



RED: Effective Trajectory Representation Learning with Comprehensive Information

Silin Zhou
University of Electronic Science and
Technology of China
zhousilinx@gmail.com

Shuo Shang*
University of Electronic Science and
Technology of China
jedi.shang@gmail.com

Lisi Chen
University of Electronic Science and
Technology of China
lchen012@e.ntu.edu.sg

Christian S. Jensen
Aalborg University
csj@cs.aau.dk

Panos Kalnis
KAUST
panos.kalnis@kaust.edu.sa

ABSTRACT

Trajectory representation learning (TRL) maps trajectories to vectors that can then be used for various downstream tasks, including trajectory similarity computation, trajectory classification, and travel-time estimation. However, existing TRL methods often produce vectors that, when used in downstream tasks, yield insufficiently accurate results. A key reason is that they fail to utilize the comprehensive information encompassed by trajectories. We propose a self-supervised TRL framework, called RED, which effectively exploits multiple types of trajectory information. Overall, RED adopts the Transformer as the backbone model and masks the constituting paths in trajectories to train a masked autoencoder (MAE). In particular, RED considers the moving patterns of trajectories by employing a *Road-aware masking strategy* that retains key paths of trajectories during masking, thereby preserving crucial information of the trajectories. RED also adopts a *spatial-temporal-user joint Embedding* scheme to encode comprehensive information when preparing the trajectories as model inputs. To conduct training, RED adopts *Dual-objective task learning*: the Transformer encoder predicts the next segment in a trajectory, and the Transformer decoder reconstructs the entire trajectory. RED also considers the spatial-temporal correlations of trajectories by modifying the attention mechanism of the Transformer. We compare RED with 9 state-of-the-art TRL methods for 4 downstream tasks on 3 real-world datasets, finding that RED can usually improve the accuracy of the best-performing baseline by over 5%.

PVLDB Reference Format:

Silin Zhou, Shuo Shang, Lisi Chen, Christian S. Jensen, and Panos Kalnis. RED: Effective Trajectory Representation Learning with Comprehensive Information. PVLDB, 18(2): 80 - 92, 2024.
doi:10.14778/3705829.3705830

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/slzhou-xy/RED>.

*The corresponding author.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 18, No. 2 ISSN 2150-8097.
doi:10.14778/3705829.3705830

1 INTRODUCTION

With the proliferation of GPS-enabled devices (e.g., smartphones, navigators, and digital watches), large amounts of trajectories are collected, recording the movements of pedestrians or vehicles. These trajectories serve as the foundation for many applications such as traffic prediction [49], urban planning [33], and transportation optimization [38]. However, as sequences of timestamped locations, trajectories require specific techniques for management and analysis. For example, trajectory similarity computation often relies on dynamic programming [3, 9], resulting in computational costs that increase quadratically with trajectory length.

Recently, *trajectory representation learning* (TRL), which maps each trajectory to a vectors embedding, has attracted attention as a general preprocessing technique [20, 31, 42, 43]. The advantage is that the learned vectors can be used directly in many downstream tasks, including trajectory classification [24], travel time estimation [27], and trajectory similarity computation [51], with standard vector processing techniques. For instance, trajectory similarity can be calculated as the distance between the vectors of two trajectories, with a cost that does not increase with trajectory length.

Early TRL methods [23, 45] usually target a specific task. For example, Traj2vec [45] uses Recurrent Neural Networks (RNNs) to map trajectories to vectors and tailors model design for trajectory clustering. The problem of these methods is that their vectors do not work well when used in other downstream tasks. Subsequent TRL methods [10, 13, 41] utilize self-supervised learning (SSL) due to the strong generalization capabilities of SSL. In particular, SSL relies on a generic pre-training task to distill data information, e.g., by manually masking a portion of the trajectories and learning to recover the trajectories. For example, Trembr [13] uses RNNs [19] with an encoder-decoder architecture. The encoder embeds trajectories into vectors, and the decoder recovers trajectories from vectors. Recent TRL methods [20, 31, 43] adopt contrastive learning (CL). CL conducts data augmentations to generate positive and negative samples for each trajectory, and the model is trained to make positive trajectory pairs more similar than negative pairs. For instance, START [20] features several data augmentation techniques, such as randomly masking segments in the trajectories and trimming the trajectories to get sub-trajectories. However, data augmentation techniques may not generalize across datasets because their performance depends strongly on the dataset under consideration.

A trajectory encompasses multiple kinds of information, including road, user, spatial-temporal, travel, and movement, but existing

Table 1: Trajectory information utilized by TRL methods.

Model	Road	User	Spatial	Temporal	Travel	Movement
T2vec [23]	✗	✗	✗	✗	✗	✗
Traj2vec [45]	✗	✗	✗	✗	✗	✓
Trembr [13]	✓	✗	✗	✓	✗	✗
PIM [41]	✓	✗	✓	✗	✗	✗
Toast [10]	✓	✗	✓	✗	✓	✓
JCLRNT [31]	✓	✗	✓	✗	✗	✗
START [20]	✓	✓	✓	✓	✓	✗
LightPath [43]	✓	✗	✗	✗	✓	✗
RED (ours)	✓	✓	✓	✓	✓	✓

TRL methods do not utilize them comprehensively, resulting in trajectory representations with reduced accuracy. In particular, *road information*, obtained by map matching [39], tells the road segments each trajectory passes and is less noisy than the raw trajectory points. *User information* ties each trajectory to a user and encodes user preference. *Spatial information* refers to the spatial locations of the road segments, and a road segment usually has multiple adjacent road segments with high transition probabilities [16]. *Temporal information* captures the progress over time when moving along a trajectory and can be used to estimate travel time and analyze traffic flow [20]. *Travel semantics* [10] refers to the overall travel statistics of a trajectory, such as travel time and distance. *Movement semantics* [45] refers to the local properties of trajectories, such as whether congestion is encountered at a location.

Table 1 summarizes how trajectory information is utilized by existing TRL methods. Road information is ignored by a few methods, e.g., Traj2vec [45] models trajectories as point sequences without considering the underlying road properties. Spatial information is usually considered by using the road network. For example, Toast [10] and PIM [41] apply the graph neural network to capture the topology of the road network of a city. Temporal information is typically used to encode model inputs, e.g., Trembr [13] includes the timestamp of each road in a trajectory as input to its RNN. Some methods incorporate the travel semantics into trajectories to extract global information. For instance, Toast [10], START [20], and LightPath [43] utilize the learned trajectory representations to recover raw trajectories. Only a few methods consider the moving semantics for local information, e.g., Traj2vec and Toast use context windows to extract short sequences from trajectories. START [20] is the only existing method that exploits user information.

To go beyond existing studies, we propose RED, a self-supervised learning method that aims to improve TRL by using more comprehensive trajectory information. For overall architecture, RED adopts Transformer as the backbone model and masks the segments of path trajectories to train a masked autoencoder. Regarding trajectory information, we first consider and analyze the *local driving pattern* of road path trajectories and propose a novel road-aware masking strategy. In particular, we split a path trajectory into the key path and the mask path, where the key path contains most of the semantic information of a trajectory to assist in the final trajectory representation. Then, we design a spatial-temporal-user joint embedding scheme to fuse *spatial*, *temporal*, and *user* information of road paths for model inputs. We also co-encode the *time* and

segment type to capture periodic information of travel. To train RED, we utilize two training tasks, including next segment prediction and trajectory reconstruction. Next segment prediction predicts the next road segment from previous key paths of a trajectory, while trajectory reconstruction reconstructs the original trajectory using the learned representations of the key paths and mask paths. Both tasks exploit the *travel* information in the trajectory representations. Moreover, we also introduce two strategies to improve trajectory modeling: virtual tokens and a spatial-temporal correlation enhanced attention module. Virtual tokens solve the path misalignment issues when using the Transformer for trajectories, and the attention module enhances Transformer attention to extract the spatial-temporal information of trajectories.

To evaluate RED, we set experiments on 3 real-world datasets and compare with 9 state-of-the-art TRL methods, includes supervised and self-supervised methods, as well as 7 non-learning trajectory similarity computation methods. We consider 4 popular downstream tasks: travel time estimation, trajectory classification, trajectory similarity computation, and most similar trajectory retrieval. The results show that RED outperforms all existing methods in terms of accuracy across the tasks and datasets. In particular, compared with the best-performing baseline, the average accuracy improvement of RED at travel time estimation, trajectory classification, and trajectory similarity computation are 7.03%, 12.11%, and 20.02%, respectively. An ablation study suggests that all the types of trajectory information we utilize contribute to the accuracy of TRL and that our model designs are effective.

To summarize, we make three main contributions.

- We identify design limitations of existing TRL methods, including the use of random masking, insufficient utilization of trajectory information, and limited supervision signals for model training.
- We propose RED as a more effective self-supervised learning framework for TRL. RED features a road-aware mask, spatial-temporal-user embedding, and dual-objective learning to address the limitations of existing methods.
- We report on an extensive evaluation and comparison of RED with state-of-the-art TRL methods. The results show that RED is capable of higher accuracy than the existing methods across important downstream tasks.

2 PROBLEM AND BACKGROUND

2.1 Preliminaries

Definition 1 (GPS Trajectory). A GPS trajectory, denoted as \mathcal{T}_{gps} , is a sequence of points collected at a fixed sampling time interval rate. Each point in \mathcal{T}_{gps} takes the form of $p_i = (x_i, y_i, t_i)$, where x_i , y_i , and t_i denote longitude, latitude, and timestamp, respectively.

Definition 2 (Road Network). A road network is modeled as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where \mathcal{V} denotes the road segment set in the road network, and $\mathcal{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is the adjacency matrix that represents the connectivity between road segments. $\mathcal{A}[i, j] = 1$ if and only if road segments v_i and v_j are directly connected, otherwise $\mathcal{A}[i, j] = 0$. Under this definition, a trajectory can be extracted as a sequence of road segments that it passes through.

