rec are four recent baselines.

(2) Hyperbolic methods: **HME** [9], **HICF** [39], **HGSR** [42]. Note that HICF and HGSR were not originally designed for POI recommendation, and we thus adapted them for our task. For HGSR, we establish user-user connections based on similarities of user preferences.

5.2 Experimental Results

5.2.1 Overall Performance. We first compare HMST with 9 baselines to investigate its capability for the next POI recommendation. The overall results are displayed in Table 2, showing the evaluation metrics on three datasets. The baselines are divided into two groups: models in Euclidean space and models in hyperbolic space. We have several observations. First, our proposed HMST model outperforms all Euclidean models. Note that the state-of-the-art POI recommendation methods are based on Euclidean spaces, such as the GETNext, AGRAN, CLSPRec, and MCN4Rec. These methods adopt graph and sequence modeling approaches, achieving significant improvements over previous methods. HMST builds on a similar principle and outperforms them due to two key aspects: our method comprehensively captures latent hierarchical patterns through hyperbolic representations and models various transition semantics via rotation operations. Second, our approach also demonstrates significant advantages when compared to methods based on hyperbolic space. Though HICF and HGSR are the latest hyperbolic recommendation methods, they are not specifically designed for POI recommendations and cannot incorporate the complex dependencies in user check-in trajectories, resulting in limited performance. Third, our model achieves much better results than the hyperbolic next POI recommendation approach HME. This is because HME focuses solely on modeling node embeddings in hyperbolic space, which is lacking in depicting multi-semantic transitions and capturing dynamic relationships effectively. Overall, HMST outperforms all baselines on three datasets, highlighting the effectiveness of using rotation operations in hyperbolic space for the next POI recommendation.

5.2.2 Ablation Study. In this section, we conduct a comprehensive ablation study to investigate the effectiveness of our design principles. Specifically, we conduct the following ablation settings: (1) **w/o hyperbolic**. Implement HMST in Euclidean space. (2) **w/o UPI**. Remove user-POI interaction. (3) **w/o MOP**. Remove the whole multi-semantic one-hop transition module. (4) **w/o MSP**. Remove the whole multi-semantic sequential transition module. (5) **W/O rotation**. Remove all rotation operations. (6) **w/o C&R**. Remove all category and region transition information.

The results of ablation studies are reported in Table 3 where metrics Acc@5 and MRR are reported. Taking the complete HMST as the baseline, we make the following observations. The result **w/o hyperbolic** shows that both metrics drop significantly because the underlying hierarchical structures cannot be effectively captured in Euclidean space, indicating the advantages of hyperbolic space. The result **w/o UPI** highlights the importance of modeling user preferences. When user preferences are not effectively captured, the model's performance declines significantly. The result

Table 2: Overall performance comparison in Acc@k and MRR on three datasets. The best scores are in boldface and the second best scores are underlined.

	NYC			TKY			CA		
	Acc@5	Acc@10	MRR	Acc@5	Acc@10	MRR	Acc@5	Acc@10	MRR
LSTM	0.3276	0.3895	0.2299	0.3276	0.4087	0.2431	0.2100	0.2641	0.1442
PLSPL	0.4399	0.5323	0.2961	0.3983	0.4804	0.2934	0.2642	0.3343	0.1847
GETNext	<u>0.4997</u>	<u>0.6017</u>	<u>0.3478</u>	<u>0.4432</u>	<u>0.5383</u>	<u>0.3148</u>	<u>0.3050</u>	0.3853	<u>0.2073</u>
AGRAN	0.4598	0.5508	0.3175	0.4075	0.4967	0.2907	0.2859	<u>0.3879</u>	0.1990
CLSPRec	0.4252	0.5018	0.3132	0.3287	0.3721	0.2463	0.2595	0.3432	0.1917
MCN4Rec	0.4577	0.5368	0.3297	0.4278	0.5052	0.3047	0.2833	0.3589	0.1946
HME	0.4120	0.5409	0.2746	0.3538	0.4460	0.2241	0.2788	0.3650	0.1869
HICF	0.3447	0.4727	0.2161	0.2597	0.3407	0.1677	0.2534	0.3367	0.1675
HGSR	0.4437	0.5512	0.2908	0.3274	0.4542	0.2264	0.2446	0.3214	0.1690
HMST	0.5388	0.6233	0.3567	0.4714	0.5633	0.3315	0.3287	0.4200	0.2309

Table 3: The results of ablation study on NYC and CA datasets. The best scores are in boldface.

