

Efficient Methods for Accurate Sparse Trajectory Recovery and Map Matching

Wei Tian

Department of Computing
Hong Kong Polytechnic University
wei.tian@connect.polyu.hk

Jieming Shi*

Department of Computing
Smart Cities Research Institute
Hong Kong Polytechnic University
jieming.shi@polyu.edu.hk

Man Lung Yiu

Department of Computing
Hong Kong Polytechnic University
csmlyiu@comp.polyu.edu.hk

Abstract—Real-world trajectories are often sparse with low-sampling rates (i.e., long intervals between consecutive GPS points) and misaligned with road networks, yet many applications demand high-quality data for optimal performance. To improve data quality with sparse trajectories as input, we systematically study two related research problems: *trajectory recovery on road network*, which aims to infer missing points to recover high-sampling trajectories, and *map matching*, which aims to map GPS points to road segments to determine underlying routes. Capturing latent patterns in complex sparse trajectory data on road networks is challenging, especially with large-scale datasets.

In this paper, we present efficient methods TRMMA and MMA for accurate trajectory recovery and map matching, respectively, where MMA serves as the first step of TRMMA. In MMA, we carefully formulate a classification task to map a GPS point from sparse trajectories to a road segment over a small candidate segment set, rather than the entire road network. We develop techniques in MMA to generate effective embeddings that capture the patterns of GPS data, directional information, and road segments, to accurately align sparse trajectories to routes. For trajectory recovery, TRMMA focuses on the segments in the route returned by MMA to infer missing points with position ratios on road segments, producing high-sampling trajectories efficiently by avoiding evaluation of all road segments. Specifically, in TRMMA, we design a dual-transformer encoding process to cohesively capture latent patterns in trajectories and routes, and an effective decoding technique to sequentially predict the position ratios and road segments of missing points. We conduct extensive experiments to compare TRMMA and MMA with numerous existing methods for trajectory recovery and map matching, respectively, on 4 large real-world datasets. TRMMA and MMA consistently achieve the best result quality, often by a significant margin. Moreover, TRMMA and MMA are highly efficient during training and inference, being up orders of magnitude faster than the next best competitors. The implementation is at <https://github.com/derekwtian/TRMMA>.

Index Terms—Trajectory Recovery, Map Matching

I. INTRODUCTION

GPS trajectories, i.e., sequences of timestamped GPS points, are increasingly produced on road networks due to the widespread use of mobile devices. High-quality trajectory data are crucial for many applications, including vehicle navigation [1], [2], travel time estimation [3], [4], trajectory similarity search [5]–[7], and traffic flow predication [8], [9].

*Corresponding author.

Ideally, trajectories should have high sampling rates with short intervals between GPS points (e.g., every 15 seconds) [10] and be aligned with road segments [11]. However, real trajectories are often sparse and of low quality, characterized by (i) low sampling rates with long intervals between consecutive GPS points (e.g., every 2–6 minutes [10], [12]), and (ii) inaccurate GPS coordinates deviating from road segments due to device or signal limits [11], [13]. Sparse trajectories lose significant travel information on road networks, negatively impacting the aforementioned applications.

To mitigate the issues of sparse trajectories and improve data quality, we study two closely related research problems in a principled way: *trajectory recovery on road network* [10], [14] and *map matching* [13]. Briefly, map matching seeks to identify the segments of GPS points and find the route of a trajectory, while trajectory recovery is to infer the missing points of a sparse trajectory to achieve a target high-sampling rate. These two tasks enhance data quality and are important in applications such as navigation [15], travel time estimation [4], and traffic flow analysis [8]. We use Fig. 1 to illustrate the two tasks. In the road network G shown in Fig. 1, the three GPS points p_1, p_2, p_3 (black dots) form a sparse trajectory T with a low-sampling rate. Observe that the GPS points are not on any segment. Map matching aims to identify T 's underlying route \mathcal{R} , a sequence of road segments e_1, e_2, e_4, e_6 (blue arrows). Given a target high-sampling rate ϵ , trajectory recovery on road network is to infer the missing points of sparse trajectory T . The inferred points a_1, a_2, a_3, a_4, a_5 (red crosses) are aligned onto road segments with position ratios indicating their relative positions. These inferred points form a high-sampling trajectory \mathcal{T}_ϵ for T . Please see Section II for detailed definitions and discussion.

Effective map-matching and trajectory recovery are both challenging, especially with large datasets containing millions of trajectories on road networks. Capturing the latent patterns in GPS points, trajectory sequences, and road segments is non-trivial and computationally expensive. Early map matching methods [16], [17] mainly adopt Hidden Markov Model (HMM) with heuristics to identify the route of a GPS trajectory, which suffer from degraded performance, especially over sparse GPS trajectories [17]. Recent map matching studies [11], [13], [18] develop end-to-end methods incorporating

learning techniques to extract patterns from historical trajectories. Nevertheless, the map matching methods can only find routes but not increase the sampling rates of trajectories.

As for trajectory recovery on road network, it has recently attracted research attention [10], [14], [19]. For instance, RNTrajRec [10] considers road network structure and the surrounding subgraphs of observed GPS points, utilizing grids and graph neural networks, for trajectory recovery. However, these methods [10], [14] consider all segments in a road network as candidates to infer missing points, leading to sub-optimal performance and intensive computational overhead. Other approaches [20]–[22] work on trajectory recovery in free space as a grid and tend to underperform when adapted to road networks. Despite the close relationship between trajectory recovery and map matching, existing studies are often dedicated to one problem.

Given a sparse trajectory T , a natural methodology for trajectory recovery on a road network G is to first use map matching as a procedure to align the GPS points in T onto road segments in G , obtaining route \mathcal{R} , and subsequently recover the missing points over \mathcal{R} . The literature contains rough discussions on this idea but predominantly emphasizes its ineffectiveness [10], [19], caused by the severe sparsity of input raw GPS trajectories. However, we believe this critique warrants further thorough investigation.

Hence, in this paper, we solve the two problems together, fulfilling the potential of the promising but undervalued idea above, and present efficient and accurate methods: TRMMA for Trajectory Recovery on road networks, with MMA for Map Matching as a procedure invoked in TRMMA.

For a sparse trajectory T , TRMMA starts by invoking MMA to map the GPS points in T to their segments in a road network G . The entire G may contain numerous segments and trajectory T is sparse, posing great challenges for MMA to be accurate and fast. To address these challenges, in MMA, we conduct empirical analysis and formulate the task of identifying the segment of a GPS point p_i as a classification problem over a small candidate segment set, rather than the entire G . We then design techniques to produce effective embeddings for p_i and its candidate segments by holistically considering GPS data, directional information, road network context, and the sequential patterns of T . The embeddings are used to map a GPS point to a segment in MMA. Over the predicted segments of the points in T , MMA deduces the underlying route \mathcal{R} . MMA achieves high accuracy in real-world datasets, serving as a solid basis for TRMMA to conduct trajectory recovery. Specifically, TRMMA restores the missing points along the route \mathcal{R} of T to produce a high-sampling \mathcal{T}_ϵ that satisfies a specified sampling rate. Note that TRMMA concentrates solely on the segments in \mathcal{R} to infer missing points, avoiding the expensive evaluation of all segments within road network G , thereby improving efficiency. TRMMA requires only the sparse trajectory T and its route \mathcal{R} as inputs. In TRMMA, we design a dual-transformer encoding module to produce effective embeddings that capture the patterns in T and \mathcal{R} , synergized via an attention mechanism. Leveraging




Fig. 1: Example of Trajectory Recovery on Road Network and Map Matching

the embeddings, TRMMA uses a multitask decoding process to sequentially predict missing points, with classification over the candidate segments in \mathcal{R} and regression to estimate position ratios on these segments. We conduct comprehensive experiments across large-scale trajectory datasets on road networks, showing that TRMMA and MMA consistently surpass existing methods for trajectory recovery and map matching, respectively, in terms of result quality and efficiency. For instance, on a large PT dataset, TRMMA only needs 0.88 seconds per 1000 trajectory recoveries, whereas the competitor with runner-up quality costs 18.17 seconds; TRMMA trains in 5.49 minutes per epoch, while the competitor needs 109.7 minutes per epoch.

In summary, we make the following contributions.

- We methodically study two related research problems, map matching and trajectory recovery on road networks, and develop efficient and accurate methods MMA and TRMMA.
- In MMA for map matching, we formulate the task of mapping a GPS point to a segment as a classification problem over a small set of candidate segments, and develop novel embedding techniques to align sparse trajectories to routes.
- For sparse trajectory recovery, TRMMA focuses on the segments in the routes returned by MMA and captures the intrinsic patterns of trajectories and routes via effective encoding and decoding processes to predict missing points.
- The superiority of TRMMA and MMA, in terms of efficiency and effectiveness, is evaluated over massive trajectory data on road networks.

II. PROBLEM FORMULATION

Definition 1 (Road Network): A road network is modeled as a directed graph $G = (V, E)$, where V is the set of nodes, and E is the set of directed edges. A node $v \in V$ represents an intersection or a road end. A directed edge $e = (u, v) \in E$ is a road segment from entrance node u to exit node v . Let $n = |E|$ and $m = |V|$ denote the number of road segments and intersections, respectively.

Definition 2 (Trajectory): A trajectory T is a sequence of GPS points with timestamps, i.e., $T = \langle p_1, p_2, \dots, p_\ell \rangle$, where ℓ is the sequence length. A GPS point is $p_i = \langle lat, lng, t_i \rangle$ consisting of latitude lat , longitude lng and timestamp t_i .

Fig. 1 shows a trajectory $T = \langle p_1, p_2, p_3 \rangle$ with three GPS points (black dots) and a road network G . Note that a GPS point p_i may not lie exactly on any road segment in G .

Definition 3 (Route): A route \mathcal{R} on road network G is a sequence of road segments forming a *path* on G . Typically, the consecutive road segments in \mathcal{R} are different.

Fig. 1 shows a route $\mathcal{R} = \langle e_1, e_2, e_4, e_6 \rangle$ (blue arrows). The problem of map matching is defined below.

Definition 4 (Map Matching [13]): Given a road network G and a trajectory T , map matching is to map the GPS points in T onto the road segments of G to find the route \mathcal{R} of T .

In Fig. 1, route $\mathcal{R} = \langle e_1, e_2, e_4, e_6 \rangle$ is the map-matched result of trajectory T , where p_1, p_2, p_3 are mapped to segments e_1, e_4, e_6 respectively, with e_2 connecting e_1 and e_4 .

In what follows, we explain the definitions related to trajectory recovery on road network. A point on a road segment, referred to as a map-matched point, is defined as follows.

Definition 5 (Map-matched point [14]): A map-matched point $a_j = (e, r, t)$ denotes a point on segment e at time t , with position ratio $r \in [0, 1)$ representing a_j 's relative position from the entrance of segment $a_j.e$ to its total length.

Fig. 1 shows 5 map-matched points (red crosses), e.g., $a_3 = \langle e_4, 0.4, t \rangle$ on segment e_4 with position ratio $r = 0.4$. Given the GPS coordinates of the entrance and exit of a segment $a_j.e$, it is easy to get the GPS of a_j by interpolation.

A *sampling rate* ϵ is the time interval between two consecutive points of a trajectory. A *sparse trajectory* has a low sampling rate, i.e., long intervals between consecutive points. For example, trajectory $T = \langle p_1, p_2, p_3 \rangle$ shown as black dots in Fig. 1 illustrates a sparse trajectory.