Definition 3 (Path Trajectory). A path trajectory \mathcal{T} is a time-ordered sequence of road segments that is generated from \mathcal{T}_{gps} by map matching. That is, $\mathcal{T} = \langle \tau_1, \tau_2, \dots, \tau_{|\mathcal{T}|} \rangle$ contains the $|\mathcal{T}|$ road

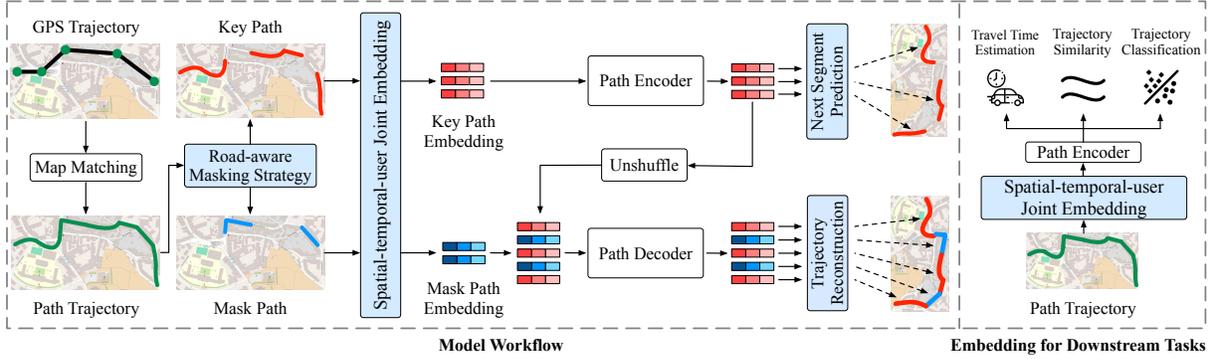


Figure 1: Overall architecture of the RED framework.

segments passed by \mathcal{T}_{gps} . Each element $\tau_i = (v_i, t_i) \in \mathcal{T}$ models that the trajectory passes road segment v_i at timestamp t_i .

2.2 Problem Statement

Given a set of path trajectories $\mathcal{D} = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_{|\mathcal{D}|}\}$ and a road network graph \mathcal{G} , we aim to compute a generic vector representation $\tilde{p}_i \in \mathbb{R}^l$ for each path trajectory $\mathcal{T}_i \in \mathcal{D}$, where l is the dimension of the trajectory vector representation. We expect these vectors to achieve a high accuracy for various downstream tasks:

- **Trajectory similarity computation:** Given trajectories \mathcal{T}_a and \mathcal{T}_b , trajectory similarity computation calculates a score capturing the similarity between \mathcal{T}_a and \mathcal{T}_b .
- **Most similar trajectory retrieval:** Given a query trajectory \mathcal{T}_a and a trajectory dataset \mathcal{D} , this task finds the trajectory $\mathcal{T}_b \in \mathcal{D}$ that is the most similar to \mathcal{T}_a .
- **Trajectory classification:** Given a trajectory \mathcal{T}_a , this task assigns \mathcal{T}_a to a category, e.g., a user ID.
- **Travel time estimation:** Given trajectory \mathcal{T}_a without temporal information, this task predicts the travel time of \mathcal{T}_a .

2.3 Background on Machine Learning

Transformer. The Transformer [34] architecture has demonstrated impressive effectiveness at NLP. It models a token sequence, where a *token* is a word identifier, that can be converted into a learnable vector. A token does not need to be associated with a word. Thus, a language model BERT [11] designs a [CLS] token to represent the summary of a sentence. The Transformer architecture encompasses multiple stacked blocks, each block contains a multi-head self-attention and a feed-forward network. Multi-head self-attention learns the inter-relationships between different elements in a sequence. A feed-forward network further enhances feature extraction and the model’s expressiveness. Compared to RNN-based methods [19], the Transformer can accommodate billions of model parameters by stacking blocks and computes a long token sequence in parallel, without the need to iterate through time steps.

An input token sequence embedding is given by $\mathbf{X} \in \mathbb{R}^{n \times l}$, where n is the sequence length, and l is the dimensionality of vectors. The vanilla Transformer first adds a position encoding to \mathbf{X} and then transforms \mathbf{X} to a query matrix $\mathbf{Q} = \mathbf{X}\mathbf{W}^Q \in \mathbb{R}^{n \times l}$, a key matrix $\mathbf{K} = \mathbf{X}\mathbf{W}^K \in \mathbb{R}^{n \times l}$, and a value matrix $\mathbf{V} = \mathbf{X}\mathbf{W}^V \in \mathbb{R}^{n \times l}$, $\mathbf{W}^Q \in \mathbb{R}^{l \times l}$,

$\mathbf{W}^K \in \mathbb{R}^{l \times l}$, and $\mathbf{W}^V \in \mathbb{R}^{l \times l}$ are learnable matrices. Next a self-attention correlation matrix is given by $\mathbf{A} = \frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{l}} \in \mathbb{R}^{n \times n}$, where $\alpha_{i,j} \in \mathbf{A}$ is the attention score between element- i and element- j of \mathbf{X} . The sequence output is computed as follows:

$$\mathbf{X}' = \text{softmax}(\mathbf{A})\mathbf{V}, \quad (1)$$

where \mathbf{X}' is the input to the subsequent feed-forward network.

Masked Autoencoder. A masked autoencoder [17] is an encoder-decoder model from computer vision (CV) that uses the Transformer. An image is first divided into small patches, where each patch can be seen as a word and can be placed in the Transformer for training. A masked autoencoder randomly masks partial patches and has its encoder act on visible patches. A lightweight decoder is used to reconstruct the raw image pixels based on mask tokens and potential representations from the encoder. This reconstruction task is an effective and meaningful self-supervised task.

In our study, we view a path trajectory as a sentence, and we view each segment as a word that is transformed into a token. Our base model is a masked autoencoder. We use prior trajectory knowledge to partition mask data, and we introduce comprehensive trajectory information to improve the accuracy of trajectory representation.

3 METHOD OVERVIEW

The left plot of Figure 1 shows the workflow of RED during training, which involves three key modules, i.e., *road-aware masking*, *spatial-temporal-user joint embedding*, and *dual-objective learning*.

Road-aware Masking. As the first step, the map matching algorithm [39] is conducted to transform a GPS trajectory to a path trajectory. Then, the road-aware masking strategy splits a path trajectory \mathcal{T} into *key path set* \mathcal{T}^k and *mask path set* \mathcal{T}^m , where $\mathcal{T} = \mathcal{T}^k \cup \mathcal{T}^m$, based on the sampling rate and driving pattern of the trajectory. The key paths in \mathcal{T}^k encompass the driving patterns of the trajectory while the mask paths in \mathcal{T}^m are less crucial and thus are masked for the encoder. In comparison, existing methods adopt the random masking strategy, which randomly masks the paths. Trajectory information will be lost if some key paths are masked. Moreover, random masking also requires extensive tuning of the mask ratio while our road-aware masking does not.

Joint Embedding. Next, each segment of a trajectory is encoded as a raw embedding to serve as framework input, which can be

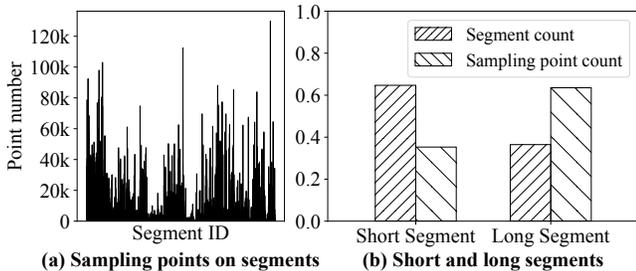


Figure 2: Trajectory sample statistics of the Porto dataset.

expressed as $\mathbf{X} = \text{Emb}(\mathcal{T})$. The spatial-temporal-user joint embedding integrates comprehensive information of the trajectory, including spatial features from the road network graph, temporal features from the traffic and travel patterns over time, user features from user ID, and segment features from the road types.

Dual-objective learning. We train RED with two objectives on to provide sufficient supervision signals for learning. In particular, the encoder maps the key paths in \mathcal{T}^k to key path embedding $\widehat{\mathbf{X}}^k = \text{Encoder}(\mathbf{X}^k)$. As the key paths capture the crucial driving patterns of a trajectory, we use the next segment prediction objective \mathcal{L}^{nsp} for the encoder to predict the next key path given the previous key paths. This resembles next token prediction in NLP. The decoder takes the path embeddings generated by the encoder, merges the key paths and mask paths according to their original order in the trajectory, and reconstructs the trajectory using the embeddings, i.e., $\widehat{\mathbf{X}}^{\mathcal{T}} = \text{Decoder}(\text{Unshuffle}(\widehat{\mathbf{X}}^k, \mathbf{X}^m))$. Thus, we use the trajectory reconstruction objective \mathcal{L}^{tr} to train the decoder. Therefore, the overall objective function is:

$$\mathcal{L} = \lambda_1 \mathcal{L}^{nsp}(\widehat{\mathbf{X}}^k) + (1 - \lambda_1) \mathcal{L}^{tr}(\widehat{\mathbf{X}}^{\mathcal{T}}), \quad (2)$$

where λ_1 controls the weight of the loss terms. To tackle sparsity on trajectories, i.e., GPS points recorded by the device during the driving process are discontinuous, inconsistent, or missing, we incorporate comprehensive information, more effective supervision signals, and a GNN for the road network.

As shown in the right of Figure 1, during inference, we only use path encoder and feed the embeddings of the complete trajectory as input to get the trajectory vector for downstream tasks.

4 KEY DESIGNS OF RED

We proceed to detail the key innovations of RED, which includes the road-aware masking strategy, the spatial-temporal-user joint embedding, the dual-objective task learning, and the enhanced trajectory modeling techniques.

4.1 Road-aware Masking Strategy

Mask-based self-supervised learning is popular for TRL, which first masks some road segments of a path trajectory, and then learns to reconstruct the masked road segments. However, existing methods use the random masking strategy, which has two problems. First, it may discard essential road segments of a trajectory, leading to some key information loss and inaccurate trajectory representations. Second, the masking ratio requires extensive tuning, and one

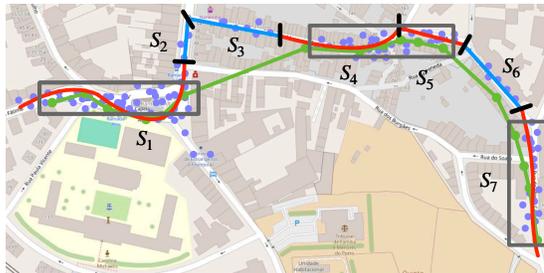


Figure 3: Illustration of the road-aware masking.

masking ratio may not suit different trajectories. We introduce a *road-aware masking strategy* that avoids the two problems.