	NYC		CA	
	Acc@5	MRR	Acc@5	MRR
(1) w/o hyperbolic	0.4976	0.3270	0.2945	0.2150
(2) w/o UPI	0.5231	0.3493	0.3185	0.2190
(3) w/o MOP	0.5182	0.3461	0.2888	0.2077
(4) w/o MSP	0.4477	0.3173	0.3105	0.2114
(5) w/o rotation	0.5190	0.3317	0.3172	0.2273
(6) w/o C&R	0.4765	0.3237	0.3120	0.2253
HMST	0.5388	0.3567	0.3287	0.2309

w/o MOP is associated with the multi-semantic one-hop transition module. Removing the modeling of global one-hop transitions causes a drop in both metrics, indicating that global POI transitions effectively capture movement trends. Similarly, the result **w/o MSP** is related to the multi-semantic sequential transition module. When the entire sequential transition modeling is removed, all metrics significantly decline. This highlights the importance of dynamic dependencies and sequential transition semantics in user trajectories for personalized POI recommendations. Moreover, the result **w/o rotation** demonstrates that rotation operations in hyperbolic space effectively capture complex transition semantics. The final ablation setting, **w/o C&R**, represents the results without incorporating additional transition semantics. The performance decline observed across both datasets highlights the importance of fully modeling transition semantics for the next POI recommendation tasks.

In summary, the results of the ablation study suggest that the best experimental outcomes are achieved when the model is complete. This underscores the effectiveness of each individual module for the next POI recommendation.

6 Conclusion

In this study, we explore the use of rotation operations in hyperbolic space to address limitations in existing POI recommendation methods. Our findings demonstrate that rotation transformations

effectively capture multiple transition semantics, enabling the modeling of complex transition patterns in check-in data. Based on this, we propose a novel hyperbolic rotation-based recommendation model, HMST, marking the first exploration of hyperbolic rotations for the next POI recommendation tasks. Our approach first incorporates user preference information by learning user-POI interactions to construct expressive user and POI hyperbolic embeddings. It then applies rotation operations to model location, category, and region transition semantics in both one-hop and sequential transitions. The final prediction module simultaneously predicts the next POI, category, and region. Extensive experiments on three real-world datasets demonstrate the effectiveness of our proposed model in capturing complex transition semantics and achieving accurate recommendation results.

References

- [1] Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. 2020. Low-Dimensional Hyperbolic Knowledge Graph Embeddings. In *ACL*.
- [2] Weize Chen, Xu Han, Yankai Lin, Hexu Zhao, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2022. Fully Hyperbolic Neural Networks. In *ACL*.
- [3] Yankai Chen, Menglin Yang, Yingxue Zhang, Mengchen Zhao, Ziqiao Meng, Jianye Hao, and Irwin King. 2022. Modeling scale-free graphs with hyperbolic geometry for knowledge-aware recommendation. In *WSDM*.
- [4] Chen Cheng, Haiqin Yang, Michael R Lyu, and Irwin King. 2013. Where you like to go next: Successive point-of-interest recommendation. In *IJCAI*.
- [5] Yao Dong, Lei Wang, Ji Xiang, Xiaobo Guo, and Yuqiang Xie. 2022. RotateCT: Knowledge Graph Embedding by Rotation and Coordinate Transformation in Complex Space. In *COLING*.
- [6] Yuntao Du, Xinjun Zhu, Lu Chen, Baihua Zheng, and Yunjun Gao. 2022. Hagk: Hierarchy-aware knowledge gated network for recommendation. In *SIGIR*.
- [7] Chenghua Duan, Wei Fan, Wei Zhou, Hu Liu, and Junhao Wen. 2023. CLSPRec: Contrastive Learning of Long and Short-term Preferences for Next POI Recommendation. In *CIKM*.
- [8] Shanshan Feng, Lisi Chen, Kaiqi Zhao, Wei Wei, Xuemeng Song, Shuo Shang, Panos Kalnis, and Ling Shao. 2022. ROLE: Rotated Lorentzian Graph Embedding Model for Asymmetric Proximity. *TKDE* (2022).
- [9] Shanshan Feng, Lucas Vinh Tran, Gao Cong, Lisi Chen, Jing Li, and Fan Li. 2020. Hme: A hyperbolic metric embedding approach for next-poi recommendation. In *SIGIR*.
- [10] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *KDD*. 855–864.
- [11] Naicheng Guo, Xiaolei Liu, Shaoshuai Li, Mingming Ha, Qiongxi Ma, Binfeng Wang, Yunan Zhao, Linxun Chen, and Xiaobo Guo. 2023. Hyperbolic Contrastive Graph Representation Learning for Session-based Recommendation. *TKDE* (2023).