Definition 6 (Map-Matched ϵ -Sampling Trajectory \mathcal{T}_ϵ [10], [14]): A map-matched trajectory with ϵ -sampling rate, \mathcal{T}_ϵ , is a sequence of map-matched points in Definition 5, i.e., $\mathcal{T}_\epsilon = \langle a_1, a_2, \dots, a_{\ell_\epsilon} \rangle$, where $a_j = (e, r, t)$ and time interval between consecutive points satisfies $a_{j+1}.t - a_j.t = \epsilon$, for $1 \leq j < \ell_\epsilon$.

In Fig. 1, $\mathcal{T}_\epsilon = \langle a_1, a_2, a_3, a_4, a_5 \rangle$ serves as an example of map-matched ϵ -sampling trajectory. Then, trajectory recovery on road network is defined below.

Definition 7 (Trajectory Recovery on Road Network [10], [14]): Given a sparse GPS trajectory $T = \langle p_1, p_2, \dots, p_\ell \rangle$, and a target high sampling rate ϵ , trajectory recovery aims to recover the map-matched ϵ -sampling trajectory $\mathcal{T}_\epsilon = \langle a_1, a_2, \dots, a_{\ell_\epsilon} \rangle$.

For instance, in Fig. 1, for the sparse trajectory $T = \langle p_1, p_2, p_3 \rangle$, the goal is to recover its map-matched ϵ -sampling trajectory $\mathcal{T}_\epsilon = \langle a_1, a_2, a_3, a_4, a_5 \rangle$.

Differences between the Two Tasks. Trajectory recovery involves inferring missing points of a sparse trajectory to meet the target ϵ sampling rate and matching all points onto the road network G simultaneously. In Fig. 1, a_2 and a_3 in \mathcal{T}_ϵ are the inferred points to satisfy the ϵ sampling rate, while a_1, a_4, a_5 are map-matched for GPS points p_1, p_2, p_3 of the sparse trajectory T . In contrast, map matching aims to find the route $\mathcal{R} = \langle e_1, e_2, e_4, e_6 \rangle$ of T . Another difference is that the route \mathcal{R} from map matching is a path on G from a source to a destination, whereas segments of map-matched points in \mathcal{T}_ϵ can be disconnected, with consecutive points possibly on the same segment. For instance, in Fig. 1, $\mathcal{T}_\epsilon = \langle a_1, a_2, a_3, a_4, a_5 \rangle$ has segment sequence $\langle e_1, e_4, e_4, e_4, e_6 \rangle$, where e_1 and e_4 are

TABLE I: Frequently used notations

Notation	Description
\mathcal{D}	Historical trajectory data \mathcal{D} .
$G = (V, E), n$	A road network G with road segment set E and intersection set V . n is the number of segments $ E $.
$e = (u, v)$	A road segment e from node u to node v .
p_i	A GPS point with coordinate and timestamp
$a_i = (e, r, t)$	A map-matched point on segment e with position ratio r and time t
$T = \langle p_1, \dots, p_\ell \rangle$	A sparse trajectory with ℓ GPS points.
$\mathcal{R}, \ell_{\mathcal{R}}$	A route, i.e., a path with $\ell_{\mathcal{R}}$ segments.
ϵ	The target sampling rate
$\mathcal{T}_\epsilon = \langle a_1, \dots, a_{\ell_\epsilon} \rangle$	The recovered map-matched ϵ -sampling trajectory with sequence length ℓ_ϵ .
\mathcal{C}_{p_i}, k_c	The candidate segment set of GPS point p_i with size k_c .
$P(c_j p_i)$	The probability that c_j is the segment of p_i
$\mathbf{M}, \mathbf{M}[i, :], \mathbf{v}$	A matrix \mathbf{M} , its i -th row vector $\mathbf{M}[i, :]$, and a vector \mathbf{v} .

not connected and a_2, a_3, a_4 are all on segment e_4 . Table I shows the frequently used notations in the paper.

III. RELATED WORK

Trajectory Recovery on Road Network. While trajectory recovery on road networks is related to map matching, recent studies often develop specialized end-to-end methods by learning from historical trajectories [10], [14], [19]. MTrajRec [14] introduces a sequence-to-sequence model with multitask learning to interpolate missing points in sparse trajectories while adhering to road network constraints. RNTrajRec [10] considers road network structure and the surrounding subgraphs of observed GPS points for trajectory recovery. MM-STGED [19] presents a graph-based encoder-decoder for trajectory recovery. The methods in [10], [14] typically project the predicted locations of missing points across all segments in a road network, which can be large. Instead, our method TRMMA only recovers the missing points on the segments in the route returned by our map-matching method MMA. The segments in the route are much fewer than the segments in G , rendering the efficiency of our designs, as demonstrated in the experiments.

Trajectory Recovery in Free Space. Another set of studies focuses on recovering trajectories in free space [20]–[26], such as social media check-in data, without the constraints of road networks. Elshrif et al. [23] introduce a heuristic search algorithm to estimate locations between two distant consecutive points within a trajectory. DHTR [20] utilizes a BiLSTM with spatial and temporal attention mechanisms and incorporates a Kalman filter-based calibration component to reconstruct missing GPS points. Chen et al. [21] develop TERI, a method for trajectory recovery with irregular time intervals. Xia et al. [24] employ self-attention mechanisms within and across trajectories to infer unobserved locations. Sun et al. [25] integrate a graph neural networks to model individual trajectories to capture transition patterns. These studies focus on free space, and exhibit suboptimal performance when adapted to road network constraints.

Map Matching. There is a rich collection of studies on map matching [11], [13], [16]–[18], [27]–[33]. A pioneer study [17] leverages HMM to find the most likely road route. Zheng et al. [29] introduce a history-based route inference system. Lou et al. [27] construct a candidate graph using the spatial and topological structures of the road network and the temporal features of trajectories. These methods usually exhibit degraded performance with sparse GPS trajectory data [17]. FMM [28] enhances HMM with a set of acceleration mechanisms. A recent trend is to extract patterns from historical trajectories to conduct map matching [11], [13], [18], [32], [33]. LHMM [11] enhances the HMM model by incorporating knowledge obtained from neural networks into learned probabilities, while DMM [18] adopts a recurrent neural network with a reinforcement learning scheme. GraphMM [13] is a graph-centric approach to incorporate graph neural networks and conditional models to leverage road and trajectory graph topology for map matching. DeepMM [32] is an end-to-end deep learning method with statistical and graph-based data augmentation techniques. However, these methods incur significant computational costs for training and inference, as shown in experiments. Overall, existing map matching methods struggle with effectiveness, efficiency, or both when dealing with sparse trajectories.

Trajectory Representation Learning. A plethora of studies have focused on embedding trajectory data into low-dimensional vectors [34]–[40]. These embeddings encapsulate both the spatial and sequential characteristics of trajectories and can serve as input of trajectory recovery tasks. TrajGAT [37] introduces a graph-based attention to capture long-term dependencies by representing each trajectory as a graph. TrajCL [38] proposes a dual-feature, self-attention-based trajectory encoder that adaptively incorporates both structural and spatial features. ST2Vec [39] produces time-aware trajectory representations by integrating spatial and temporal elements. In our experiments, we adopt these methods [37]–[39] with a decoding process for trajectory recovery. However, these approaches often yield suboptimal quality in the experiments, highlighting the need for specialized technical designs tailored to trajectory recovery on road networks.

IV. MMA: MAP MATCHING

As mentioned, our trajectory recovery method TRMMA utilizes the route \mathcal{R} of trajectory T obtained by our map matching method MMA to recover the map-matched ϵ -sampling trajectory \mathcal{T}_ϵ of T . Hence, we first present MMA in this section, while developing TRMMA in Section V.

Given a sparse trajectory $T = \langle p_1, \dots, p_\ell \rangle$, the major task of MMA is mapping every GPS point p_i to its corresponding segment e_i in road network G , such as mapping p_2 to e_4 in Fig. 1. Due to the sparsity of T , this task is non-trivial. After conducting an empirical analysis, we formulate it as a classification problem over a small candidate set of segments near p_i , rather than expensively considering all segments in G (Section IV-A). Then we develop point embedding and candidate segment embedding techniques in MMA to exploit




Fig. 2: The ratio of GPS points with their segments in their top- k_c nearest segments when varying k_c from 1 to 10.

the sequential nature of T , the properties of candidate segments, and the characteristics of road network for effective classification (Section IV-B). Once the GPS points of T are accurately mapped onto segments, MMA adopts elementary route planning methods to link disconnected segments into a route \mathcal{R} to return.

A. MMA Formulation

Intuitively, the road segment where p_i is located should be in its vicinity, making it unnecessary to consider all road segments in G , which can be numerous. However, the nearest segment to p_i may not always be the one where p_i actually resides. We conduct an empirical analysis to validate.

Given historical trajectory data \mathcal{D} for training, for every GPS point p_i in every trajectory in \mathcal{D} , we obtain its top- k_c nearest segments in road network G , and calculate the ratio of all points with their ground-truth road segment among their top- k_c nearest segments. Fig. 2 reports the ratio curves when varying k_c from 1 to 10 across the four datasets used in our experiments (data description in Section VI-A). When $k_c = 1$, the ratios of GPS points with their nearest segments as their actual segments are relatively low on all datasets, approximately 0.7 for three of the four datasets. This indicates the inaccuracy of solely relying on the nearest segment of a GPS point for map matching. More importantly, observe that when k_c increases to 10, the ratio values approach 1 across all datasets. This supports the intuition that the segment of a GPS point is typically nearby. Specifically, with probability close to 1, the true segment of a GPS point is within the top- k_c nearest segment (e.g., $k_c = 10$).

According to [41], civilian GPS devices have horizontal errors of $7.0m$ at a 95% confidence level and $30m$ at a 99% confidence level. For the datasets in Fig. 2, the average distance between a GPS point and its 10th nearest segment is $82.36m$ (PT), $122.03m$ (XA), $90.81m$ (BJ), and $95.85m$ (CD). These exceed the $30m$ threshold, confirming that the actual segment is within the top- $k_c = 10$ nearest segments with high probability. We also calculated the average distance to the true segment (and nearest segment): $5.56(4.67)m$, $4.96(4.23)m$, $6.84(5.21)m$, and $4.20(3.16)m$ for PT, XA, BJ, and CD, respectively. These distances are within the $7.0m$ range at the 95% confidence level, consistent with the observation that when $k_c = 1$, the ratio of GPS points with their nearest segments as actual segments is about 0.7 for 3 out of 4 datasets.

Therefore, we only need to investigate these segments as candidates to identify the segment of p_i , without tediously looking at all segments in G , where $k_c \ll |V|$. In Definition 8, we define the candidate segment set \mathcal{C}_{p_i} of p_i as the set of its top- k_c nearest segments in G .

Definition 8 (Candidate Segment Set of a GPS Point): Given a road network G and a GPS point p_i , the candidate segment set \mathcal{C}_{p_i} contains the top- k_c nearest segments ranked by the perpendicular distance from p_i to the segments in G .

The usage of the candidate set \mathcal{C}_{p_i} of p_i is two-fold. Besides acting as the candidate pool, the semantics of all segments in \mathcal{C}_{p_i} also server as the context of p_i to enhance its embedding, to be explained shortly. We can efficiently get \mathcal{C}_{p_i} of p_i via a top- k_c query of p_i over an R-tree [42] index of road segments.