To motivate the road-aware masking strategy, we illustrate the sampling statistics of real-world trajectories in Figure 2. The dataset is Porto, and the sampling interval is about 15 seconds for consecutive points in a trajectory. Since the average length of the road segments in the road network is about 85 meters, we call road segments below 85 meters as short segments and the converse long segments. Figure 2(a) shows that the distribution of the sampling points over the road segments is highly skewed, i.e., some segments have many sampling points while some segments have only a few sampling points. Intuitively, segments with many sampling points are important because they are passed by many trajectories and may take an important role in the road network (e.g., connecting two important areas). Figure 2(b) shows that there are more short road segments than long road segments in the road network but more sampling points reside on the long road segments than the short segments. This suggests that the long segments are more important than the short segments because the trajectories spend more time on them. Overall, Figure 2 suggests that the road segments are different in their importance. Thus, by treating all road segments equally, random masking may lose important segments and hinder model training.

Driven by the observations from Figure 2, our road-aware masking strategy treats road segments with many sampling points and long road segments as *key paths*, which will not be masked. To avoid parameter tuning, we call a segment hot (resp. long) when its sampling points (resp. length) are larger than the average over the road segments. Figure 3 provides a running example of road-aware masking. The original trajectory is plotted in green color. After map matching, we obtain 7 road segments, i.e., $\langle s_1, s_2, \dots, s_7 \rangle$. The key paths, i.e., $\langle s_1, s_4, s_5, s_7 \rangle$, are marked in red, while the mask paths, i.e., $\langle s_2, s_3, s_6 \rangle$ are marked in blue. The key paths either have many sampling points (e.g., s_1) or a large length (e.g., s_7). s_6 is a mask path because it has only a few sampling points, while s_2 and s_3 are mask paths because they are generated by map matching and do not contain real sampling points from the trajectory.

4.2 Spatial-temporal-user Joint Embedding

Path trajectories encompass a wealth of information, including user ID, road segment, and time information. All of this information is important when learning a versatile representation for trajectories. Consequently, we propose a *spatial-temporal-user joint embedding* to encode all the information for model input. Formally, given a

segment v_i of a path trajectory, we use Equation 3 for encoding:

$$\mathbf{x}_i = \mathbf{h}_i + \mathbf{t}_i + \mathbf{u}_i, \quad (3)$$

where \mathbf{h}_i , \mathbf{t}_i , and \mathbf{u}_i denote the spatial, time and user encodings of segment v_i , which are described next.

Spatial Encoding \mathbf{h}_i . The road segments of a city form a complex structure, which can be captured by the road network graph \mathcal{G} . We observe that the properties of a segment are affected by itself and its adjacent segments. Therefore, we apply a graph neural network (GNN) to learn embeddings for road segments based on the road network graph \mathcal{G} . In particular, GNN computes an embedding for each node in the graph by aggregating the embeddings of its neighbors. We choose the Graph Attention Network (GAT) [35], which adopts an attention mechanism for neighbor aggregation.

We feed multiple attributes of a segment as the initial input feature of GAT, including the maximum speed limit, average travel time, segment direction, out-degree, in-degree, segment length, and segment type. To be specific, for the maximum speed limit, average travel time, segment direction, and segment length, we apply the min-max normalization. Regarding segment type, we categorize each segment into eight classes, i.e., [*living street, motorway, primary, residential, secondary, tertiary, trunk, unclassified*], and we employ one-hot encoding. All these attributes are concatenated to obtain the input features f_i for segment v_i . Formally, given the initial feature f_i of segment v_i and road network graph \mathcal{G} , GAT computes segment embedding as $\mathbf{h}_i = \text{GAT}(f_i, \mathcal{G})$. We use a 3-layer GAT model to consider 3-hop neighbors for each segment.

Time Encoding \mathbf{t}_i . We consider two types of time regularities for trajectories. (i) Trajectory patterns vary considerably at different times of a day or a week. For instance, many trajectories move towards office areas on weekday mornings, while evenings usually observe a surge in trajectories heading back home. (ii) The traffic patterns of different road segment types differ at the same time. For example, segments in commercial areas show elevated trajectory volumes during the daytime, whereas segments in residential areas observe increased trajectory data during non-work hours.

To account for (i), we introduce a time encoding to capture periodicity. Inspired by Time2vec [21], given the timestamp t_i of segment v_i in a trajectory, we first transform t_i into a 6-dimension vector [*hour, minute, second, year, month, day*], then we learn timestamp t_i as a vector \mathbf{e}_i^t of dimension $l/2$ as follows:

$$\mathbf{e}_i^t = \text{FC}_1(t_i) \parallel \sin(\text{FC}_2(t_i)), \quad (4)$$

where $\text{FC}_1(\cdot)$ and $\text{FC}_2(\cdot)$ are linear layers to learn time embedding, and $\sin(\cdot)$ function helps capture time periodic behaviors. \parallel is a concatenate operation for vectors by channel-dimension. To account for (ii), we design a segment type encoding. In particular, we employ a learnable type encoding matrix $\mathbf{E}^p \in \mathbb{R}^{|T| \times \frac{l}{2}}$, where each row corresponds to a segment type in $T = \{\textit{living street, motorway, primary, residential, secondary, tertiary, trunk, unclassified}\}$. Given the type of a segment v_i , we look up \mathbf{E}^p to obtain its type encoding \mathbf{e}_i^p . Then we concatenate the time encoding \mathbf{e}_i^t and segment type encoding \mathbf{e}_i^p , and use a fully connected network to interact two encodings $\mathbf{t}_i = \text{FC}(\mathbf{e}_i^t \parallel \mathbf{e}_i^p) \in \mathbb{R}^d$.

User Encoding \mathbf{u}_i . Different drivers have different travel patterns in their trajectories. E.g., an office worker may follow the same

route on all weekdays, while taxi drivers try to pick up passengers in crowded areas every day. By combining spatial information and temporal information (joint embedding) with trajectory attributes (dual-objective task learning), the user encoding can capture the specific behaviors of individual users, including travel time, trajectory length, and high-frequency access areas, which benefits tasks such as trajectory similarity computation, travel time estimation, and trajectory classification. To encode user information, we employ a learnable user encoding matrix $\mathbf{E}^u \in \mathbb{R}^{|U| \times l}$, where $|U|$ is the number of users, and each row of \mathbf{E}^u corresponds to a user.

4.3 Dual-objective Task Learning

Here, we present the two learning objectives of RED.

Path Encoder. The path encoder is designed to capture key information of a path trajectory and to predict the next segment for learning local information of the trajectory. Section 4.1 shows that the key paths of a trajectory carry much more information than the mask paths, suggesting that a trajectory can be described by the key paths. Thus, we require the path encoder to learn information from the key paths. Formally, given the key path set of a trajectory $\mathcal{T}^k = \langle v_1^k, v_2^k, \dots, v_{|\mathcal{T}^k|}^k \rangle$, the corresponding key path embedding $\mathbf{X}^k = \langle \mathbf{x}_1^k, \mathbf{x}_2^k, \dots, \mathbf{x}_{|\mathcal{T}^k|}^k \rangle$ is first generated by the spatial-temporal-user joint embedding, then position encoding [34] is added to \mathbf{X}^k to introduce the position information of each segment in the path trajectory. We feed \mathbf{X}^k to the Transformer of the encoder and denote $\widehat{\mathbf{X}}^k = \text{Transformer}(\mathbf{X}^k)$ as the output.

To extract the prediction for the next segment from encoder output, we use a fully connected layer $\mathbf{Y}^k = \text{FC}(\widehat{\mathbf{X}}^k)$, where $\mathbf{Y}^k \in \mathbb{R}^{|\mathcal{T}^k| \times |\mathcal{V}|}$. Each row of \mathbf{Y}^k is the predicted distribution of a key segment, and each column of \mathbf{Y}^k corresponds to one segment (with $|\mathcal{V}|$ being the number of all segments). Then we use cross-entropy loss to compute errors between the predicted value and the ground-truth segment ID for the next segment prediction task as follows:

$$\mathcal{L}_{\mathcal{T}^k}^{nsp} = -\frac{1}{|\mathcal{T}^k|} \sum_{v_i \in \mathcal{T}^k} \log \frac{\exp(\mathbf{Y}_{v_i}^k)}{\sum_{v_j \in \mathcal{V}} \exp(\mathbf{Y}_{v_j}^k)}, \quad (5)$$

We average the above loss over the trajectories in a mini-batch to obtain the final next segment prediction loss \mathcal{L}^{nsp} .

Note that we predict the next key path in the encoder, so we use *causal self-attention* here to avoid information leakage.

Path Decoder. The path decoder is designed to restore a raw path trajectory by decoding the embeddings of its segments generated by the encoder. However, the path encoder only encodes the key paths and ignores the mask paths. Although the mask paths carry less information about the trajectory, they are indispensable when restoring a complete trajectory. Thus, we use another Transformer model as the decoder to combine the key paths and the mask paths. We first transform all mask path segments to learnable embedding \mathbf{X}^m , and combine \mathbf{X}^m with the outputs of the path encoder for the key paths. Then we conduct an *unshuffle* operation to restore the mask path embeddings to their original positions in path trajectory \mathcal{T} . This process can be formulated as $\mathbf{X}^T = \text{Unshuffle}(\text{Combine}(\widehat{\mathbf{X}}^k, \mathbf{X}^m))$.