- [12] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* (1997).
- [13] Zheng Huang, Jing Ma, Yushun Dong, Natasha Zhang Foutz, and Jundong Li. 2022. Empowering next poi recommendation with multi-relational modeling. In *SIGIR*.
- [14] Max Kochurov, Rasul Karimov, and Serge Kozlukov. 2020. Geoopt: Riemannian optimization in pytorch. *arXiv preprint arXiv:2005.02819* (2020).
- [15] Marc Law, Renjie Liao, Jake Snell, and Richard Zemel. 2019. Lorentzian Distance Learning for Hyperbolic Representations. In *ICML*.
- [16] Anchen Li, Bo Yang, Huan Huo, Hongxu Chen, Guandong Xu, and Zhen Wang. 2022. Hyperbolic neural collaborative recommender. *TKDE* (2022).
- [17] Shuzhe Li, Wei Chen, Bin Wang, Chao Huang, Yanwei Yu, and Junyu Dong. 2024. MCN4Rec: Multi-level Collaborative Neural Network for Next Location Recommendation. *TOIS* (2024).
- [18] Defu Lian, Yongji Wu, Yong Ge, Xing Xie, and Enhong Chen. 2020. Geography-aware sequential location recommendation. In *KDD*.
- [19] Nicholas Lim, Bryan Hooi, See-Kiong Ng, Yong Liang Goh, Renrong Weng, and Rui Tan. 2022. Hierarchical multi-task graph recurrent network for next poi recommendation. In *SIGIR*.
- [20] Nicholas Lim, Bryan Hooi, See-Kiong Ng, Yong Liang Goh, Renrong Weng, and Rui Tan. 2023. Learning Hierarchical Spatial Tasks with Visiting Relations for Next POI Recommendation. *TORS* (2023).
- [21] Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. Predicting the next location: A recurrent model with spatial and temporal contexts. In *AAAI*.
- [22] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. 2021. Hyperbolic deep neural networks: A survey. *PAMI* (2021).
- [23] Hongliang Qiao, Shanshan Feng, Xutao Li, Huiwei Lin, Han Hu, Wei Wei, and Yunming Ye. 2023. RotDiff: A Hyperbolic Rotation Representation Model for Information Diffusion Prediction. In *CIKM*.
- [24] Ali Sadeghian, Mohammadreza Armandpour, Anthony Colas, and Daisy Zhe Wang. 2021. Chronor: Rotation based temporal knowledge graph embedding. In *AAAI*.
- [25] Pablo Sánchez and Alejandro Bellogin. 2022. Point-of-interest recommender systems based on location-based social networks: a survey from an experimental perspective. *ACM Computing Surveys (CSUR)* (2022).
- [26] Jianing Sun, Zhaoyue Cheng, Saba Zuberi, Felipe Pérez, and Maksims Volkovs. 2021. Hgcf: Hyperbolic graph convolution networks for collaborative filtering. In *WWW*.
- [27] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space. In *ICLR*.
- [28] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. Line: Large-scale Information Network Embedding. In *Proceedings of the 24th international conference on world wide web*. 1067–1077.
- [29] Lucas Vinh Tran, Yi Tay, Shuai Zhang, Gao Cong, and Xiaoli Li. 2020. HyperML: A Boosting Metric Learning Approach in Hyperbolic Space for Recommender Systems. *WSDM* (2020).
- [30] En Wang, Yiheng Jiang, Yuanbo Xu, Liang Wang, and Yongjian Yang. 2022. Spatial-temporal interval aware sequential POI recommendation. In *ICDE*.
- [31] Hao Wang, Defu Lian, Hanghang Tong, Qi Liu, Zhenya Huang, and Enhong Chen. 2021. Hypersorec: Exploiting hyperbolic user and item representations with multiple aspects for social-aware recommendation. *TOIS* (2021).