After defining the candidate segment set \mathcal{C}_{p_i} for a GPS point p_i , we propose to formulate the task of finding p_i 's segment as a classification problem over the candidates in \mathcal{C}_{p_i} . Below, for a GPS point p_i , we state the class labels of the candidate segments in \mathcal{C}_{p_i} and the classification problem.

Class Labels in \mathcal{C}_{p_i} . For a GPS point p_i of a trajectory, a candidate segment $c \in \mathcal{C}_{p_i}$ has class label $y_c = 1$, if c is the ground-truth segment of p_i ; otherwise, the class label of c is $y_c = 0$. At most one segment in \mathcal{C}_{p_i} has class label 1. It is possible that all candidates are in class 0 if the ground truth is not included in \mathcal{C}_{p_i} , which is rare when $k_c = 10$.

Problem Statement. Given training trajectory data \mathcal{D} on road network G , the goal is to train a classification model to accurately identify the true segment e_i (class label 1) of point p_i from its candidate set \mathcal{C}_{p_i} . MMA computes the probability $P(c|p_i)$ for every candidate segment $c \in \mathcal{C}_{p_i}$ and the candidate with the highest probability is the segment of p_i .

B. MMA Method

In Algorithm 1, MMA first maps every GPS point p_i in T to a segment e_i (Lines 1-9) and then constructs route \mathcal{R} (Lines 10-13). As mentioned, we formulate the task of mapping p_i to a segment as a classification problem over \mathcal{C}_{p_i} . To achieve accurate classification, we design a point embedding module and a candidate segment embedding module, and adopt a cross entropy loss as objective, as illustrated in Fig. 3.

Candidate Segment Embedding. As shown in the bottom part of Fig. 3, we aim to generate an effective embedding \mathbf{c}_j for each candidate segment $c_j \in \mathcal{C}_{p_i}$ of point p_i in the sparse trajectory T . The embedding \mathbf{c}_j should intuitively reflect the connectivity of c_j in the road network G and its relationship with point p_i . To preserve the first aspect, we use network embedding Node2Vec [43] to pre-learn segment embeddings $\mathbf{W}_G \in \mathbb{R}^{n \times d_0}$ for all n segments in G . Then we initialize the learnable parameters \mathbf{W}^C of a fully-connected (FC) layer with these embeddings. This FC layer is subsequently trained to produce the representations \mathbf{e}_{c_j} of c_j in G . For every $c_j \in \mathcal{C}_{p_i}$, the FC layer transforms its one-hot id vector $\mathbf{1}_{c_j}$ to a dense representation \mathbf{e}_{c_j} via learnable weights \mathbf{W}^C ,

$$\mathbf{e}_{c_j} = \mathbf{1}_{c_j} \mathbf{W}^C \quad (1)$$

Algorithm 1: MMA for Map Matching

Input: A sparse trajectory $T = \langle p_1, \dots, p_\ell \rangle$, road network G
Output: Route \mathcal{R}

- 1 Get embedding $\mathbf{z}_i^{(2)}$ for every p_i in Eq.(3) preserving sequential patterns in T ;
- 2 **foreach** GPS point p_i in T **do**
- 3 Get candidate set \mathcal{C}_{p_i} with size k_c for p_i by Definition 8;
- 4 **foreach** candidate segment $c_j \in \mathcal{C}_{p_i}$ **do**
- 5 $\mathbf{c}_j \leftarrow$ Candidate Segment Embedding by Eq.(1)(2);
- 6 $\alpha_{j,i} \leftarrow$ Eq.(7) with $\mathbf{z}_i^{(2)}$ and \mathbf{c}_j ;
- 7 $\mathbf{p}_i \leftarrow \mathbf{z}_i^{(2)} + \sum_{\forall c_j \in \mathcal{C}_{p_i}} \alpha_{j,i} \cdot \mathbf{c}_j$;
- 8 $\forall c_j \in \mathcal{C}_{p_i}, P(c_j|p_i) \leftarrow \text{sigmoid}(\mathbf{c}_j \cdot \mathbf{p}_i)$;
- 9 p_i is mapped to segment $e_i \leftarrow \arg \max_{c_j \in \mathcal{C}_{p_i}} P(c_j|p_i)$;
- 10 $\mathcal{R} \leftarrow \langle \rangle$;
- 11 **for** $i = 1, \dots, \ell - 1$ **do**
- 12 Get the route \mathcal{R}_i between p_i 's segment e_i and p_{i+1} 's segment e_{i+1} ;
- 13 $\mathcal{R}.\text{append}(\mathcal{R}_i)$;
- 14 **return** \mathcal{R} ;

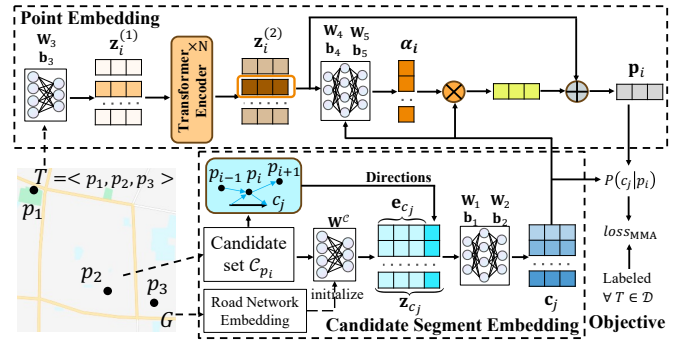


Fig. 3: Map GPS Points to Segments

where $\mathbf{W}^C \in \mathbb{R}^{n \times d_0}$ contains the learnable parameters and initially $\mathbf{W}^C = \mathbf{W}_G$, and n -dimensional one-hot vector $\mathbf{1}_{c_j} \in \{0, 1\}^n$ has all elements are 0, except 1 at the position corresponding to the id of c_j .

For the second aspect, we consider the directional relationship between the candidate $c_j \in \mathcal{C}_{p_i}$ and p_i in the sparse trajectory T . Specifically, as shown in Fig. 3, for every segment c_j , treated as a vector from its entrance to its exit, we compute four cosine similarity values: its similarity with the vector from the entrance of c_j to p_i , from p_i to the exit of c_j , from p_{i-1} to p_i in T , and p_i to p_{i+1} in T . We concatenate these four similarity values and with \mathbf{e}_{c_j} to get the embedding $\mathbf{z}_{c_j} \in \mathbb{R}^{(d_0+4)}$. This embedding \mathbf{z}_{c_j} is then fed into a Multi-layer Perceptron (MLP) in Eq.(2) to produce the final candidate embedding $\mathbf{c}_j \in \mathbb{R}^{d_2}$. Briefly, MLP is a fully connected feed-forward neural network, and it often combines with nonlinear activation function, e.g., $ReLU(x) = \max(0, x)$, to bring non-linearity into the model, to alleviate the vanishing gradient problem [44].

$$\mathbf{c}_j = ReLU(\mathbf{z}_{c_j} \mathbf{W}_1 + \mathbf{b}_1) \mathbf{W}_2 + \mathbf{b}_2 \quad (2)$$

where $\mathbf{W}_1 \in \mathbb{R}^{(d_0+4) \times d_1}$, $\mathbf{b}_1 \in \mathbb{R}^{d_1}$, $\mathbf{W}_2 \in \mathbb{R}^{d_1 \times d_2}$, and $\mathbf{b}_2 \in \mathbb{R}^{d_2}$.

As shown in Fig. 3, the candidate embedding \mathbf{c}_j serves two purposes: training the classification objective, and used as the context of p_i to enhance point embedding \mathbf{p}_i developed below.

Point Embedding. In the top part of Fig. 3, we generate the embedding \mathbf{p}_i of a GPS point p_i by considering the GPS information of p_i , the sequential features of T , and the embeddings of its candidate segments in \mathcal{C}_{p_i} . The idea is that each GPS point p_i has a unique set of candidate segments \mathcal{C}_{p_i} , which should be reflected in the embedding \mathbf{p}_i . Initially, point p_i has a vector $\mathbf{z}_i^{(0)}$ containing its min-max normalized latitude, longitude and timestamp. $\mathbf{z}_i^{(0)}$ is fed into an FC to get $\mathbf{z}_i^{(1)} = \mathbf{z}_i^{(0)}\mathbf{W}_3 + \mathbf{b}_3$, where $\mathbf{W}_3 \in \mathbb{R}^{3 \times d_2}$, $\mathbf{b}_3 \in \mathbb{R}^{d_2}$ are the learnable parameters.

Let $\mathbf{Z}_1 = [\mathbf{z}_1^{(1)}, \mathbf{z}_2^{(1)}, \dots, \mathbf{z}_\ell^{(1)}]$ be the embeddings of all GPS points in T . As depicted in Fig. 3, we adopt a transformer [45] to convert \mathbf{Z}_1 to $\mathbf{Z}_2 = [\mathbf{z}_1^{(2)}, \mathbf{z}_2^{(2)}, \dots, \mathbf{z}_\ell^{(2)}]$, capturing the sequential features in T with two stacked transformer layers,

$$\mathbf{Z}_2 = \text{Trans}(\mathbf{Z}_1) \quad (3)$$

We explain the details of a transformer [45], which will also be used in Section V. A transformer layer comprises two sub-layers: a multi-head self-attention mechanism and a position-wise feed-forward network (FFN). Self-attention allows the model to focus on relevant parts of the input sequence selectively. Multi-head attention (MHAttn) in Eq.(4) maps the input sequence representations to output representations using a scaled dot-product function across multiple heads:

$$\begin{aligned} \text{MHAttn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{concat}[\text{head}_1, \dots, \text{head}_h] \cdot \mathbf{W}^O \\ \text{head}_i &= \text{Attention}(\mathbf{Q}\mathbf{W}_i^Q, \mathbf{K}\mathbf{W}_i^K, \mathbf{V}\mathbf{W}_i^V) \\ \text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) &= \text{softmax}(\mathbf{Q}\mathbf{K}^T / \sqrt{d})\mathbf{V} \end{aligned} \quad (4)$$

where \mathbf{Q} , \mathbf{K} and \mathbf{V} represent the query, key, and value matrix respectively, $\mathbf{W}_i^Q, \mathbf{W}_i^K, \mathbf{W}_i^V \in \mathbb{R}^{d_{in} \times d_{in}/h}$ are projection matrices for the i -th head, h is the number of attention heads, and $\mathbf{W}^O \in \mathbb{R}^{d_{in} \times d_{out}}$ are parameters for the output, and d is the feature dimension of \mathbf{K} .

An FFN is two-layer MLP with ReLU activation in Eq.(5).

$$\text{FFN}(\mathbf{X}) = \text{ReLU}(\mathbf{X}\mathbf{W}_x + \mathbf{b}_x)\mathbf{W}_y + \mathbf{b}_y \quad (5)$$

where $\mathbf{W}_x, \mathbf{b}_x, \mathbf{W}_y, \mathbf{b}_y$ are learnable parameters.

Given input \mathbf{X} , a transformer layer applies MHAttn to \mathbf{X} itself, followed by an FFN, both with residual connection and layer normalization in Eq.(6).

$$\begin{aligned} \text{Trans}(\mathbf{X}) &= \text{LayerNorm}(\mathbf{X}' + \text{FFN}(\mathbf{X}')) \\ \mathbf{X}' &= \text{LayerNorm}(\mathbf{X} + \text{MHAttn}(\mathbf{X}, \mathbf{X}, \mathbf{X})) \end{aligned} \quad (6)$$

where MHAttn is defined in Eq.(4) and FFN is Eq.(5).