The decoder takes \mathbf{X}^T as input and computes hidden embedding output as $\widehat{\mathbf{X}}^T = \text{Transformer}(\mathbf{X}^T)$. To restore the raw information of trajectory \mathcal{T} , we design a trajectory reconstruction task that

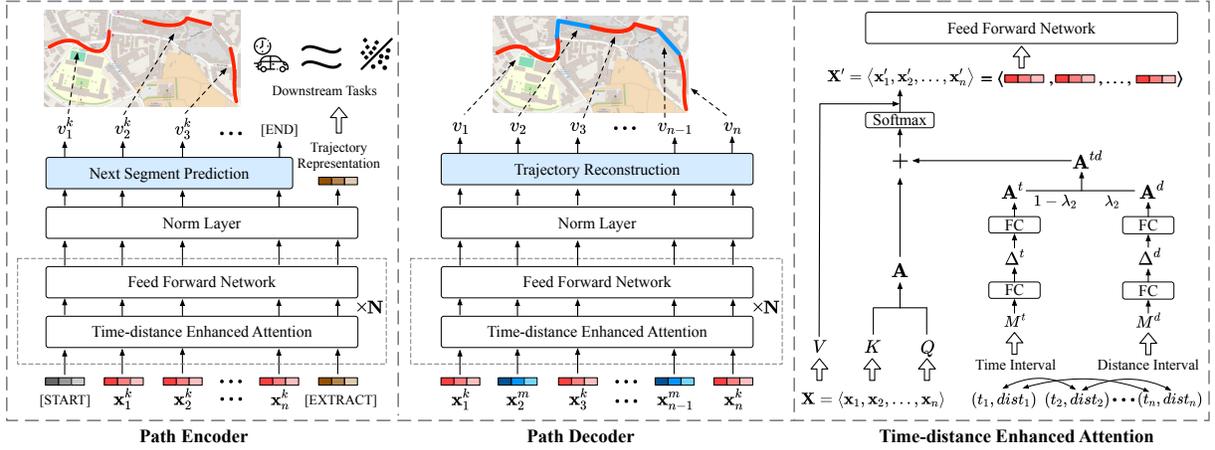


Figure 4: The enhanced Transformer of the encoder and decoder.

aims to reconstruct the whole path trajectory from the decoder output $\widehat{\mathcal{T}}$. To achieve this, we first use a fully connected layer to obtain the predictions $Y^{\mathcal{T}} = FC(\widehat{\mathcal{T}})$, where $Y^{\mathcal{T}} \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{V}|}$. Each row of $Y^{\mathcal{T}}$ is the predicted distribution of a segment of \mathcal{T} , and each column of $Y^{\mathcal{T}}$ corresponds to a segment. Then we use cross-entropy loss for a complete trajectory \mathcal{T} as follows:

$$\mathcal{L}_{\mathcal{T}}^{tr} = -\frac{1}{|\mathcal{T}|} \sum_{v_i \in \mathcal{T}} \log \frac{\exp(Y_{v_i}^{\mathcal{T}})}{\sum_{v_j \in \mathcal{V}} \exp(Y_{v_j}^{\mathcal{T}})}, \quad (6)$$

We also average the above loss over all trajectories in a mini-batch to obtain the final trajectory reconstruction loss \mathcal{L}^{tr} .

Existing TRL methods consider limited supervised information in the training task, e.g., START [20] only reconstructs mask segments of a path trajectory, and Traj2vec [45] recovers trajectory points. RED with dual-objective task explores more supervised information and thus can improve the accuracy of trajectory representations.

4.4 Enhancing Transformer for Trajectory

To use Transformer as the model for the path encoder and decoder, two issues need to be addressed. First, the next segment prediction task is different from next token prediction of language in subtle ways, and directly using Transformer will encounter errors in the input and output (which we call segment misalignment). Second, the Transformer relies on the self-attention among the token embeddings while trajectories contain spatial-temporal information. To utilize such spatial-temporal information, we need to integrate it into Transformer attention. To fix segment misalignment, we introduce *virtual tokens* for the encoder. To utilize spatial-temporal information of trajectories, we design a *time-distance enhanced attention* to enhance the Transformer for the decoder. Figure 4 gives an overview of the enhanced path encoder and decoder.

4.4.1 Virtual Token. The next segment prediction task of path encoder can encounter segment misalignment. For instance, given path trajectory $\langle v_1, v_2, v_3, v_4, v_5 \rangle$, the key path is $\langle v_1, v_3, v_5 \rangle$ and the mask path is $\langle v_2, v_4 \rangle$. To predict the next key segments, the input of the path encoder is $\langle v_1, v_3 \rangle$, and the expected output of the path

encoder is $\langle v_3, v_5 \rangle$. Therefore, the last key segment v_5 is not used as input and the first key segment v_1 is not used as output. When we combine the output of the encoder with the mask path $\langle v_2, v_4 \rangle$, we will get an incomplete trajectory input $\langle v_2, v_3, v_4, v_5 \rangle$ (i.e., missing v_1) for the decoder, which causes incomplete trajectory reconstruction. To solve these problems, we introduce three virtual tokens, i.e., a start token, an end token, and an extract token.

Start Token. To complete the first segment in the path encoder output, we introduce a start token [START] at the initiation of all trajectories. However, utilizing [START] to predict the first key segment of a path trajectory cannot consider the topology of the road network. To address this problem, we add [START] as a virtual node in the road network graph \mathcal{G} . Specifically, the [START] node is connected to all the nodes of \mathcal{G} with edges, and this augmented graph is also used to generate the spatial encoding.

End Token. The end token [END] serves as the expected output of the path encoder when taking all key segments of a trajectory as input to predict the next segment.

Extract Token. To extract the overall information of a trajectory, we introduce an extract token [EXTRACT] at the end of each trajectory. The encoder output embedding for [EXTRACT] is used as the trajectory representation.

To illustrate the virtual tokens, we take the same path trajectory $\langle v_1, v_2, v_3, v_4, v_5 \rangle$ as an example. Still, the key path is $\langle v_1, v_3, v_5 \rangle$ and the mask path is $\langle v_2, v_4 \rangle$. By adding virtual tokens, the input of the path encoder becomes $\langle [\text{START}], v_1, v_3, v_5, [\text{EXTRACT}] \rangle$. The output of the path encoder is to predict $\langle v_1, v_3, v_5, [\text{END}] \rangle$, and the output of [EXTRACT] is the trajectory representation. Therefore, the input and output of the path encoder is complete.

4.4.2 Time-distance Enhanced Attention. With spatial-temporal-user joint embedding, a path trajectory $\mathcal{T} = \langle \tau_1, \tau_2, \dots, \tau_{|\mathcal{T}|} \rangle$ is transformed to embedding $\mathbf{X} = \langle \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{|\mathcal{T}|} \rangle \in \mathbb{R}^{|\mathcal{T}| \times l}$. Then we can use the vanilla Transformer to model trajectory representation. However, directly applying Equation 1 neglects the essential spatial-temporal information in trajectories. Spatially, a path trajectory consists of continuous road segments, and each segment is related most closely to its adjacent segments. Consequently, two segments

that are distant in a path trajectory should have small attention scores, while two segments that are close should have large attention scores. Temporally, the sampling timestamps of the segments of a trajectory are consecutive. Two segments with a large sampling time difference should have small attention scores, while the converse is true for two segments with a small sampling time difference. To ensure these properties, we introduce the time-distance enhanced attention to consider the time interval and distance interval between the road segments in a trajectory. In particular, we construct a time-distance attention matrix \mathbf{A}^{td} to incorporate spatial-temporal information to the vanilla Transformer as follows:

$$\mathbf{X}' = \text{softmax}(\mathbf{A} + \mathbf{A}^{td})\mathbf{V}, \quad (7)$$

where $\mathbf{A}^{td} = (1 - \lambda_2)\mathbf{A}^t + \lambda_2\mathbf{A}^d$. \mathbf{A}^t is the time interval correlation, \mathbf{A}^d is the distance interval correlation, and λ_2 is a hyper parameter. **Time Interval Correlation.** Follow by START [20], consider a path trajectory $\mathcal{T} = \langle \tau_1, \tau_2, \dots, \tau_{|\mathcal{T}|} \rangle$, where $\tau_i = (v_i, t_i)$ denotes the segment v_i and timestamp t_i that the trajectory passes through. We first construct a time interval matrix $M^t \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$, where element $m_{i,j}^t \in M^t$ is the difference $|t_i - t_j|$ between the sampling timestamps of segments v_i and v_j . Then, we introduce a transformation $f(m_{i,j}^t) = \frac{1}{\log(e+g_t(m_{i,j}^t))}$ to make larger values smaller to

capture correlation based on time interval, where $g_t(m_{i,j}^t) = \frac{m_{i,j}^t}{60}$ maps seconds to minutes to avoid excessive time differences between the segments. In the path encoder, the inputs are key paths of a trajectory, so the time interval matrix is a sub-matrix of M^t .

To incorporate the information of M^t into the trajectory representation, we use a fully connected layers to transform each element $m_{i,j}^t \in M^t$ to a vector in the embedding space as $\delta_{i,j}^t = FC(m_{i,j}^t)$, where the $\delta_{i,j}^t \in \mathbb{R}^{\frac{l}{2}}$. Consequently, we obtain the time interval hidden embedding $\Delta^t \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}| \times \frac{l}{2}}$ for \mathcal{T} . Then, we use another fully connected layer to transform Δ^t to time interval correlation matrix $\mathbf{A}^t \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$, where each element $\alpha_{i,j}^t \in \mathbf{A}^t$ is a time interval correlation value, and $\alpha_{i,j}^t = FC(\delta_{i,j}^t)$.

Distance Interval Correlation. Like time interval correlation, we first construct a travel distance interval matrix $M^d \in \mathbb{R}^{|\mathcal{T}| \times |\mathcal{T}|}$ for a path trajectory \mathcal{T} . In particular, we use the travel length between two segments as $m_{i,j}^d \in M^d$. Then another transformation $f(m_{i,j}^d) = \frac{1}{\log(e+g_d(m_{i,j}^d))}$ is used to ensure that two distant segments have

lower correlation, where $g_d(m_{i,j}^d) = \frac{m_{i,j}^d}{1000}$ transforms distance from meters to kilometers to avoid excessive values. We also use a fully connected layer to map the transformed distances to the final distance interval correlation matrix \mathbf{A}^d .

4.5 Discussion

Complexity Analysis. The training time of RED encompasses three terms: i.e., road-aware masking, graph neural network, and trajectory encoding. The road-aware masking has time complexity $\mathcal{O}(|\mathcal{D}| \cdot |\mathcal{T}|)$, where $|\mathcal{T}|$ is the average trajectory length and $|\mathcal{D}|$ is the cardinality of the trajectory dataset. The GNN training has complexity $\mathcal{O}(|\mathcal{V}| \cdot l + |\mathcal{E}| \cdot l)$, where l is the embedding dimensionality, $|\mathcal{V}|$ and $|\mathcal{E}|$ are the numbers of nodes and edges in the road network. The complexity of encoding the trajectories is

Table 2: Statistics of the experiment datasets.