- [32] Shicheng Wang, Shu Guo, Lihong Wang, Tingwen Liu, and Hongbo Xu. 2023. HDNR: A Hyperbolic-Based Debiased Approach for Personalized News Recommendation. In *SIGIR*.
- [33] Zhaobo Wang, Yanmin Zhu, Haobing Liu, and Chunyang Wang. 2022. Learning graph-based disentangled representations for next POI recommendation. In *SIGIR*.
- [34] Zhaobo Wang, Yanmin Zhu, Chunyang Wang, Wenze Ma, Bo Li, and Jiadi Yu. 2023. Adaptive Graph Representation Learning for Next POI Recommendation. In *SIGIR*.
- [35] Yuxia Wu, Ke Li, Guoshuai Zhao, and Xueming Qian. 2020. Personalized long-and short-term preference learning for next POI recommendation. *TKDE* (2020).
- [36] Jiayi Xie and Zhenzhong Chen. 2023. Hierarchical transformer with spatio-temporal context aggregation for next point-of-interest recommendation. *TOIS* (2023).
- [37] Xiaodong Yan, Tengwei Song, Yifeng Jiao, Jianshan He, Jiaotuan Wang, Ruopeng Li, and Wei Chu. 2023. Spatio-Temporal Hypergraph Learning for Next POI Recommendation. In *SIGIR*.
- [38] Dingqi Yang, Daqing Zhang, Vincent W Zheng, and Zhiyong Yu. 2014. Modeling user activity preference by leveraging user spatial temporal characteristics in LBSNs. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* (2014).
- [39] Menglin Yang, Zhihao Li, Min Zhou, Jiahong Liu, and Irwin King. 2022. Hicf: Hyperbolic informative collaborative filtering. In *KDD*.
- [40] Menglin Yang, Min Zhou, Jiahong Liu, Defu Lian, and Irwin King. 2022. HRCF: Enhancing collaborative filtering via hyperbolic geometric regularization. In *WWW*.
- [41] Song Yang, Jiamou Liu, and Kaiqi Zhao. 2022. GETNext: trajectory flow map enhanced transformer for next POI recommendation. In *SIGIR*.
- [42] Yonghui Yang, Le Wu, Kun Zhang, Richang Hong, Hailin Zhou, Zhiqiang Zhang, Jun Zhou, and Meng Wang. 2023. Hyperbolic Graph Learning for Social Recommendation. *TKDE* (2023).
- [43] Feiyu Yin, Yong Liu, Zhiqi Shen, Lisi Chen, Shuo Shang, and Peng Han. 2023. Next POI recommendation with dynamic graph and explicit dependency. In *AAAI*.
- [44] Fuqiang Yu, Lizhen Cui, Wei Guo, Xudong Lu, Qingzhong Li, and Hua Lu. 2020. A Category-Aware Deep Model for Successive POI Recommendation on Sparse Check-in Data. In *WWW*.
- [45] Meng Yuan, Fuzhen Zhuang, Zhao Zhang, Deqing Wang, and Jin Dong. 2023. Knowledge-based Multiple Adaptive Spaces Fusion for Recommendation. In *RecSys*.
- [46] Quan Yuan, Gao Cong, Zongyang Ma, Aixin Sun, and Nadia Magnenat Thalmann. 2013. Time-aware point-of-interest recommendation. In *SIGIR*.
- [47] Hongyu Zang, Dongcheng Han, Xin Li, Zhifeng Wan, and Mingzhong Wang. 2021. Cha: Categorical hierarchy-based attention for next poi recommendation. *TOIS* (2021).
- [48] Lu Zhang, Zhu Sun, Ziqing Wu, Jie Zhang, Yew Soon Ong, and Xinghua Qu. 2022. Next point-of-interest recommendation with inferring multi-step future preferences. In *IJCAI*.
- [49] Lu Zhang, Zhu Sun, Jie Zhang, Horst Kloeden, and Felix Klanner. 2020. Modeling hierarchical category transition for next POI recommendation with uncertain check-ins. *Information Sciences* (2020).
- [50] Pengpeng Zhao, Anjing Luo, Yanchi Liu, Jiajie Xu, Zhixu Li, Fuzhen Zhuang, Victor S Sheng, and Xiaofang Zhou. 2020. Where to go next: A spatio-temporal gated network for next poi recommendation. *TKDE* (2020).