As explained, the candidate segments \mathcal{C}_{p_i} of p_i represent the specific context of p_i , and this should be reflected in the embedding of p_i . To achieve this, we design an attention mechanism to learn the attention score $\alpha_{j,i}$ of every candidate embedding \mathbf{c}_j relative to $\mathbf{z}_i^{(2)}$ of p_i . This is done in Eq.(7) via a two-layer MLP over the concatenation of $\mathbf{z}_i^{(2)}$ and

\mathbf{c}_j to get intermediate $s_{j,i}$, which is normalized by softmax normalization to get $\alpha_{j,i}$.

$$\begin{aligned} s_{j,i} &= \text{ReLU}(\text{concat}[\mathbf{z}_i^{(2)}, \mathbf{c}_j]\mathbf{W}_4 + \mathbf{b}_4)\mathbf{W}_5 + \mathbf{b}_5 \\ \alpha_{j,i} &= \frac{\exp(s_{j,i})}{\sum_{k=1}^{k_c} \exp(s_{k,i})} \end{aligned} \quad (7)$$

where $\mathbf{W}_4 \in \mathbb{R}^{2d_2 \times d_3}$, $\mathbf{b}_4 \in \mathbb{R}^{d_3}$, $\mathbf{W}_5 \in \mathbb{R}^{d_3 \times 1}$, $\mathbf{b}_5 \in \mathbb{R}^1$ are the learnable parameters.

The attention scores are then used to aggregate all \mathbf{c}_j of \mathcal{C}_{p_i} and $\mathbf{z}_i^{(2)}$, to get the final embedding \mathbf{p}_i of point p_i ,

$$\mathbf{p}_i = \mathbf{z}_i^{(2)} + \sum_{\forall \mathbf{c}_j \in \mathcal{C}_{p_i}} \alpha_{j,i} \cdot \mathbf{c}_j \quad (8)$$

Given the embedding \mathbf{p}_i of point p_i and the embedding \mathbf{c}_j of a candidate segment $c_j \in \mathcal{C}_{p_i}$, the probability $P(c_j|p_i)$ that c_j is the segment of p_i is computed as the inner product of \mathbf{c}_j and \mathbf{p}_i normalized by $\text{sigmoid}(x) = \frac{1}{1+\exp(-x)}$ into $(0, 1)$.

$$P(c_j|p_i) = \text{sigmoid}(\mathbf{c}_j \cdot \mathbf{p}_i) \quad (9)$$

Objective. During offline training stage, a candidate $c_j \in \mathcal{C}_{p_i}$ has class label $y_{c_j} = 1$ or 0. MMA employs a binary cross entropy loss for a point p_i ,

$$L_{p_i} = - \sum_{\forall \mathbf{c}_j \in \mathcal{C}_{p_i}} (y_{c_j} \log P(c_j|p_i) + (1 - y_{c_j}) \log (1 - P(c_j|p_i)))$$

On training data \mathcal{D} , the overall training loss of MMA is,

$$L_{\text{MMA}} = \frac{1}{|\mathcal{D}|} \sum_{\forall T \in \mathcal{D}} \sum_{\forall p_i \in T} L_{p_i} \quad (10)$$

Algorithm. Algorithm 1 shows the pseudocode MMA for forward execution to recover the route \mathcal{R} of a sparse trajectory T on road network G . From Lines 1-9, MMA maps each GPS point p_i to its corresponding segment. At Line 1, we first incorporate the sequential pattern in T into embedding $\mathbf{z}_i^{(2)}$ for every GPS point p_i , to be used later. Then, for every p_i , we obtain its candidate segment set \mathcal{C}_{p_i} and generate candidate segment embedding \mathbf{c}_j for each $c_j \in \mathcal{C}_{p_i}$ (Lines 2-5). The attention between p_i and each candidate segment c_j is obtained via Eq.(7) with $\mathbf{z}_i^{(2)}$ and \mathbf{c}_j as input (Lines 6). Using the attention scores, we derive the final point embedding \mathbf{p}_i via Eq.(8) (Line 7). Then we compute probability $P(c_j|p_i)$ using embeddings \mathbf{c}_j and \mathbf{p}_i (Line 8), and map p_i to the segment with the highest probability in \mathcal{C}_{p_i} (Line 9). Since the matched segments of two consecutive points p_i and p_{i+1} may not be connected, we construct the routes \mathcal{R}_i between the segments of all possible p_i and p_{i+1} when necessary (Line 10-12) and combine them to get the final \mathcal{R} (Line 13). Note that MMA is orthogonal to the route planning method used at Line 12. In our experiments, for our methods and baselines requiring route planning as a routine, we use the same DA-based method from [2] that relies on basic statistical counts.

Complexity. Let d be the dimension of embeddings involved. T has ℓ GPS points. Line 1 costs $O(\ell^2 d)$ time to run the transformer. To get \mathcal{C}_{p_i} of p_i , a top- k_c query over R-tree takes $O(\log n)$ time in average case. MLPs at Lines 5-6 costs $O(d^2)$ time, and k_c is constant. Lines 2-8 are executed ℓ times. The

route planning routine at Line 12 has $O(l'\tilde{deg})$ time [2], where l' is the max route length limit and \tilde{deg} is the max degree in G , which is typically small. Regarding l' same as l , MMA costs $O(\ell^2 d + \ell d^2 + \ell \log n + \ell^2 \tilde{deg})$ time, where $\log n$ is sub-linear to n and l , d , and \tilde{deg} are usually small, indicating the efficiency of MMA.

V. TRMMA: SPARSE TRAJECTORY RECOVERY

After obtaining the route \mathcal{R} of a sparse trajectory T , TRMMA focuses on the segments in \mathcal{R} to infer missing map-matched points a_i , thereby recovering the map-matched ϵ -sampling trajectory \mathcal{T}_ϵ for T . The segments in \mathcal{R} are typically much fewer than those in G , making TRMMA efficient, especially compared to existing methods [10], [14] that evaluate all segments in G for trajectory recovery.

As shown in Algorithm 2, TRMMA first invokes MMA to get \mathcal{R} (Line 1). For each GPS point p_i in T , we get its map-matched point a_i by projecting p_i to its segment e_i (Lines 2-4). Then TRMMA utilizes the sparse trajectory T , a_i of every p_i in T , and route \mathcal{R} together to get \mathcal{T}_ϵ satisfying the target ϵ -sampling rate from Lines 5 to 16, using a proposed DualFormer encoding module and a multitask decoding module illustrated in Fig. 4. Specifically, we design the DualFormer encoding module to integrate T and \mathcal{R} sequences to produce expressive embeddings. Subsequently, we develop the decoding module that utilizes the embeddings to estimate missing points a_j , including $a_j.e$ and $a_j.r$, to get \mathcal{T}_ϵ . Briefly, we regard the segments in \mathcal{R} as the candidates to predict $a_j.e$, while treating the estimation of $a_j.r$ as a regression problem.

DualFormer Encoding. The DualFormer encoding module employs one transformer to encode the sparse trajectory T with sequence length ℓ into representations $\mathbf{T} \in \mathbb{R}^{\ell \times d_h}$, and uses another transformer to encode the route \mathcal{R} of sequence length $\ell_{\mathcal{R}}$ into representations $\mathbf{R} \in \mathbb{R}^{\ell_{\mathcal{R}} \times d_h}$. An attention mechanism is then applied to \mathbf{T} and \mathbf{R} to capture the intrinsic relationship between the points in T and the segments in \mathcal{R} . This process yields the final embeddings $\mathbf{H} \in \mathbb{R}^{\ell_{\mathcal{R}} \times d_h}$, which serve as the input for the decoding process.

In particular, for all points p_i in T , we first initialize an embedding matrix $\mathbf{T}_0 \in \mathbb{R}^{\ell \times d_h}$, the i -th row of which is the initial encoding of p_i containing information of p_i and its map-matched point a_i , including min-max normalized $p_i.lat, p_i.lng, p_i.t, a_i.r$ and the id embedding of segment $a_i.e$. \mathbf{T}_0 goes through one FC layer to get \mathbf{T}_1 , and then a transformer is applied over \mathbf{T}_1 to capture the sequential patterns of T into $\mathbf{T} \in \mathbb{R}^{\ell \times d_h}$,

$$\mathbf{T} = \text{Trans}_T(\mathbf{T}_1), \quad \mathbf{T}_1 = \mathbf{T}_0 \mathbf{W}_6 + \mathbf{b}_6, \quad (11)$$

where $\mathbf{W}_6 \in \mathbb{R}^{d_4 \times d_h}$, $\mathbf{b}_6 \in \mathbb{R}^{d_h}$ are the learnable parameters.

Meanwhile, we encode the sequence of segments in route \mathcal{R} via another transformer. Specifically, for all $\ell_{\mathcal{R}}$ segments e_j in \mathcal{R} , we apply an FC layer to get $\mathbf{R}_1 \in \mathbb{R}^{\ell_{\mathcal{R}} \times d_h}$, where the j -th row is an embedding of the j -th segment e_j in \mathcal{R} . We then adopt a transformer over \mathbf{R}_1 to get $\mathbf{R} \in \mathbb{R}^{\ell_{\mathcal{R}} \times d_h}$, preserving the sequential patterns in \mathcal{R} ,

$$\mathbf{R} = \text{Trans}_{\mathcal{R}}(\mathbf{R}_1), \quad \mathbf{R}_1 = \mathbf{1}_{\mathcal{R}} \mathbf{W}_7 + \mathbf{b}_7, \quad (12)$$

Algorithm 2: TRMMA for Trajectory Recovery on Road Network

Input: A sparse trajectory $T = \langle p_1, p_2, \dots, p_\ell \rangle$, a target sampling rate ϵ , road network G
Output: Recovered map-matched ϵ -sampling trajectory $\mathcal{T}_\epsilon = \langle a_1, a_2, \dots, a_{\ell_\epsilon} \rangle$

- 1 Invoke MMA (Algorithm 1) to get the route \mathcal{R} of T ;
- 2 **foreach** GPS point p_i in T with segment e_i in \mathcal{R} **do**
- 3 $a_i.e \leftarrow e_i$; $a_i.t \leftarrow p_i.t$;
- 4 Get $a_i.r$ by orthogonal projection of p_i to e_i ;
- 5 $\mathbf{H} \leftarrow$ Invoke DualFormer encoding module with input T , all a_i of p_i in T , route \mathcal{R} (Eq.(11)-(14));
- 6 $\mathbf{h}_0 \leftarrow$ mean pooling over \mathbf{H} ;
- 7 $\mathcal{T}_\epsilon \leftarrow \langle a_1 \rangle$;
- 8 **foreach** a_i of p_i in T **do**
- 9 $n_i \leftarrow \lfloor (a_{i+1}.t - a_i.t) / \epsilon \rfloor$;
- 10 **for** $j = 1, \dots, n_i$ **do**
- 11 Invoke decoding process with \mathbf{h}_j and \mathbf{H} as input of Eq.(15) to get $w_{k,j}$ for segment $e_k \in \mathcal{R}[a_{j-1}.e, :]$;
- 12 $\forall e_k \in \mathcal{R}[a_{j-1}.e, :]$, get $P(e_k|a_j)$ by Eq.(16);
- 13 $a_j.e \leftarrow \arg \max_{e_k \in \mathcal{R}[a_{j-1}.e, :]} P(e_k|a_j)$ by Eq.(17);
- 14 $a_j.r \leftarrow$ Eq.(18);
- 15 $\mathcal{T}_\epsilon.append(a_j)$;
- 16 $\mathcal{T}_\epsilon.append(a_{i+1})$;
- 17 **return** \mathcal{T}_ϵ ;