Dataset	Porto	Rome	Chengdu
#Users	439	313	13,715
#Path trajectories	859,986	80,203	4,413,602
#Road Segments	10,537	49,230	45,534
Avg. trajectory length	37	56	26

$\mathcal{O}(|\mathcal{D}| \cdot |\mathcal{T}^k|^2 \cdot l \cdot L_e + |\mathcal{D}| \cdot |\mathcal{T}|^2 \cdot l \cdot L_d)$, where $|\mathcal{T}^k|$ is the number of key segments in trajectory \mathcal{T} and $|\mathcal{T}^k| < |\mathcal{T}|$, L_e and L_d are the numbers of path encoder and path decoder layers. To conduct model inference for a trajectory dataset \mathcal{D}' , we use the path encoder to compute the vector representations. The complexity of this is $\mathcal{O}(|\mathcal{D}'| |\mathcal{T}|^2 \cdot l \cdot L_e)$. The resulting trajectory representations take $\mathcal{O}(l \cdot |\mathcal{D}'|)$ space. Using the representations, computing the similarity of two trajectories has time cost $\mathcal{O}(l)$. In contrast, traditional trajectory similarity algorithms (e.g., DTW, EDR) have cost $\mathcal{O}(|\mathcal{D}| \cdot |\mathcal{T}|^2)$ because they use dynamic programming.

Limitations. RED targets trajectories on the road network (i.e., involving map matching and road segments) and needs to be adjusted for other types of trajectories (e.g., POI trajectories, animal and human flow trajectories). Moreover, RED focuses on trajectory-based tasks (e.g., similarity computation and travel time estimation) and is unsuitable for road network related tasks that do not work at the trajectory level, e.g., road classification and flow estimation.

5 EXPERIMENTAL EVALUATION

We evaluate and compare RED with the state-of-the-art baselines on four downstream tasks, trajectory classification, travel time estimation, trajectory similarity computation, and most similar trajectory retrieval on three real-world datasets.

5.1 Experimental Settings

Datasets. We use three real-world datasets: Porto¹, Rome², and Chengdu³. Key statistics are given in Table 2. We use map data from OpenStreetMap⁴. This data includes road IDs, lengths, road types, speed limits, etc. Trajectory points include latitude, longitude, timestamp, user ID, etc., which are sampled on average every 15, 5, and 30 seconds in Porto, Rome, and Chengdu, respectively. We remove trajectory points that are not within the chosen map boundaries and that are drifting. We apply the FMM [39] map matching method, which is widely used for road network-based methods [20, 31] and is efficient and accurate, to obtain path trajectories. We also remove path trajectories with fewer than 6 segments and set the maximum path trajectory length to 256. The overall proportion of training, validation, and testing data for Porto and Chengdu is set to [0.6, 0.2, 0.2], while for Rome, it is set to [0.8, 0.1, 0.1].

Baselines. We compare RED with 9 state-of-the-art TRL methods, including the four *RNN-based methods* **Traj2vec** [45], **T2vec** [23], **Trembr** [13], and **PIM** [41] and the five *Transformer-based methods* **Transformer** [34], **BERT** [11], **PIM-TF** [41], **Toast** [10], and

¹<https://www.kaggle.com/c/pkdd-15-predict-taxi-service-trajectory-i>

²<https://ieee-dataport.org/open-access/crawdad-Rometaxi>

³<https://www.pkbigdata.com/common/zhzgbCmptDetails.html>

⁴<https://www.openstreetmap.org>

Table 3: Accuracy at travel time estimation. Best-performing existing methods are underlined, and the bottom row reports RED’s accuracy improvement over the best existing method.

	Porto			Rome			Chengdu		
	MAE↓	MAPE(%)↓	RMSE↓	MAE↓	MAPE(%)↓	RMSE↓	MAE↓	MAPE(%)↓	RMSE↓
Traj2vec	1.624	20.005	3.137	5.933	45.292	7.188	1.397	16.833	2.114
T2vec	1.608	19.906	3.043	5.873	44.975	7.104	1.400	17.072	2.136
Trembr	1.576	19.391	2.892	5.336	44.220	6.907	1.374	16.635	1.937
PIM	1.684	20.137	3.205	5.992	45.937	7.294	1.622	20.897	2.796
Transformer	1.844	21.325	3.790	6.684	47.131	9.982	1.707	22.458	3.031
BERT	1.707	20.243	3.275	6.114	46.142	7.358	1.453	17.520	2.544
PIM-TF	2.018	22.456	3.898	6.941	48.020	8.313	1.833	23.561	3.163
Toast	1.801	20.952	3.334	6.411	46.864	7.997	1.391	16.984	2.110
START	<u>1.506</u>	<u>18.012</u>	<u>2.799</u>	<u>4.342</u>	<u>43.823</u>	<u>6.282</u>	<u>1.278</u>	<u>15.961</u>	<u>1.889</u>
RED	1.468	16.716	2.651	3.798	41.658	5.358	1.212	15.648	1.720
Improve.	2.52%	7.20%	5.29%	12.53%	4.94%	14.71%	5.16%	1.96%	8.95%

Table 4: Accuracy at trajectory classification. Best-performing existing methods are underlined, and the bottom row reports RED’s accuracy improvement over the best existing method.

	Porto			Rome			Chengdu		
	Mi-F1↑	Ma-F1↑	Recall@5↑	Mi-F1↑	Ma-F1↑	Recall@5↑	Accuracy ↑	Precision ↑	F1↑
Traj2vec	0.074	0.052	0.223	0.069	0.047	0.185	0.809	0.817	0.820
T2vec	0.083	0.062	0.238	0.074	0.053	0.190	0.794	0.803	0.814
Trembr	0.087	0.069	0.243	0.081	0.062	0.208	0.811	0.836	0.848
PIM	0.069	0.041	0.207	0.060	0.034	0.174	0.764	0.787	0.791
Transformer	0.043	0.035	0.107	0.041	0.034	0.098	0.767	0.794	0.800
BERT	0.077	0.054	0.229	0.075	0.058	0.177	0.810	0.837	0.852
PIM-TF	0.035	0.028	0.083	0.032	0.026	0.079	0.754	0.762	0.786
Toast	0.075	0.048	0.215	0.070	0.044	0.206	0.824	0.841	0.860
START	<u>0.098</u>	<u>0.076</u>	<u>0.251</u>	<u>0.095</u>	<u>0.079</u>	<u>0.238</u>	<u>0.825</u>	<u>0.849</u>	<u>0.879</u>
RED	0.107	0.085	0.262	0.120	0.106	0.284	0.835	0.865	0.885
Improve.	9.18%	11.84%	4.38%	26.32%	34.18%	19.33%	1.21%	1.88%	0.68%

START [20]. For trajectory similarity tasks, we further introduce the seven heuristic methods **SSPD** [4], **Hausdorff** [1], **Fréchet** [2], **DTW** [3], **LCSS** [36], **ERP** [8], and **EDR** [8].

Implementation. We implement RED using PyTorch 2.0. We use a machine running Ubuntu 20.04 and with an NVIDIA RTX GeForce 4090 GPU. We set the embedding size l to 128 in all methods, the number of GAT layers to 3, and the number of encoder and decoder layers to 6. The attention heads are [8, 16, 1] for GAT, and 8 for the encoder and decoder. The dropout ratio is 0.1. We set $\lambda_1 = 0.1$ for the dual-objective tasks, and $\lambda_2 = 0.5$ for time-distance correlation. We pre-train RED using the AdamW [30] optimizer. The batch sizes are 64, 32, and 32, the training epochs are 10, 30, and 5, and the learning rates are $1e-4$ for Porto and Chengdu, $2e-4$ for Rome.

Downstream Task Settings. For all downstream tasks, we only use the encoder of RED and use complete trajectories as input. For trajectory classification and travel-time estimation, we use AdamW [30] to fine-tune the encoder. The batch size is 64, the training epoch is 30, and the learning rate is $1e-4$. For trajectory classification on Porto and Rome, we classify trajectories according to user IDs and remove the user encoding to avoid information leakage during fine-tuning. For Chengdu, we classify trajectories based on whether the vehicle is carrying passengers and keep the user encoding during fine-tuning. We add an extra linear layer to

get predictions and use cross-entropy as the loss function. For travel time estimation, we remove all time information, including the time encoding as well as the time correlation, except for the departure time, to avoid time information leakage. We also add an extra linear layer to predict travel time and use mean square error (MSE) as the loss function. The travel times of all methods are reported in minutes. For trajectory similarity computation, we conduct two types of experiments without fine-tuning, including trajectory similarity computation and most similar trajectory retrieval, and we use the inner-product of trajectory vectors as the similarity score.

Performance Metrics. For travel time estimation, we report the mean absolute error (MAE), mean absolute percentage error (MAPE), and root mean square error (RMSE). For trajectory multi-classification, we report the Micro-F1 (Mi-F1), Macro-F1 (Ma-F1), and Recall@5 values. For the trajectory binary-classification, we report the F1-score, accuracy, and precision. For trajectory similarity tasks, we use Hit Ratio i.e., the HR@1, HR@5, HR@10 values, for trajectory similarity computation and the mean rank (MR) for most similar trajectory retrieval. Ideally, the MR value should be equal to 1.

5.2 Main Results

The main findings for the four downstream tasks are reported at Table 3 in travel time estimation, Table 4 for trajectory classification,

Table 5: Accuracy at trajectory similarity computation. The best-performing baseline is marked with underline, and the bottom row is our accuracy improvement over the best baseline.

	Porto						Rome					
	Hausdorff distance			Fréchet distance			Hausdorff distance			Fréchet distance		
	HR@1↑	HR@5↑	HR@10↑	HR@1↑	HR@5↑	HR@10↑	HR@1↑	HR@5↑	HR@10↑	HR@1↑	HR@5↑	HR@10↑
Traj2vec	0.103	0.292	0.380	0.110	0.308	0.404	0.092	0.276	0.329	0.101	0.285	0.340
T2vec	0.129	0.330	0.431	0.140	0.352	0.458	0.121	0.304	0.361	0.135	0.307	0.356
Trembr	0.141	0.354	0.462	0.164	0.377	0.476	0.134	0.311	0.379	0.147	0.320	0.384
PIM	0.114	0.302	0.374	0.129	0.321	0.399	0.105	0.279	0.343	0.111	0.286	0.350
Transformer	0.083	0.260	0.351	0.090	0.271	0.372	0.079	0.244	0.302	0.088	0.266	0.321
BERT	0.091	0.277	0.364	0.099	0.290	0.385	0.085	0.261	0.318	0.092	0.273	0.329
PIM-TF	0.073	0.241	0.337	0.075	0.248	0.343	0.064	0.210	0.278	0.072	0.221	0.290
Toast	0.121	0.315	0.408	0.137	0.320	0.417	0.113	0.291	0.357	0.128	0.302	0.366
START	0.172	0.389	0.490	0.185	0.427	0.532	0.162	0.329	0.401	0.179	0.339	0.418
RED	0.239	0.501	0.608	0.246	0.523	0.638	0.190	0.358	0.431	0.205	0.387	0.463
Improve.	38.95%	28.79%	24.08%	32.97%	22.48%	19.92%	17.28%	8.81%	7.48%	14.53%	14.16%	10.77%

Table 6: Accuracy for most similar trajectory retrieval in terms of mean rank (MR↓), where lower values indicate better performance. The best-performing method is highlighted in bold, and the second-best method is underlined.