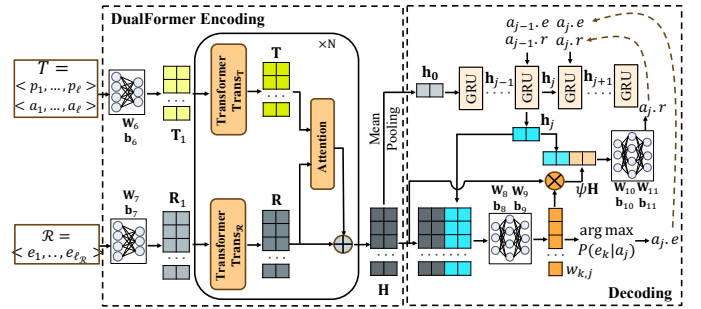


Fig. 4: TRMMA to Recover Missing Points on Route

where $\mathbf{1}_{\mathcal{R}} \in \mathbb{R}^{\ell_{\mathcal{R}} \times n}$ are the one-hot vectors of the segments in \mathcal{R} , $\mathbf{W}_7 \in \mathbb{R}^{n \times d_h}$, $\mathbf{b}_7 \in \mathbb{R}^{d_h}$ are the learnable parameters.

Now the question is how to fuse \mathbf{T} and \mathbf{R} together to obtain the final representation $\mathbf{H} \in \mathbb{R}^{\ell_{\mathcal{R}} \times d_h}$ that preserves not only the sequences T and \mathcal{R} but also their relationships. \mathbf{H} should have $\ell_{\mathcal{R}}$ rows, each corresponding to a segment in \mathcal{R} , since the decoding module only focuses on the $\ell_{\mathcal{R}}$ segments in \mathcal{R} as candidates to recover missing points. To this end, we design an attention mechanism to produce \mathbf{H} . Intuitively, from the view of a segment e_j in route \mathcal{R} , the points in the sparse trajectory T should have distinct influence on it. Thus, for every segment e_j in \mathcal{R} with embedding $\mathbf{R}[j, :]$, we compute its attention score $\beta_{i,j}$ for every point p_i with embedding $\mathbf{T}[i, :]$ in Eq.(13).

$$\beta_{i,j} = \frac{\exp(b_{i,j})}{\sum_{k=1}^{\ell} \exp(b_{k,j})}, \quad b_{i,j} = \mathbf{R}[j, :] \cdot \mathbf{T}[i, :] \quad (13)$$

The representation of the trajectory T w.r.t. segment e_j in \mathcal{R} is $\sum_{p_i \in T} \beta_{i,j} \mathbf{T}[i, :]$. Then the final embedding of $\mathbf{H}[j, :]$ is

$$\mathbf{H}[j, :] = \mathbf{R}[j, :] + \sum_{p_i \in T} \beta_{i,j} \mathbf{T}[i, :], \quad (14)$$

where $\mathbf{H} \in \mathbb{R}^{|\mathcal{R}| \times d_h}$ contains the output embeddings.

Decoding. With the embeddings above as input, the decoding process predicts a missing a_j by two concurrent tasks: classifying the segment $a_j.e$ from the candidate segments in \mathcal{R} , and regressing the position ratio $a_j.r \in [0, 1]$.

In the classification task of the decoding module, we determine the segment $a_j.e$ for the missing point a_j . A straightforward approach would be to map the hidden state \mathbf{h}_j to a probability vector matching the number of road segments in the road network G , and then apply a mask that sets the probabilities of segments in \mathcal{R} to 1 and all others to 0, selecting the segment with the highest probability as $a_j.e$. However, this method is inefficient for large road networks [10], [14]. Moreover, since sparse trajectories produce routes of varying lengths, we opt for binary classification on the segments within \mathcal{R} to predict $a_j.e$, in contrast to existing methods that use multi-classification over a fixed set of $|E|$ segments in G .

As shown in Fig. 4, we use GRU [46] as the sequential model to produce the hidden state \mathbf{h}_j , which will be used to guide the recovery of the missing a_j . The initial hidden state \mathbf{h}_0 is derived by mean pooling the row vectors of \mathbf{H} . When predicting a_j , for the k -th segment in \mathcal{R} with k -th row $\mathbf{H}[k, :]$ as its embedding, we concatenate the embedding with the hidden state \mathbf{h}_j . This concatenated vector is then processed by a two-layer MLP classifier to produce a scalar value $w_{k,j}$.

$$w_{k,j} = \text{ReLU}(\text{concat}[\mathbf{H}[k, :], \mathbf{h}_j] \mathbf{W}_8 + \mathbf{b}_8) \mathbf{W}_9 + \mathbf{b}_9 \quad (15)$$

where $\mathbf{W}_8 \in \mathbb{R}^{2d_h \times d_h}$, $\mathbf{b}_8 \in \mathbb{R}^{d_h}$, $\mathbf{W}_9 \in \mathbb{R}^{d_h \times 1}$, $\mathbf{b}_9 \in \mathbb{R}^1$ are the learnable parameters.

The scalar value $w_{k,j}$ serves two purposes. First, we can get the probability that segment e_k is the segment of a_j with hidden state \mathbf{h}_j via a sigmoid function over $w_{k,j}$,

$$P(e_k|a_j) = \text{sigmoid}(w_{k,j}) \quad (16)$$

As shown in Algorithm 2 (Lines 10-15), we recover missing points in \mathcal{T}_ϵ sequentially. The segments of the recovered points in \mathcal{T}_ϵ must adhere the segment order specified in route \mathcal{R} . In other words, if the missing point a_j is recovered after $a_{j'}$, then the segment $a_j.e$ should not precede $a_{j'}.e$ in route \mathcal{R} . Hence, when predicting segment $a_j.e$, we select the segment with the highest probability $P(e_k|a_j)$ among all segments in the sub-Route of \mathcal{R} starting from $a_{j-1}.e$, denoted as $\mathcal{R}[a_{j-1}.e, :]$, where a_{j-1} is the point just recovered before a_j .

$$a_j.e = \arg \max_{e_k \in \mathcal{R}[a_{j-1}.e, :]} P(e_k|a_j) \quad (17)$$

The second usage of $w_{k,j}$ is for the regression task to predict the position ratio $a_j.r$. We apply softmax over all $w_{k,j}$ values for $k = 1, \dots, \ell_{\mathcal{R}}$, where $\ell_{\mathcal{R}}$ is the sequence length of route \mathcal{R} , to get $\psi_{k,j}$, which is an attention value between $\mathbf{H}[k, :]$ and \mathbf{h}_j . Let attention vector be $\psi_j = [\psi_{1,j}, \dots, \psi_{\ell_{\mathcal{R}},j}]$. We then apply a weighted row sum in \mathbf{H} with each row weighted by $\psi_{k,j}$, to get $\psi_j \mathbf{H}$. Finally, we concatenate \mathbf{h}_j and $\psi_j \mathbf{H}$ to go through an MLP with sigmoid to get the predicted ratio $a_j.r$,

$$a_j.r = \text{sigmoid}(\text{ReLU}(\text{concat}[\mathbf{h}_j, \psi_j \mathbf{H}] \mathbf{W}_{10} + \mathbf{b}_{10}) \mathbf{W}_{11} + \mathbf{b}_{11}), \quad (18)$$

where $\mathbf{W}_{10} \in \mathbb{R}^{2d_h \times d_h}$, $\mathbf{b}_{10} \in \mathbb{R}^{d_h}$, $\mathbf{W}_{11} \in \mathbb{R}^{d_h \times 1}$, $\mathbf{b}_{11} \in \mathbb{R}^1$ are the learnable parameters.

Lastly, in the decoding module shown in Fig. 4, the current hidden state \mathbf{h}_j , the predicted segment and ratio of a_j are then used as the input to obtain the next hidden state \mathbf{h}_{j+1} for recovering the next missing point a_{j+1} , until \mathcal{T}_ϵ is obtained for the input sparse trajectory T .

Objective. During offline training, a_j has ground-truth segment $a_j.e$, and a segment e_k in route \mathcal{R} has class label $y_{k,j} = 1$ if $e_k = a_j.e$, class label $y_{k,j} = 0$ otherwise. The training loss of predicting $a_j.e$ among all segments in \mathcal{R} is

$$L_{seg} = - \sum_{\forall a_j \in \mathcal{T}_\epsilon} \sum_{\forall e_k \in \mathcal{R}} (y_{k,j} \log P(e_k|a_j) + (1 - y_{k,j}) \log (1 - P(e_k|a_j))) \quad (19)$$

The prediction of position ratio adopts mean absolute error,

$$L_r = \sum_{\forall a_j \in \mathcal{T}_\epsilon} |a_j.r - a_j.\hat{r}| \quad (20)$$

where $a_j.r$ is the predicted ratio and $a_j.\hat{r}$ is the ground truth.

The overall loss of TRMMA is the weighted combination of the two loss functions in Eq.(19) and (20),

$$L_{\text{TRMMA}} = \frac{1}{|\mathcal{D}|} \sum_{\forall \mathcal{T}_\epsilon \in \mathcal{D}} (L_{seg} + \lambda L_r) \quad (21)$$

where λ is a hyperparameter to trade off the two tasks.

Algorithm. In Algorithm 2, after obtaining a_i for every p_i in T over \mathcal{R} (Lines 1-4), we compute the DualFormer encoding \mathbf{H} by Eq.(11)-(14) (Line 5) and get the initial state \mathbf{h}_0 by mean pooling (Line 6). After initializing \mathcal{T}_ϵ at Line 7, for every pair of consecutive points a_i and a_{i+1} , we calculate the number n_i of missing points necessary to maintain an ϵ -sampling rate based on their time interval and ϵ (Lines 8-9). From Lines 10 to 15, TRMMA invokes the decoding process to recover n_i missing points p_j s by obtaining probability $P(e_k|a_j)$ for candidate segments e_k (Lines 11-12), selecting the candidate with the highest probability as $a_j.e$ (Line 13), and predicting position ratio $a_j.r$ (Line 14). The recovered a_j is appended to \mathcal{T}_ϵ at Line 15, and a_{i+1} of GPS point p_j is appended at Line 16. The final \mathcal{T}_ϵ is returned at Line 17.

Complexity. The first step of TRMMA is MMA with $O(\ell^2 d + \ell d^2 + \ell \log n + \ell^2 \tilde{deg})$ time. For the encoding and decoding processes in TRMMA, the time complexity is as follows. Let d be the dimension of embeddings involved. In the encoding process, we run a transformer over T in $O(\ell^2 d)$ time, and another transformer over \mathcal{R} in $O(\ell_{\mathcal{R}}^2 d)$ time, where $\ell_{\mathcal{R}}$ is the sequence length of \mathcal{R} . We compute attention between T and \mathcal{R} in $O(\ell \ell_{\mathcal{R}} d)$, and the time of MLPs is $O(d^2)$. In the decoding process of TRMMA, we need to recover at most ℓ_ϵ points of \mathcal{T}_ϵ , which costs $O(\ell_\epsilon (\ell_{\mathcal{R}} + d^2))$. The time of encoding and decoding in TRMMA is $O(\ell^2 d + \ell_{\mathcal{R}}^2 d + \ell \ell_{\mathcal{R}} d + \ell_\epsilon (\ell_{\mathcal{R}} + d^2))$, which simplifies to $O(\ell^2 d + \ell d^2)$ if using ℓ to represent $\ell, \ell_{\mathcal{R}}, \ell_\epsilon$. Thus, the overall time complexity of TRMMA is $O(\ell^2 d + \ell d^2 + \ell \log n + \ell^2 \tilde{deg})$.

VI. EXPERIMENTS

After the experimental setup in Section VI-A, we evaluate the efficiency and effectiveness of trajectory recovery in Section VI-B followed by experimental analysis in Section VI-C. Then we conduct efficiency and effectiveness evaluations and experimental analysis of map matching in Section VI-D.

A. Experimental Setup

Datasets. Table II lists the statistics of the 4 real-world trajectory data, including taxi data on Porto District in Portugal (PT) [47] and Beijing (BJ) [48], and DiDi ride-sharing data on Xi'an (XA) and Chengdu (CD) that have been widely used in the literature [10], [19], [38]. All datasets can be downloaded from our code repository. All datasets are large with millions of trajectories. Table II shows the ϵ -sampling rate in seconds, number of GPS points, length, and travel time of trajectories. The time period, the number of segments and intersections and the area of road networks are also provided. Note that the areas of all datasets already cover the major urban regions of cities, comparable to or larger than existing studies [10]. Road networks are from OpenStreetMap [49].

Note that the ϵ sampling rate of each dataset in Table II is regarded as the target high-sampling rate. For an ϵ -sampling trajectory, we generate its sparse trajectory by randomly sampling the points in it, so that the resulting sparse trajectory T has average interval ϵ/γ , where γ is a ratio in $(0, 1)$ to control the level of sparsity. By default, γ is 0.1, meaning that the time interval of sparse trajectories is 10 times larger than the high-sampling trajectories, *i.e.*, 150s on PT, 120s on XA, 600s on BJ, and 120s on CD. We also vary γ from 0.1 to 0.5 to study the effect under different sparsity levels. To get ground truth, for a raw high-sampling GPS trajectory with ϵ sample rate, we use map matching from [28] to align it with the road network, to get ground-truth route. For each GPS point p_i in the raw trajectory, we draw a perpendicular line from p_i to its ground-truth segment. The intersection point is the map-matched point a_i for p_i , thus obtaining the ground-truth map-matched ϵ -sampling trajectory $\hat{\mathcal{T}}_\epsilon$. For a dataset, we randomly split its trajectories into training, validation, testing sets with ratio in 40%, 30%, and 30%.

Competitors. For *trajectory recovery*, we compare TRMMA with baselines in 3 categories. (i) Trajectory recovery on road networks: MTrajRec [14], RNTrajRec [10] and MM-STGED [19]; Linear that first maps a sparse trajectory to road network by FMM [28] and then applies linear interpolation. (ii) Free-space trajectory recovery methods DHTR [20] and TERI [21] that are extended to road networks by substituting their grid structures with road segments. (iii) Trajectory representation learning methods followed by a decoder to recover trajectories, including TrajGAT+Dec [37], TrajCL+Dec [38], and ST2Vec+Dec [39]. We follow [10], [19] to compare this category of methods, and use the decoder in [14]. For *map matching*, we compare MMA with the latest methods LHMM [11] and GraphMM [13], FMM [28], DeepMM [32], RNTrajRec modified to only return routes, and Nearest that maps a GPS point to its nearest segment.

TABLE II: Dataset Statistics

	Porto (PT)	Xi'an (XA)	Beijing (BJ)	Chengdu (CD)
# of trajectories	1,013,437	1,426,950	1,176,097	2,382,422
ϵ sampling rate (s)	15	12	60	12
Avg # of points	40.21	69.36	31.59	54.32
Avg length (m)	4,180.41	5,049.27	6494.78	4,397.41
Avg travel time (s)	585.12	816.44	845.95	636.37
Time period	2013/7/1- 2014/6/30	2016/10/1- 2016/10/31	2009/3/2- 2009/3/25	2016/10/1- 2016/10/31
Area (km ²)	11.7 × 5.2	9.1 × 8.5	29.6 × 30.0	10.4 × 10.8
# of segments	11,491	5,699	65,276	9,255
# of intersections	5,330	2,579	28,738	3,973

Implementations. In our methods, k_c is set to 10 as analyzed in Section IV-A. We set dimension $d_0 = 64$ in Eq.(1), $d_2 = 64$ in Eq.(2) for candidate segment embedding and point embedding in MMA. We stack two layers and four heads of transformer in Eq.(3). As for the hidden dimension of MLP in Eq.(2) and (7), we set $d_1 = 128$ and $d_3 = 256$ respectively. We set the input dimension of transformer $d_h = 64$ in Eq.(11) and (12), and the FFN dimension in transformer is 512. We stack four DualFormer layers with four heads in TRMMA. The learning rate and batch size are 1e-3 and 512 respectively. Our methods are implemented in Python 3.8 with PyTorch 1.13. Source codes of all competitors are in Python and obtained from the respective authors. We follow their suggested settings to tune optimal parameters. All methods are trained to converge. For any method requiring route planning as a subroutine, we adopt the DA-based route in [2] for fair comparison. Existing studies on route planning [1], [2] indicate that it is possible but rare for the output route from a source e_i to fail in reaching destination e_{i+1} , with a very low probability, such as 0.06% on PT in experiments. This has a negligible impact on the overall map-matching quality. In such a case, a practical solution is to use the fastest route to connect e_i and e_{i+1} . All experiments are conducted on a Linux machine powered by Intel Xeon® Gold 6226R 2.90GHz CPU and NVIDIA GTX 3090 GPU with 24GB video memory.

Evaluation Metrics. (i) *Efficiency metrics.* We evaluate a method by inference time to recover trajectories and training time per epoch. (ii) *Effectiveness metrics.* For trajectory recovery of $T = \langle p_1, p_2, \dots, p_\ell \rangle$ with map-matched ϵ -sampling trajectory $\mathcal{T}_\epsilon = \langle a_1, a_2, \dots, a_{\ell_\epsilon} \rangle$, we follow [10], [14] to use Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) to evaluate the distance error of the recovered map-matched points over ground truth,

$$MAE(\mathcal{T}_\epsilon, \hat{\mathcal{T}}_\epsilon) = \frac{1}{\ell_\epsilon} \sum_{i=1}^{\ell_\epsilon} |d(a_i, \hat{a}_i)|; RMSE(\mathcal{T}_\epsilon, \hat{\mathcal{T}}_\epsilon) = \sqrt{\frac{1}{\ell_\epsilon} \sum_{i=1}^{\ell_\epsilon} (d(a_i, \hat{a}_i))^2}, \quad (22)$$

where $d(a_i, \hat{a}_i)$ is the road network distance between predicted a_i and ground truth \hat{a}_i .

To evaluate if a method recovers the segments of a trajectory, we adopt 4 popular metrics [10], [14]: Precision, Recall, F1 score and Accuracy. For *trajectory recovery*, let S represent the segments of all a_i in \mathcal{T}_ϵ and \hat{S} represent the segments in the ground-truth $\hat{\mathcal{T}}_\epsilon$, and the 4 metrics are calculated as follows. When evaluating *map matching*, S becomes the returned

TABLE III: Effectiveness of Trajectory Recovery. *Higher* recall, precision, F1, and accuracy (in percentage), and *lower* MAE and RMSE (in meters) represent better performance. Best is in bold and runner-up is underlined.

Methods	PT						XA					
	Recall	Precision	F1	Accuracy	MAE	RMSE	Recall	Precision	F1	Accuracy	MAE	RMSE
Linear	66.42	65.85	65.83	39.54	127.6	170.1	85.65	86.58	85.73	66.26	94.2	127.1
DHTR	<u>69.84</u>	73.96	71.52	47.92	135.4	181.7	<u>85.91</u>	91.92	88.47	69.39	162.2	211.2
TERI	67.76	72.11	69.35	43.23	180.5	249.6	83.32	90.59	86.15	60.73	222.5	301.2
TrajGAT+Dec	56.44	74.21	63.45	39.83	188.6	251.8	75.06	88.78	80.25	60.37	203.3	265.1
TrajCL+Dec	60.11	77.61	67.18	43.67	152.2	204.8	75.76	89.01	80.99	62.56	154.9	204.4
ST2Vec+Dec	61.49	76.99	67.80	43.59	149.1	200.1	76.38	87.58	80.69	62.35	158.1	207.7
MTrajRec	66.24	77.33	70.93	49.72	112.1	151.5	82.58	92.18	86.65	71.19	105.9	140.3
MM-STGED	67.52	78.54	72.19	50.19	112.9	153.8	84.01	93.26	87.94	73.69	98.4	132.8
RNTrajRec	67.29	<u>79.52</u>	<u>72.48</u>	<u>52.22</u>	<u>102.6</u>	<u>140.6</u>	84.73	<u>93.76</u>	<u>88.61</u>	<u>74.79</u>	<u>93.1</u>	<u>126.5</u>
TRMMA	72.07	80.92	75.87	57.83	84.10	121.8	86.89	95.09	90.44	78.95	68.1	103.1

Methods	BJ						CD					
	Recall	Precision	F1	Accuracy	MAE	RMSE	Recall	Precision	F1	Accuracy	MAE	RMSE
Linear	50.28	54.13	51.54	37.35	325.5	431.3	82.66	81.82	81.77	58.17	106.2	141.5
DHTR	54.41	59.61	56.16	43.77	486.7	629.4	83.14	87.22	84.68	63.84	168.3	222.3
TERI	<u>56.61</u>	59.34	57.23	44.34	451.5	592.1	81.62	86.07	83.15	57.99	216.6	294.7
TrajGAT+Dec	47.95	58.64	51.29	39.41	476.5	605.4	74.42	87.56	80.05	57.95	200.4	264.2
TrajCL+Dec	52.63	64.39	57.02	43.04	397.1	509.2	75.12	87.79	80.11	60.14	152.6	204.3
ST2Vec+Dec	51.36	<u>62.98</u>	55.67	41.89	423.5	543.3	75.46	88.18	80.49	60.43	155.1	206.9
MTrajRec	53.35	62.44	56.68	43.58	375.1	477.2	83.34	91.24	86.65	68.42	104.8	141.1
MM-STGED	55.49	62.98	58.19	45.21	415.4	551.3	83.81	92.01	87.25	69.78	103.1	140.5
RNTrajRec	55.39	64.38	58.78	46.22	318.2	413.7	84.17	93.26	88.05	71.78	95.1	131.8
TRMMA	62.15	66.53	63.62	53.71	234.3	327.1	85.86	93.95	89.29	75.28	75.1	114.7

route \mathcal{R} while \hat{S} is ground-truth route; following existing methods [11], [18], we use Recall, Precision, and F1, as well as Jaccard similarity $Jaccard(S, \hat{S}) = \frac{|S \cap \hat{S}|}{|S \cup \hat{S}|}$.

$$Recall(S, \hat{S}) = \frac{|S \cap \hat{S}|}{|\hat{S}|} \quad F1(S, \hat{S}) = \frac{2pre(S, \hat{S})rec(S, \hat{S})}{pre(S, \hat{S}) + rec(S, \hat{S})}$$

$$Precision(S, \hat{S}) = \frac{|S \cap \hat{S}|}{|\hat{S}|} \quad Accuracy(S, \hat{S}) = \frac{1}{\ell_\epsilon} \sum_{i=1}^{\ell_\epsilon} \mathbf{1}\{s_i = \hat{s}_i\}$$

For each metric, we calculate the metric score per trajectory and report the average over all testing trajectories.