	Porto						Rome					
	$p = 0.1$ on diverse dataset			various p on 100k dataset			$p = 0.1$ on diverse dataset			various p on 40k dataset		
	10k	50k	100k	$p = 0.2$	$p = 0.3$	$p = 0.4$	10k	20k	40k	$p = 0.2$	$p = 0.3$	$p = 0.4$
Hausdorff	35.42	83.12	132.2	237.3	312.0	464.1	50.18	68.19	103.6	236.8	262.4	340.5
Fréchet	45.10	88.95	159.1	210.5	372.1	469.6	187.6	212.6	437.8	55.55	73.62	88.70
DTW	1.752	3.067	3.694	6.298	13.01	30.60	1.985	2.946	4.832	10.33	24.64	59.43
LCSS	92.67	312.8	421.6	965.9	1349	2431	125.1	147.2	350.2	706.7	1143	1645
EDR	146.2	447.2	667.3	1090	1508	2202	176.8	232.1	455.5	852.3	1421	1830
ERP	32.13	89.15	274.1	393.0	424.5	572.6	86.32	107.7	331.1	446.7	444.0	520.8
SSPD	2.034	3.816	4.754	6.855	14.72	32.27	3.122	5.078	8.948	25.33	36.93	102.1
Traj2vec	1.842	2.693	2.911	5.123	10.72	18.36	1.801	2.034	2.739	4.911	8.234	14.34
T2vec	1.740	2.351	2.400	5.940	10.34	17.00	1.661	1.839	2.177	4.267	7.963	15.72
Trembr	1.509	2.010	2.130	4.536	9.681	12.74	1.488	1.721	1.990	3.891	6.706	10.44
PIM	4.331	8.937	10.98	40.66	67.50	161.1	3.912	4.771	8.324	36.42	63.00	92.64
Toast	4.997	9.310	12.43	44.21	84.45	132.6	3.949	4.801	7.330	29.65	48.51	78.50
START	1.232	1.720	1.847	3.251	6.241	8.831	<u>1.114</u>	<u>1.222</u>	<u>1.445</u>	<u>3.524</u>	<u>5.583</u>	7.268
RED	<u>1.420</u>	<u>1.893</u>	<u>1.996</u>	<u>4.158</u>	<u>7.497</u>	<u>10.77</u>	1.096	1.171	1.343	2.955	4.370	<u>7.560</u>

Table 5 for trajectory similarity computation, and Table 6 for most similar trajectory retrieval. For brevity, we only show results on two datasets for some experiments.

Performance at Travel Time Estimation. The results for travel time estimation in Table 3 show that RED achieves the best performance on all metrics on three datasets. Among the existing methods, Trembr and START perform much better than other methods because these two methods exploit temporal information. On the contrary, disregarding temporal information yields inaccurate predictions, e.g., Toast and PIM. Our model not only utilizes temporal correlations in trajectories but also includes a pre-trained time embedding, i.e., Time2vec [21], which makes the input of the model carry more temporal features. In addition, we consider the behaviors of the road segment types in different time periods so that the features of road segments and the time are fully integrated.

Performance at Trajectory Classification. The results for trajectory classification are shown in Table 4. RED performs the best because of introducing more comprehensive trajectory information, e.g., spatial information, temporal information, user encoding, and etc. Moreover, RED’s road-aware masking strategy preserves the

crucial information of trajectories, while its graph neural networks and use of static attributes of related road segments as input. The use of comprehensive information can improve the accuracy of trajectory representation and greatly enhance the performance of trajectory classification. On the Chengdu dataset, the performance improvement of RED is small because the binary classification task is simple, making it difficult for any method to excel.

Performance at Trajectory Similarity Computation. We choose two commonly used trajectory similarity heuristics, i.e., Hausdorff distance and Fréchet distance, to obtain the most similar trajectory ground truths in the test set. The results of Table 5 show that RED achieves the best performance. RED utilizes comprehensive information on trajectories to improve the expressive of learned trajectory representations and thus supports the high accuracy of trajectory similarity computation.

Performance at Most Similar Trajectory Retrieval. Following the method [6], we create experimental data including a query set Q and a database set D from datasets. Then we downsample the points of each trajectory (if the method uses a road network, we apply detour method [31]), in the query set Q at ratio p to obtain

Table 7: Accuracy for trajectory classification on the Geolife dataset.

Method	SVM [18]	RF [5]	TimeLSTM [52]	GRU-D [7]	STGN [48]	STGRU [28]	TrajODE [25]	Trajformer [24]	RED
Accuracy↑	49.88	56.16	66.68	71.71	75.60	73.15	85.25	<u>85.45</u>	86.89

Table 8: Average time (in μ s) to compute the representation for a trajectory.

Method	RNN-based Method				Transformer-based Method					RED
	Traj2vec	T2vec	Trembr	PIM	Transformer	BERT	PIM-TF	Toast	START	
Time	302.25	291.90	297.33	346.00	137.80	140.83	151.93	205.92	217.80	213.30
Rank	9	7	8	10	1	2	3	4	6	5

Table 9: Average time (in μ s) to compute one trajectory similarity of RED and different heuristic algorithms.

	Hausdorff	Fréchet	DTW	LCSS	EDR	ERP	SSPD	RED
Porto	58.48	215.55	66.87	70.68	64.62	216.52	60.30	6.14
Rome	229.85	1198.52	367.13	389.87	353.49	1194.93	232.08	

a sub-trajectory set Q' , and then we add Q' to D to obtain the expanded database set D' . By construction the trajectories in Q and the trajectories in Q' have high similarity. For each trajectory in query set Q , the task aims to find the corresponding sub-trajectory in D' . We set $|Q| = 1,000$, and then set $|D| = 100,000$ in Porto and $|D| = 40,000$ in Rome. Table 6 shows the results.

Varying the database size $|D|$. We fix the downsampling rate p at 0.1, and vary the database size $|D|$ from 10,000 to 100,000 in Porto and from 10,000 to 40,000 in Rome. The results show that the average performance of the TRL methods is much higher than that of the heuristic algorithms. Because a sub-trajectory has the same semantics as the raw trajectory and TRL methods rely on trajectory semantics, they achieve better performance. Second, the performance of RED is slightly lower than START on Porto and slightly higher than that of START on Rome. The average length of a trajectory in Rome is longer than that of Porto (see Tabel 2), and through the learning of key paths, our model can capture semantic information of trajectories well in the case of long trajectories.

Varying downsampling rate p . We downsample trajectories in $|Q|$ by different ratios $p \in [0.2, 0.4]$, while fixing $|D|$ at 100,000 for Porto and 40,000 for Rome. The results show that the average performance of the TRL methods is much higher than that of the heuristic algorithms. The heuristic algorithms perform very poorly at the high downsampling rate $p = 0.4$, while the TRL methods still perform well using their learned semantic information. RED performs slightly below START on Porto and slightly above START on Rome. Because START uses trimming data augmentation during training, it relies on different scales to trim raw trajectories, and this makes START insensitive to downsampling rates.

Generalization of RED. In addition to the vehicle trajectories, we also apply RED to human mobility trajectories using the Geolife [50] dataset. When performing trajectory classification [24, 25], Table 7 shows that RED achieves the best accuracy. We further apply RED to sparse trajectories, i.e., POI trajectories, for next POI

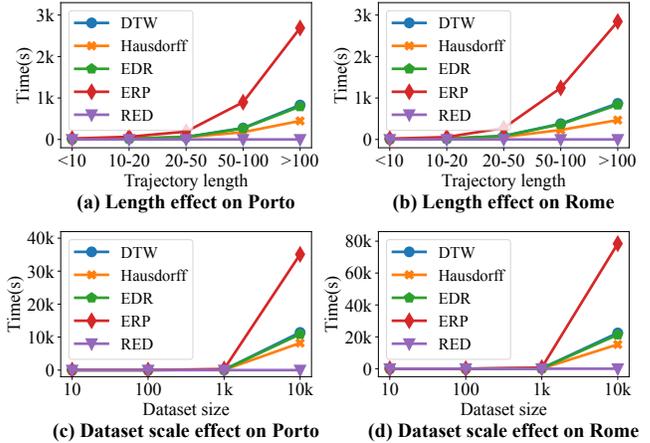


Figure 5: Trajectory similarity computation time (in s) when varying the trajectory length and dataset size.

recommendation [37, 40] and route recommendation [14, 22]. Here, RED achieves the second-best or the third-best accuracy.

5.3 Efficiency Study

We study the efficiency on Porto and Rome. In Table 8, we report the average inference time for a trajectory. The inference time is crucial for the efficiency of downstream tasks. The results show that the Transformer-based methods have the shorter inference time than the RNN-based methods. This is because the self-attention in the Transformer can be computed in parallel, while RNNs need to iterate over the paths in a trajectory. The inference time of RED is comparable to those of the Transformer-based methods, which usually perform the best in the accuracy experiments. This suggests that good accuracy comes from tailored model designs instead of high model complexity.

Computing the similarity between two trajectories is the basic operation in retrieval and clustering. In Table 9, we report the average time of a trajectory similarity computation. We do not compare RED with other TRL methods because they all transform similarity computation into vector computation and thus have identical computation time when using the same embedding dimensionality. Instead, we compare with heuristic trajectory similarity measures such as DTW, LCSS, etc. The results show that RED is much faster

Table 10: Accuracy of RED when disabling key designs.

	Porto			Rome		
	MAPE↓ (TTE)	HR@10↑ (TS)	Recall@5↑ (TC)	MAPE↓ (TTE)	HR@10↑ (TS)	Recall@5↑ (TC)
w/o NSP	18.239	0.466	0.259	42.029	0.277	0.260
w/o TE	18.681	0.385	0.252	44.835	0.336	0.274
w/o VT	18.537	0.394	0.257	43.933	0.301	0.264
w/o TD	17.587	0.539	0.261	41.904	0.353	0.279
RED	16.716	0.638	0.262	41.658	0.463	0.284

than the heuristic measures. This is because the vector computation of RED has linear complexity while the heuristic measures use dynamic programming, which has quadratic complexity.

In addition, we explore the impact of trajectory length and dataset size on the efficiency of similarity computation. Figure 5 reports the total time to compute the similarity of trajectory pairs in two datasets. For the trajectory length experiments, we partition trajectories into five sets based on their lengths: [0~10, 10~20, 20~50, 50~100, 100~]. Within each set, we randomly sampled 1,000 trajectories and compute the similarity between each pair. Figures 5(a) and (b) show that the heuristic algorithms exhibit runtimes that increase notably with longer trajectories. Conversely, RED employs parallel computation on vectors, substantially reducing the computational time. For the dataset size experiments, we vary the dataset size across [10, 100, 1k, 10k] and compute the similarity of each trajectory pair. Figures 5(c) and (d) show again that the heuristic algorithms are inefficient. In contrast, parallel computation on vectors is highly efficient and is insensitive to data size fluctuations.

5.4 Ablation Study and Model Designs

Ablation Study. To gain insight into the contributions of each sub-module in RED, we conduct ablation study on Porto and Rome. Table 10 shows the results. We denote travel time estimation as TTE, trajectory classification as TC, and trajectory similarity as TS. For brevity, we only report on one metric for each task. In trajectory similarity, we show the performance using Fréchet distance.

- **w/o NSP:** We remove the next segment prediction task, and keep only the trajectory reconstruction task.
- **w/o time encoding (TE):** We remove the time encoding and keep only the segment types.
- **w/o virtual token (VT):** We remove all virtual tokens and the virtual node, and take the prediction of the last position of the encoder as trajectory representation.
- **w/o time-distance (TD):** We remove the time-distance correlation in self-attention in the encoder and the decoder.

Table 10 indicates a pronounced impact at HR@10 performance in trajectory similarity computation. This metric relies on pre-trained trajectory vectors without fine-tuning, necessitating robust generalization of representations. Although reductions in MAPE and Recall@5 performance are also observed, they are comparatively less significant than that of HR@10.

Additionally, the time encoding, NSP, and virtual token mechanisms play pivotal roles in RED. Time encoding captures crucial temporal information, notably enhancing trajectory representation. NSP, tasked with predicting the next segment, is a challenging yet pivotal component. The virtual token addresses data misalignment

Table 11: Accuracy of RED with tuned random mask ratio and road-aware masking strategy.

	Porto			Rome		
	MAPE↓ (TTE)	HR@10↑ (TS)	Recall@5↑ (TC)	MAPE↓ (TTE)	HR@10↑ (TS)	Recall@5↑ (TC)
Ratio 0.1	17.260	0.591	0.259	43.307	0.435	0.274
Ratio 0.3	17.258	0.602	0.262	43.107	0.452	0.282
Ratio 0.5	17.267	0.535	0.262	<u>42.331</u>	0.471	0.281
Ratio 0.7	17.290	0.502	<u>0.260</u>	42.972	0.465	<u>0.283</u>
Ratio 0.9	17.310	0.310	0.240	43.752	0.329	0.276
RED	16.716	0.638	0.262	41.658	<u>0.463</u>	0.284

in the encoder, facilitating accurate prediction of the start position, and thereby preventing incomplete model predictions.

Furthermore, incorporating the time-distance yields notable enhancements by injecting spatial-temporal information into attention computation, increasing the awareness of neighboring paths within each trajectory path.

Road-aware Masking Strategy. The impact of varying random masking ratios, as illustrated in Table 11, is profound. A masking ratio of 0.9 degrades performance severely. Excessive masking impedes RED’s ability to extract information from its input, while overly sparse masking oversimplifies the task and prevents the model from learning. The low MAPE variations across different masks are due to the preservation of critical departure and arrival times, pivotal for accurate travel time estimation.

In contrast, the road-aware masking strategy consistently yields superior performance across metrics. By accounting for the driving patterns in trajectories, this strategy preserves essential road paths and improves training effectiveness. Adaptively adjusting the mask ratio based on trajectory properties ensures a diverse range of mask ratios within the training data. Notably, the automated nature of the road-aware masking strategy obviates the need for manual mask ratio adjustments, streamlining model design and testing.

Time-distance Correlation. We perform experiments to explore the interplay between time and distance. Specifically, we vary the value of λ_2 to evaluate the effect of time and distance on the trajectory representations. The results are displayed in Figure 6. We show only one metric for each task. Figure 6 shows that the model performance increases and then decreases with the value of λ_2 . Time information produces a higher effect when $\lambda_2 = 0.1$, while distance information produces a higher effect when $\lambda_2 = 0.9$. Both cases show a significant degradation in model performance, which implies that relying on time or distance information alone to influence the attention factor is not effective in learning the local nature of paths. Overall, performance is best when λ_2 is 0.5.

Performance versus Dimensionality. In this experiment, we vary the dimensionality among 16, 32, 64, 128, and 256, and the remaining hyper parameters are kept at default values. For brevity, we only show the HR@10 values for trajectory similarity computation, exhibit similar performance trends. The results in Figures 7 (a) and (b) indicate that the performance of RED increases with the increase in dimensionality and reaches an optimum when $d = 128$ and then starts to decrease. This is because when the dimensionality is small, the representations do not learn enough information, and too large a representation makes the model learn noisy information.

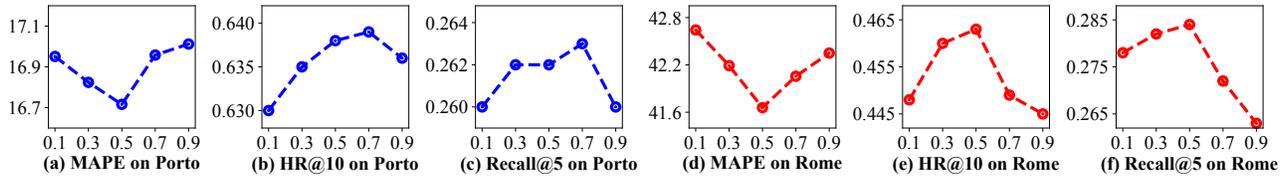


Figure 6: Influence of the time-distance attention ratio (i.e., λ_2) on the accuracy of RED.

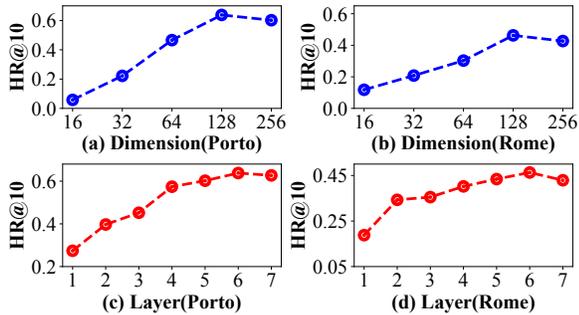


Figure 7: Effect of embedding dimension and model layers.

Performance versus Layers. We perform layer experiments, varying the number of layers among [1, 2, 3, 4, 5, 6, 7], keeping the number of encoder and decoder layers the same and the other hyper parameters are at their default values. We only show the HR@10 values for trajectory similarity computation and other tasks exhibit similar trends. The results in Figures 7 (c) and (d) show that the performance of RED increases with increasing numbers of layer and reaches an optimum when layer equals 6 for both datasets and then starts to decrease. When the number of layers is small, the model does not have the ability to extract enough features from the data, while a larger number of layers may incur overfitting.

6 RELATED WORK

Trajectory analysis is fundamental to many spatial-temporal analysis tasks and is an important aspect of urban data mining and analytics. Many methods have been proposed to learn embeddings of trajectories. For example, T2vec [23], Traj2SimVec [46], GTS [16], TrajGAT [44], and GRLSTM [51] learn trajectory representations for similarity computation. Traj2vec [45] and E²DTC [12] use the seq2seq model for trajectory clustering. Further, GM-VSAE [29] and RL4OASD [47] target anomalous trajectory detection. These methods learn trajectory representations for specific tasks. In contrast, our RED considers generic trajectory representation learning for diverse downstream tasks. Existing generic TRL methods can be classified into RNN-based and Transformer-based methods.

RNN-based Methods. Trembr [13] uses an RNN-based encoder-decoder and uses the timestamp of each road segment to learn temporal information, but it fails to capture periodic behaviors of different road types in trajectories. PIM [41] first uses node2vec [15] to learn road segment embeddings and then conducts self-supervised RNN-based mutual information maximization to learn trajectory

representation. WSCCL [42] proposes a weakly supervised contrastive learning method with curriculum learning using LSTMs [19]. HMTRL+ [26] uses GRU to model the spatial-temporal correlation of trajectories. These RNN-based methods exhibit long training times, struggle to handle long trajectories and achieve much lower accuracy than Transformer-based methods.

Transformer-based Methods. Since trajectories are sequential structures, that fit very well with the Transformer, transformer-based methods have become the main focus of the TRL research. Toast [10] proposes a traffic context aware skip-gram [32] module and a trajectory-enhanced Transformer module to learn trajectory representation. JCLRNT [31] offers domain-specific augmentations for road-road contrast and trajectory-trajectory contrast to incorporate valuable inter-relations. START [20] incorporates temporal regularities and travel semantics into generic TRL by random masking strategy. START also features data augmentation techniques involving trajectory trimming, temporal shifting, and contrastive learning. LightPath [43] uses different mask ratios to augment trajectories and then applies a teacher-student contrastive model to learn trajectory representations. However, these approaches either ignore spatial-temporal information (e.g., LightPath, Toast) or rely heavily on contrastive learning for data augmentation (e.g., JCLRNT, START, LightPath). In particular, the same data augmentation technique generally does not work well across different datasets, and also doubles the training dataset and increases the training time. In contrast, RED considers the comprehensive information of the trajectories and does not require data augmentation technologies.

7 CONCLUSION

We propose RED, a framework that utilizes self-supervised learning for trajectory representation learning. Motivated by the limitations of existing studies, RED aims to utilize comprehensive information in trajectory, including road, user, spatial, temporal, travel, and movement. For such purpose, RED features three key designs: road-aware masking strategy, spatial-temporal-user joint embedding, and dual-objective task learning. Experimental results show that RED consistently outperforms existing methods in terms of accuracy across different datasets and downstream tasks, thus advancing the state-of-the-art in trajectory representation learning.

8 ACKNOWLEDGEMENT

This paper was supported by the National Key R&D Program of China 2023YFC3305600, 2024YFE0111800, and NSFC U22B2037, U21B2046, and 62032001.