B. Effectiveness and Efficiency of Trajectory Recovery

Effectiveness. Table III reports the recall, precision, F1, accuracy, MAE, and RMSE of all methods for trajectory recovery. An overall observation is that our method TRMMA consistently achieves the best performance on all datasets under all evaluation metrics, outperforming existing methods often by a significant margin. For instance, On PT, TRMMA achieves 75.87% F1 score, while the F1 of RNTrajRec is 72.48%, indicating 3.39% improvement. Further, on PT, RNTrajRec achieves strong precision at the cost of low recall, while TRMMA can improve both metrics. on BJ, TRMMA achieves 62.15% recall score, indicating 5.54% improvement over the best competitor TERI with 56.61% recall. On XA, TRMMA has MAE 68.1m and RMSE 103.1m, achieving 26.85% and 18.5% relative improvements over RNTrajRec with 93.1m and 126.5m respectively. A similar observation can be made on BJ, where TRMMA has MAE 234.3m, significantly lower than the MAE of RNTrajRec 318.2m. Table III demonstrates that TRMMA is effective in accurately predicting the correct segments and position ratios to infer missing points to get \mathcal{T}_ϵ . This is achieved by the proposed techniques in TRMMA, including the encoding and decoding modules in Section V, and the map matching method MMA in Section IV, which

will be evaluated independently in Section VI-D. Additionally, Linear and other baselines in the category of trajectory representation learning, including TrajGAT+Dec, TrajCL+Dec, and ST2Vec+Dec, achieve inferior performance, indicating the importance of tailoring techniques for trajectory recovery. Due to space limit, in what follows, we omit them.

Efficiency. Fig. 5 reports the average *inference time* per 1000 sparse trajectory recoveries in seconds. TRMMA is faster than the competitors, and more importantly, compared with the best competitor RNTrajRec that achieves high effectiveness in Table III, TRMMA is significantly faster, often by orders of magnitude. For example, on BJ, TRMMA only needs 0.59s to recover 1000 trajectories, 75.9× faster than RNTrajRec that needs 44.79s. Although DHTR is relatively quick, it delivers lower quality in Table III. Fig. 6 reports the *training time* per epoch. TRMMA is much faster to train than the competitors. For instance, on XA dataset, our method TRMMA only needs 7.95 minutes per epoch to train, 19× faster than RNTrajRec that costs 152.95 minutes per epoch. DHTR and TERI are relatively fast to train, but produce inferior quality as reported in Table III. The efficiency of TRMMA stems from the technical designs in Section V with powerful encoding and decoding processes. TRMMA first identifies a route \mathcal{R} of T and then focuses solely on the segments in \mathcal{R} for recovery. In contrast, existing methods [10], [14] consider all segments from a road network as candidates. The efficiency of TRMMA aligns with the time complexity analyzed in Section V.

C. Experimental Analysis of Trajectory Recovery

Varying Level of Sparsity. As mentioned, given a trajectory with the target ϵ sampling rate, we generate its sparse trajectory by random sampling with time interval ϵ/γ , where γ controls the level of sparsity. We vary γ from 0.1 to 0.5, and report the accuracy results in Fig. 7. As trajectories become sparser with a smaller γ , all methods achieve degraded performance.




Fig. 5: Inference Time per 1000 Trajectory Recoveries (in seconds) (★ marks the best competitor in Table III).




Fig. 6: Training Time per Epoch (in minutes) of Trajectory Recovery Methods (★ marks the best competitor in Table III).




Fig. 7: Trajectory Recovery with Varied Levels of Sparsity
TABLE IV: Ablation Result of TRMMA by Accuracy (%)

	PT	XA	BJ	CD
TRMMA	57.83	78.95	53.71	75.28
TRMMA-HMM	53.54	76.81	49.57	70.63
TRMMA-Near	47.01	65.81	43.66	56.22
MMA+linear	43.74	68.99	41.72	62.82
Nearest+linear	35.45	58.03	33.97	47.61
TRMMA-DF	54.83	77.62	50.73	73.91
TRMMA-C	56.85	78.63	52.13	74.96
TRMMA-DI	51.02	71.47	45.83	69.15

Nevertheless, the performance gap between TRMMA and the competitors maintains under all settings on all datasets. For instance, on PT dataset, when $\gamma = 0.5$, TRMMA has 0.8392 accuracy, 4.57% higher than TERI with 0.7935 accuracy.

Ablation Study. We ablate the techniques in TRMMA and report the accuracy in Table IV. TRMMA-Near is TRMMA without MMA in Section IV, but with nearest segment of a GPS point p_i as its segment a_i . TRMMA-HMM is TRMMA without MMA, but with HMM [28]. TRMMA-DF is TRMMA without the DualFormer encoding module in Section V, *i.e.*, using \mathbf{R} as \mathbf{H} in Fig. 4. TRMMA-C is TRMMA with MMA that does not consider the candidates in \mathcal{C}_{p_i} as the context to get \mathbf{p}_i in the point embedding module in Section IV-B. TRMMA-DI

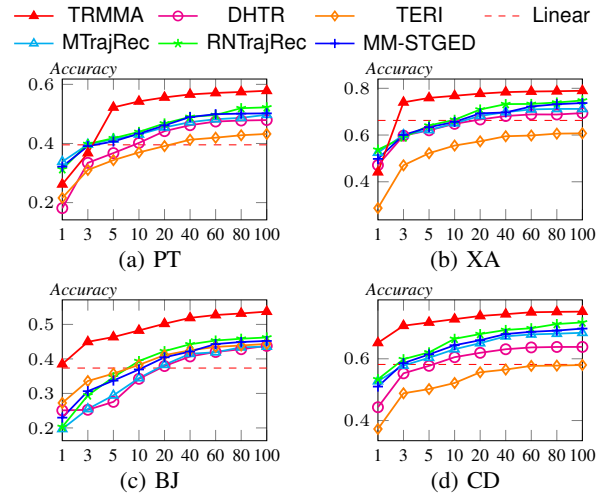


Fig. 8: Trajectory Recovery When Varying Training Data Size

is TRMMA with MMA that does not consider the directional information. MMA+linear is MMA with linear interpolation for recovery. Nearest+linear is Nearest with linear interpolation. In Table IV, the accuracy of TRMMA is higher than all ablated versions, validating the effectiveness of our techniques.

Robustness v.s. Amount of Training Data. We evaluate the robustness of TRMMA and baselines using between 1% and 100% of the total training data, and report accuracy results in Fig. 8. Note that Linear does not involve training and serves as a benchmark. Typically, reducing the amount of training data degrades a method's accuracy. The performance gap between TRMMA and the competitors maintains under most settings. In Fig. 8(c,d) for BJ and CD, TRMMA consistently outperforms existing methods under all settings. In Fig. 8(a,b) for PT and XA, TRMMA achieves the best performance after using 3% and 1% of the training data, respectively.

D. Effectiveness and Efficiency of Map Matching

Effectiveness. Table V presents the map matching results of our method, MMA, compared to competitors. For trajectory T , the metric values are based on segments in the returned route

TABLE V: Effectiveness of Map Matching. A larger value is better (in percentage). Best is in bold and runner-up is underlined.

Methods	PT				XA				BJ				CD			
	Precision	Recall	F1	Jaccard	Precision	Recall	F1	Jaccard	Precision	Recall	F1	Jaccard	Precision	Recall	F1	Jaccard
Nearest	80.42	85.42	82.42	74.55	79.01	89.79	82.69	75.03	66.81	71.86	68.20	59.93	72.29	87.24	77.32	69.10
FMM	86.34	83.71	84.74	78.08	93.60	91.85	92.49	88.84	72.51	70.36	70.69	63.82	89.14	88.39	88.34	83.94
LHMM	89.80	87.06	88.20	82.37	<u>95.53</u>	94.14	94.62	91.84	75.30	72.35	73.08	65.34	91.19	90.69	90.57	87.10
RNTrajRec	89.70	89.46	89.10	84.29	<u>93.15</u>	94.10	93.03	89.73	<u>78.82</u>	76.64	76.80	<u>70.30</u>	89.46	91.17	89.45	85.48
DeepMM	<u>91.34</u>	<u>90.95</u>	<u>90.88</u>	<u>86.22</u>	95.40	<u>95.14</u>	<u>95.06</u>	<u>92.23</u>	<u>78.29</u>	<u>77.66</u>	<u>76.99</u>	<u>69.41</u>	<u>94.99</u>	<u>94.67</u>	<u>94.58</u>	<u>91.54</u>
GraphMM	87.01	88.84	87.26	79.13	92.84	94.62	92.75	87.06	75.39	73.84	72.32	62.82	88.53	92.56	89.31	82.23
MMA	94.46	94.53	94.35	91.53	97.20	97.97	97.36	95.97	82.17	81.08	80.92	75.28	96.27	97.51	96.54	94.94




Fig. 9: Inference Time of Map Matching per 1000 Trajectories (in seconds) (★ marks the best competitor in Table V).




Fig. 10: Training Time per Epoch (in minutes) of Map Matching (★ marks the best competitor in Table V).

\mathcal{R} , not the recovered \mathcal{T}_ϵ in trajectory recovery above. As shown in Table V, MMA outperforms all methods across all four datasets under all metrics. For instance, on PT, MMA achieves a 91.53% Jaccard score, surpassing DeepMM’s 86.22% by 5.31%. On BJ, MMA has a recall of 81.08%, 3.42% higher than DeepMM’s 77.66%. This superior performance validates the effectiveness of our analysis and formulation in Section IV-A, treating it as a classification task over a small candidate segment set to identify the segment of a GPS point, and demonstrates the power of the proposed point and candidate segment embedding techniques in Section IV-B.

Efficiency. Fig. 9 reports the average inference time for map matching 1000 trajectories. Our method, MMA, is significantly faster than all other methods, often by orders of magnitude. For example, on the PT dataset, MMA requires only 0.19 seconds, whereas DeepMM requires 5.67 seconds, indicating a 29-fold speedup. Fig. 10 reports the training time per epoch, except for FMM, which does not require training. Similarly, MMA is faster to train than existing methods. The efficiency of MMA in both inference and training demonstrates the power of the proposed techniques in Section IV. Specifically, given a GPS point, we focus on a small candidate segment set to identify its segment in Section IV-A, and then develop efficient and effective embedding techniques in Section IV-B to capture intrinsic patterns for map matching.

Map Matching with Varied Levels of Sparsity. As mentioned, given a trajectory with the target ϵ sampling rate, we




Fig. 11: Map Matching with Varied Levels of Sparsity

generate its sparse trajectory by random sampling with time interval ϵ/γ , where γ controls the level of sparsity. We vary γ from 0.1 to 0.5, and report the F1 scores of map matching in Fig. 11. As trajectories become sparser with a smaller γ , all methods achieve degraded performance. Nevertheless, MMA is the best across all sparsity levels on all datasets, demonstrating the effectiveness of our techniques.