REFERENCES

- [1] Helmut Alt. 2009. The Computational Geometry of Comparing Shapes. *Efficient Algorithms, Essays Dedicated to Kurt Mehlhorn on the Occasion of His 60th Birthday* 5760 (2009), 235–248.
- [2] Helmut Alt and Michael Godau. 1995. Computing the Fréchet distance between two polygonal curves. *International Journal of Computational Geometry & Applications* 5 (1995), 75–91.
- [3] Donald J. Berndt and James Clifford. 1994. Using Dynamic Time Warping to Find Patterns in Time Series. In *AAAI*. 359–370.
- [4] Philippe C. Besse, Brendan Guillaouet, Jean-Michel Loubes, and François Royer. 2015. Review and perspective for distance based trajectory clustering. In *arXiv preprint*. <http://arxiv.org/pdf/1508.04904>.
- [5] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (2001), 5–32.
- [6] Yanchuan Chang, Jianzhong Qi, Yuxuan Liang, and Egemen Tanin. 2023. Contrastive Trajectory Similarity Learning with Dual-Feature Attention. In *ICDE*. 2933–2945.
- [7] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. 2018. Recurrent neural networks for multivariate time series with missing values. *Scientific reports* 8, 1 (2018), 6085.
- [8] Lei Chen and Raymond T. Ng. 2004. On The Marriage of Lp-norms and Edit Distance. In *VLDB*. 792–803.
- [9] Lei Chen, M. Tamer Özsu, and Vincent Oria. 2005. Robust and Fast Similarity Search for Moving Object Trajectories. In *SIGMOD*. 491–502.
- [10] Yile Chen, Xiucheng Li, Gao Cong, Zhifeng Bao, Cheng Long, Yiding Liu, Arun Kumar Chandran, and Richard Ellison. 2021. Robust Road Network Representation Learning: When Traffic Patterns Meet Traveling Semantics. In *CIKM*. 211–220.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL*. 4171–4186.
- [12] Ziquan Fang, Yuntao Du, Lu Chen, Yujia Hu, Yunjun Gao, and Gang Chen. 2021. E²DTC: An End to End Deep Trajectory Clustering Framework via Self-Training. In *ICDE*. 696–707.
- [13] Tao-Yang Fu and Wang-Chien Lee. 2020. Trembr: Exploring road networks for trajectory representation learning. *ACM Transactions on Intelligent Systems and Technology* 11, 1 (2020), 1–25.
- [14] Qiang Gao, Wei Wang, Kunpeng Zhang, Xin Yang, Congcong Miao, and Tianrui Li. 2022. Self-supervised representation learning for trip recommendation. *Knowledge-Based Systems* 247 (2022), 108791.
- [15] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. In *KDD*. 855–864.
- [16] Peng Han, Jin Wang, Di Yao, Shuo Shang, and Xiangliang Zhang. 2021. A Graph-based Approach for Trajectory Similarity Computation in Spatial Networks. In *KDD*. 556–564.
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. 2022. Masked Autoencoders Are Scalable Vision Learners. In *CVPR*. 15979–15988.
- [18] M.A. Hearst, S.T. Dumais, E. Osuna, J. Platt, and B. Scholkopf. 1998. Support vector machines. *IEEE Intelligent Systems and their Applications* 13, 4 (1998), 18–28.
- [19] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [20] Jiawei Jiang, Dayan Pan, Houxing Ren, Xiaohan Jiang, Chao Li, and Jingyuan Wang. 2023. Self-supervised Trajectory Representation Learning with Temporal Regularities and Travel Semantics. In *ICDE*. 843–855.
- [21] Seyed Mehran Kazemi, Rishab Goel, Sepehr Eghbali, Janahan Ramanan, Jaspreet Sahota, Sanjay Thakur, Stella Wu, Cathal Smyth, Pascal Poupard, and Marcus A. Brubaker. 2019. Time2vec: Learning a vector representation of time. In *arXiv preprint*. <https://arxiv.org/pdf/1907.05321>.
- [22] Ai-Te Kuo, Haiquan Chen, and Wei-Shinn Ku. 2023. BERT-Trip: Effective and Scalable Trip Representation using Attentive Contrast Learning. In *ICDE*. 612–623.
- [23] Xiucheng Li, Kaiqi Zhao, Gao Cong, Christian S. Jensen, and Wei Wei. 2018. Deep Representation Learning for Trajectory Similarity Computation. In *ICDE*. 617–628.
- [24] Yuxuan Liang, Kun Ouyang, Yiwei Wang, Xu Liu, Hongyang Chen, Junbo Zhang, Yu Zheng, and Roger Zimmermann. 2022. TrajFormer: Efficient Trajectory Classification with Transformers. In *CIKM*. 1229–1237.
- [25] Yuxuan Liang, Kun Ouyang, Hanshu Yan, Yiwei Wang, Zekun Tong, and Roger Zimmermann. 2021. Modeling Trajectories with Neural Ordinary Differential Equations. In *IJCAI*. 1498–1504.
- [26] Hao Liu, Jindong Han, Yanjie Fu, Yanyan Li, Kai Chen, and Hui Xiong. 2023. Unified route representation learning for multi-modal transportation recommendation with spatiotemporal pre-training. *The VLDB Journal* 32, 2 (2023), 325–342.
- [27] Hao Liu, Wenzhao Jiang, Shui Liu, and Xi Chen. 2023. Uncertainty-Aware Probabilistic Travel Time Prediction for On-Demand Ride-Hailing at DiDi. In *KDD*. 4516–4526.
- [28] Hongbin Liu, Hao Wu, Weiwei Sun, and Ickjai Lee. 2019. Spatio-Temporal GRU for Trajectory Classification. In *ICDM*. 1228–1233.
- [29] Yiding Liu, Kaiqi Zhao, Gao Cong, and Zhifeng Bao. 2020. Online anomalous trajectory detection with deep generative sequence modeling. In *ICDE*. 949–960.
- [30] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *ICLR*. <https://arxiv.org/pdf/1711.05101>.
- [31] Zhenyu Mao, Ziyue Li, Dedong Li, Lei Bai, and Rui Zhao. 2022. Jointly contrastive trajectory learning on road network and trajectory. In *CIKM*. 1501–1510.
- [32] Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR*. <https://arxiv.org/pdf/1301.3781>.
- [33] Simon Aagaard Pedersen, Bin Yang, and Christian S. Jensen. 2020. Fast stochastic routing under time-varying uncertainty. *The VLDB Journal* 29, 4 (2020), 819–839.
- [34] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*. 5998–6008.
- [35] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*. <https://arxiv.org/pdf/1710.10903>.
- [36] Michail Vlachos, Dimitrios Gunopulos, and George Kollios. 2002. Discovering Similar Multidimensional Trajectories. In *ICDE*. 673–684.
- [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*. 403–412.
- [38] Chengcheng Yang, Lisi Chen, Hao Wang, and Shuo Shang. 2021. Towards Efficient Selection of Activity Trajectories based on Diversity and Coverage. In *AAAI*. 689–696.
- [39] Can Yang and Gyöző Gidófalvi. 2018. Fast map matching, an algorithm integrating hidden Markov model with precomputation. *International Journal of Geographical Information Science* 32, 3 (2018), 547–570.
- [40] Song Yang, Jiamou Liu, and Kaiqi Zhao. 2022. GETNext: Trajectory Flow Map Enhanced Transformer for Next POI Recommendation. In *SIGIR*. 1144–1153.
- [41] Sean Bin Yang, Chenjuan Guo, Jilin Hu, Jian Tang, and Bin Yang. 2021. Unsupervised Path Representation Learning with Curriculum Negative Sampling. In *IJCAI*. 3286–3292.
- [42] Sean Bin Yang, Chenjuan Guo, Jilin Hu, Bin Yang, Jian Tang, and Christian S. Jensen. 2022. Weakly-supervised Temporal Path Representation Learning with Contrastive Curriculum Learning. In *ICDE*. 2873–2885.
- [43] Sean Bin Yang, Jilin Hu, Chenjuan Guo, Bin Yang, and Christian S. Jensen. 2023. LightPath: Lightweight and Scalable Path Representation Learning. In *KDD*. 2999–3010.
- [44] Di Yao, Haonan Hu, Lun Du, Gao Cong, Shi Han, and Jingping Bi. 2022. TrajGAT: A Graph-based Long-term Dependency Modeling Approach for Trajectory Similarity Computation. In *KDD*. 2275–2285.
- [45] Di Yao, Chao Zhang, Zhihua Zhu, Jian-Hui Huang, and Jingping Bi. 2017. Trajectory clustering via deep representation learning. In *IJCNN*. 3880–3887.
- [46] Hanyuan Zhang, Xinyu Zhang, Qize Jiang, Baihua Zheng, Zhenbang Sun, Weiwei Sun, and Changhu Wang. 2020. Trajectory Similarity Learning with Auxiliary Supervision and Optimal Matching. In *IJCAI*. 3209–3215.
- [47] Qianru Zhang, Zheng Wang, Cheng Long, Chao Huang, Siu-Ming Yiu, Yiding Liu, Gao Cong, and Jieming Shi. 2023. Online Anomalous Subtrajectory Detection on Road Networks with Deep Reinforcement Learning. In *ICDE*. 246–258.
- [48] Pengpeng Zhao, Anjing Luo, Yanchi Liu, Jiajie Xu, Zhixu Li, Fuzhen Zhuang, Victor S. Sheng, and Xiaofang Zhou. 2022. Where to Go Next: A Spatio-Temporal Gated Network for Next POI Recommendation. *IEEE Transactions on Knowledge and Data Engineering* 34, 5 (2022), 2512–2524.
- [49] Yusheng Zhao, Xiao Luo, Wei Ju, Chong Chen, Xian-Sheng Hua, and Ming Zhang. 2023. Dynamic Hypergraph Structure Learning for Traffic Flow Forecasting. In *ICDE*. 2303–2316.
- [50] Yu Zheng, Xing Xie, and Wei-Ying Ma. 2010. GeoLife: A Collaborative Social Networking Service among User, Location and Trajectory. *IEEE Data Engineering Bulletin* 33, 2 (2010), 32–39.
- [51] Silin Zhou, Jing Li, Hao Wang, Shuo Shang, and Peng Han. 2023. GRLSTM: Trajectory Similarity Computation with Graph-Based Residual LSTM. In *AAAI*. 4972–4980.
- [52] Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. 2017. What to Do Next: Modeling User Behaviors by Time-LSTM. In *IJCAI*. 3602–3608.