VII. CONCLUSION

We present TRMMA and MMA, efficient methods for accurate trajectory recovery and map matching, respectively, with MMA serving as the initial step for TRMMA. MMA formulates a classification task to map a GPS point to road segment within a small candidate set, and generates effective embeddings for accurate map matching. TRMMA then focuses on the segments of the route identified by MMA for recovery, avoiding the need to evaluate all road segments. It employs a dual-transformer encoding process to capture latent patterns and a decoding technique to predict position ratios and road segments of missing points. Extensive experiments on large real-world datasets demonstrate that TRMMA and MMA outperform existing methods in both result quality and efficiency.

ACKNOWLEDGMENT

The work described in this paper was supported by grants from the Research Grants Council of Hong Kong Special Administrative Region, China (No. PolyU 25201221, PolyU 15205224, PolyU 152043/23E). Jieming Shi is supported by NSFC No. 62202404, Otto Poon Charitable Foundation Smart Cities Research Institute (SCRI) P0051036-P0050643. This work is supported by Tencent Technology Co., Ltd. P0048511.

REFERENCES

- [1] J. Jain, V. Bagadia, S. Manchanda, and S. Ranu, "Neuomlr: Robust & reliable route recommendation on road networks," in *NeurIPS*, 2021, pp. 22 070–22 082.
- [2] W. Tian, J. Shi, S. Luo, H. Li, X. Xie, and Y. Zou, "Effective and efficient route planning using historical trajectories on road networks," *PVLDB*, vol. 16, no. 10, pp. 2512–2524, 2023.
- [3] H. Yuan, G. Li, Z. Bao, and L. Feng, "Effective travel time estimation: When historical trajectories over road networks matter," in *SIGMOD*, 2020, pp. 2135–2149.
- [4] H. Yuan, G. Li, and Z. Bao, "Route travel time estimation on A road network revisited: Heterogeneity, proximity, periodicity and dynamicity," *PVLDB*, vol. 16, no. 3, pp. 393–405, 2022.
- [5] H. Lan, J. Xie, Z. Bao, F. Li, W. Tian, F. Wang, S. Wang, and A. Zhang, "VRE: A versatile, robust, and economical trajectory data system," *PVLDB*, vol. 15, no. 12, pp. 3398–3410, 2022.
- [6] Z. Fang, S. Gong, L. Chen, J. Xu, Y. Gao, and C. S. Jensen, "Ghost: A general framework for high-performance online similarity queries over distributed trajectory streams," *SIGMOD*, vol. 1, no. 2, pp. 173:1–173:25, 2023.
- [7] J. Jin, P. Cheng, L. Chen, X. Lin, and W. Zhang, "Efficient non-learning similar subtrajectory search," *PVLDB*, vol. 16, no. 11, pp. 3111–3123, 2023.
- [8] Z. Fang, L. Pan, L. Chen, Y. Du, and Y. Gao, "MDTP: A multi-source deep traffic prediction framework over spatio-temporal trajectory data," *PVLDB*, vol. 14, no. 8, pp. 1289–1297, 2021.
- [9] H. Yuan, G. Cong, and G. Li, "Nuhuo: An effective estimation model for traffic speed histogram imputation on A road network," *PVLDB*, vol. 17, no. 7, pp. 1605–1617, 2024.
- [10] Y. Chen, H. Zhang, W. Sun, and B. Zheng, "Rntrajrec: Road network enhanced trajectory recovery with spatial-temporal transformer," in *ICDE*, 2023, pp. 829–842.
- [11] W. Shi, J. Xu, J. Fang, P. Chao, A. Liu, and X. Zhou, "LHMM: A learning enhanced HMM model for cellular trajectory map matching," in *ICDE*, 2023, pp. 2429–2442.
- [12] J. Yuan, Y. Zheng, C. Zhang, X. Xie, and G. Sun, "An interactive-voting based map matching algorithm," in *MDM*, 2010, pp. 43–52.
- [13] Y. Liu, Q. Ge, W. Luo, Q. Huang, L. Zou, H. Wang, X. Li, and C. Liu, "Graphmm: Graph-based vehicular map matching by leveraging trajectory and road correlations," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 1, pp. 184–198, 2024.
- [14] H. Ren, S. Ruan, Y. Li, J. Bao, C. Meng, R. Li, and Y. Zheng, "Mtrajrec: Map-constrained trajectory recovery via seq2seq multi-task learning," in *KDD*, 2021, pp. 1410–1419.
- [15] D. Shi, Y. Tong, Z. Zhou, K. Xu, Z. Wang, and J. Ye, "Graph-constrained diffusion for end-to-end path planning," in *ICLR*, 2024.
- [16] O. Pink and B. Hummel, "A statistical approach to map matching using road network geometry, topology and vehicular motion constraints," in *ITSC*, 2008, pp. 862–867.
- [17] P. Newson and J. Krumm, "Hidden markov map matching through noise and sparseness," in *SIGSPATIAL*, 2009, pp. 336–343.
- [18] Z. Shen, K. Yang, X. Zhao, J. Zou, W. Du, and J. Wu, "Dmm: A deep reinforcement learning based map matching framework for cellular data," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 10, pp. 5120–5137, 2024.
- [19] T. Wei, Y. Lin, Y. Lin, S. Guo, L. Zhang, and H. Wan, "Micro-macro spatial-temporal graph-based encoder-decoder for map-constrained trajectory recovery," *IEEE Trans. Knowl. Data Eng.*, vol. 36, no. 11, pp. 6574–6587, 2024.
- [20] J. Wang, N. Wu, X. Lu, W. X. Zhao, and K. Feng, "Deep trajectory recovery with fine-grained calibration using kalman filter," *IEEE Trans. Knowl. Data Eng.*, vol. 33, no. 3, pp. 921–934, 2021.
- [21] Y. Chen, G. Cong, and C. Anda, "TERI: an effective framework for trajectory recovery with irregular time intervals," *PVLDB*, vol. 17, no. 3, pp. 414–426, 2023.
- [22] M. Musleh and M. F. Mokbel, "A demonstration of KAMEL: A scalable bert-based system for trajectory imputation," in *SIGMOD*, 2023, pp. 191–194.
- [23] M. M. Elsharif, K. Isufaj, and M. F. Mokbel, "Network-less trajectory imputation," in *SIGSPATIAL*, 2022, pp. 8:1–8:10.
- [24] T. Xia, Y. Qi, J. Feng, F. Xu, F. Sun, D. Guo, and Y. Li, "Attnmove: History enhanced trajectory recovery via attentional network," in *AAAI*, 2021, pp. 4494–4502.
- [25] H. Sun, C. Yang, L. Deng, F. Zhou, F. Huang, and K. Zheng, "Periodic-move: Shift-aware human mobility recovery with graph neural network," in *CIKM*, 2021, pp. 1734–1743.
- [26] D. Xi, F. Zhuang, Y. Liu, J. Gu, H. Xiong, and Q. He, "Modelling of bi-directional spatio-temporal dependence and users' dynamic preferences for missing POI check-in identification," in *AAAI*, 2019, pp. 5458–5465.
- [27] Y. Lou, C. Zhang, Y. Zheng, X. Xie, W. Wang, and Y. Huang, "Map-matching for low-sampling-rate GPS trajectories," in *SIGSPATIAL*, 2009, pp. 352–361.
- [28] C. Yang and G. Gidófalvi, "Fast map matching, an algorithm integrating hidden markov model with precomputation," *Int. J. Geogr. Inf. Sci.*, vol. 32, no. 3, pp. 547–570, 2018.
- [29] K. Zheng, Y. Zheng, X. Xie, and X. Zhou, "Reducing uncertainty of low-sampling-rate trajectories," in *ICDE*, 2012, pp. 1144–1155.
- [30] P. Banerjee, S. Ranu, and S. Raghavan, "Inferring uncertain trajectories from partial observations," in *ICDM*, 2014, pp. 30–39.
- [31] Y. Gong, E. Chen, X. Zhang, L. M. Ni, and J. Zhang, "Antmapper: An ant colony-based map matching approach for trajectory-based applications," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 2, pp. 390–401, 2018.
- [32] J. Feng, Y. Li, K. Zhao, Z. Xu, T. Xia, J. Zhang, and D. Jin, "Deepmm: Deep learning based map matching with data augmentation," *IEEE Trans. Mob. Comput.*, vol. 21, no. 7, pp. 2372–2384, 2022.
- [33] L. Jiang, C. Chen, and C. Chen, "L2MM: learning to map matching with deep models for low-quality GPS trajectory data," *ACM Trans. Knowl. Discov. Data*, vol. 17, no. 3, pp. 39:1–39:25, 2023.
- [34] X. Li, K. Zhao, G. Cong, C. S. Jensen, and W. Wei, "Deep representation learning for trajectory similarity computation," in *ICDE*, 2018, pp. 617–628.
- [35] D. Yao, G. Cong, C. Zhang, and J. Bi, "Computing trajectory similarity in linear time: A generic seed-guided neural metric learning approach," in *ICDE*, 2019, pp. 1358–1369.
- [36] P. Yang, H. Wang, Y. Zhang, L. Qin, W. Zhang, and X. Lin, "T3S: effective representation learning for trajectory similarity computation," in *ICDE*, 2021, pp. 2183–2188.
- [37] D. Yao, H. Hu, L. Du, G. Cong, S. Han, and J. Bi, "Trajgat: A graph-based long-term dependency modeling approach for trajectory similarity computation," in *KDD*, 2022, pp. 2275–2285.
- [38] Y. Chang, J. Qi, Y. Liang, and E. Tanin, "Contrastive trajectory similarity learning with dual-feature attention," in *ICDE*, 2023, pp. 2933–2945.
- [39] Z. Fang, Y. Du, X. Zhu, D. Hu, L. Chen, Y. Gao, and C. S. Jensen, "Spatio-temporal trajectory similarity learning in road networks," in *KDD*, 2022, pp. 347–356.
- [40] P. Han, J. Wang, D. Yao, S. Shang, and X. Zhang, "A graph-based approach for trajectory similarity computation in spatial networks," in *KDD*, 2021, pp. 556–564.
- [41] B. A. Renfro, M. Stein, E. B. Reed, and A. Finn, "An analysis of global positioning system standard positioning service performance for 2022," *Space and Geophysics Laboratory Applied Research Laboratories, The University of Texas at Austin*. [Online]. Available: <https://www.gps.gov/systems/gps/performance/>
- [42] S. T. Leutenegger, J. M. Edgington, and M. A. López, "STR: A simple and efficient algorithm for r-tree packing," in *ICDE*, 1997, pp. 497–506.
- [43] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *KDD*, 2016, pp. 855–864.
- [44] X. Glorot, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *AISTATS*, 2011, pp. 315–323.
- [45] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NIPS*, 2017, pp. 5998–6008.
- [46] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," in *EMNLP*, 2014, pp. 103–111.
- [47] "Porto dataset," <https://www.kaggle.com/c/pkdd-15-predict-taxi-service-trajectory-i/data>.
- [48] Z. Shang, G. Li, and Z. Bao, "DITA: distributed in-memory trajectory analytics," in *SIGMOD*, 2018, pp. 725–740.
- [49] "Openstreetmap," <https://www.openstreetmap.org/>.