

Continuous Trajectory Generation Based on Two-Stage GAN

Wenjun Jiang¹, Wayne Xin Zhao⁴, Jingyuan Wang^{1, 2, 3*}, Jiawei Jiang¹

¹ School of Computer Science and Engineering, Beihang University, Beijing, China

² Pengcheng Laboratory, Shenzhen, China

³ School of Economics and Management, Beihang University, Beijing, China

⁴ Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China

Abstract

Simulating the human mobility and generating large-scale trajectories are of great use in many real-world applications, such as urban planning, epidemic spreading analysis, and geographic privacy protect. Although many previous works have studied the problem of trajectory generation, the continuity of the generated trajectories has been neglected, which makes these methods useless for practical urban simulation scenarios. To solve this problem, we propose a novel two-stage generative adversarial framework to generate the continuous trajectory on the road network, namely TS-TrajGen, which efficiently integrates prior domain knowledge of human mobility with model-free learning paradigm. Specifically, we build the generator under the human mobility hypothesis of the A* algorithm to learn the human mobility behavior. For the discriminator, we combine the sequential reward with the mobility yaw reward to enhance the effectiveness of the generator. Finally, we propose a novel two-stage generation process to overcome the weak point of the existing stochastic generation process. Extensive experiments on two real-world datasets and two case studies demonstrate that our framework yields significant improvements over the state-of-the-art methods.

Introduction

Modeling human mobility and generating synthetic yet realistic trajectories is crucial in many applications. On the one hand, synthetic trajectories are fundamental for urban planning(Ji et al. 2022b; Wang et al. 2020, 2018), epidemic spreading analysis(Ren, Wang, and Zhao 2022; Wang, Wang, and Wu 2018), and traffic control(Ji et al. 2022a; Wang et al. 2016), e.g., simulating changes in urban mobility in the presence of new infrastructure or international events. On the other hand, generating synthetic trajectories is a viable solution to protect the privacy of human mobility trajectory data (Pellungrini et al. 2017). The difficulty of open source sharing of trajectory data today is heavily due to privacy concerns, which hinders most existing data-driven studies. Using synthetic trajectories instead of real trajectories can not only preserve the utility of trajectory data, but also avoid leakage of user privacy. Therefore,

in such cases, it is important to generate synthetic trajectories with good data utility.

In the last decades, the problem of trajectory generation has been widely studied. In the early stage, the researchers aim to build model-based methods to model the regularity of human mobility (Chaoming et al. 2010; Jiang et al. 2016), such as temporal periodicity, spatial continuity. These methods assume that human mobility can be described by specific mobility patterns and thus can be modeled with finite parameters with explicit physical meanings. However, in fact, human mobility behaviors exhibit complex sequential transition regularities, which can be time-dependent and high-ordered. Thus, although these model-based methods have the advantage of being interpretable by design, their performance is limited due to the simplicity of the implemented mechanisms. To address the above limitation, the researchers propose model-free methods (Dan, Qing, and Liran 2018; Ouyang et al. 2018), mainly using neural network generative paradigms such as generative adversarial network and variational autoencoder. Unlike model-based methods, model-free methods abandon the extraction of specific human mobility patterns, and instead directly build a neural network to learn the distribution of the real data and generate trajectories from the same distribution.

However, comparing the above approaches, there are several key challenges remaining to be solved: (1) First, the problem of the continuity of the generated trajectories is ignored. The trajectories generated by current methods are not continuous routes on the road network, which makes these synthetic trajectories unusable for downstream applications like traffic simulation. (2) Second, those model-free methods without utilizing prior knowledge of human mobility fail to effectively generate continuous trajectories. (3) Third, the stochastic generation process of existing methods has the error accumulation problem, where the trajectory is randomly generated according to the probability given by the generator. However, once the generator predicts wrong, the process continues to generate under the wrong premise, which reduces the quality of the generated trajectory.

In this paper, we propose TS-TrajGen, a novel two-stage generative adversarial framework with spatial-temporal oriented designs in the generator, discriminator and generation process to tackle the above challenges. Specifically, we build the generator based on the human mobility hypothesis of the

*Corresponding author: jywang@buaa.edu.cn

A* algorithm (E., J., and Bertram 1968) to address the first challenges. The A* algorithm is a heuristic search algorithm used extensively on the road network. In the A* hypothesis, human mobility behaviors are determined by two factors: the observed cost from the source road to the candidate road, and the expected cost from the candidate road to the destination. Combining above two costs, A* algorithm evaluates which candidate road is the best candidate to be searched, and then generates a heuristic best continuous trajectory. Therefore, our generator consists of two parts: an attention-based network to learn the observed cost, and a GAT-based network to estimate the expected cost. As for the second challenge, we combine the sequential reward with the mobility yaw reward to improve the effectiveness of the generator, which distinguish the trajectory from the perspective of time series similarity and spatial similarity respectively. For the third challenge, we propose a two-stage generation process to solve the problem of error accumulation. In the first stage, we construct the structural regions on the top of the road network and then generate the regional trajectory. In the second stage, we generate the continuous trajectory in the guidance of the regional trajectory.

To the best of our knowledge, we are the first to solve the continuous trajectory generation problem on the urban road network, through combining the A* algorithm with neural network. In addition, to improve the effectiveness and efficiency of the generation, we build a discriminator combining the sequential reward with mobility yaw reward and propose a two-stage generation process. Extensive experiments on two real-world trajectory data have demonstrated the effectiveness and robustness of our proposed framework.

Preliminary

A trajectory in the road network can be defined as a time-ordered sequence $T = \{x_1, x_2, \dots, x_n\}$, where x_i is a spatial-temporal point defined as a tuple (l_i, t_i) . The l_i is the road segment ID and the t_i is the timestamp of x_i . The continuous trajectory is a trajectory, where each segment pair (x_i, x_{i+1}) is adjacent in the road network graph.

Definition 1. Continuous Mobility Trajectory Generation. Given a real-world mobility trajectory dataset, generate a continuous mobility trajectory $\hat{T} = \{\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n\}$ with a θ -parameterized generative model G .

In general, the process of continuous mobility trajectory generation can be regarded as a Markov decision process (MDP). The state s describes the current movement situation, composed of the current partial trajectory $x_{1:i}$ and the travel destination l_d . The action a is the next candidate road segment l_j to move. The agent is the generative model, which model the human movement policy $\pi(a|s)$.

Definition 2. Human Movement Policy. The human movement policy $\pi(a|s)$ is the conditional probability of taking action a in the state s , which can be formulated as follows:

$$\pi(a|s) = P(a|s) = P(l_j|x_{1:i}, l_d). \quad (1)$$

Thus, the generation process can be described as maximizing the probability of the generated trajectory as

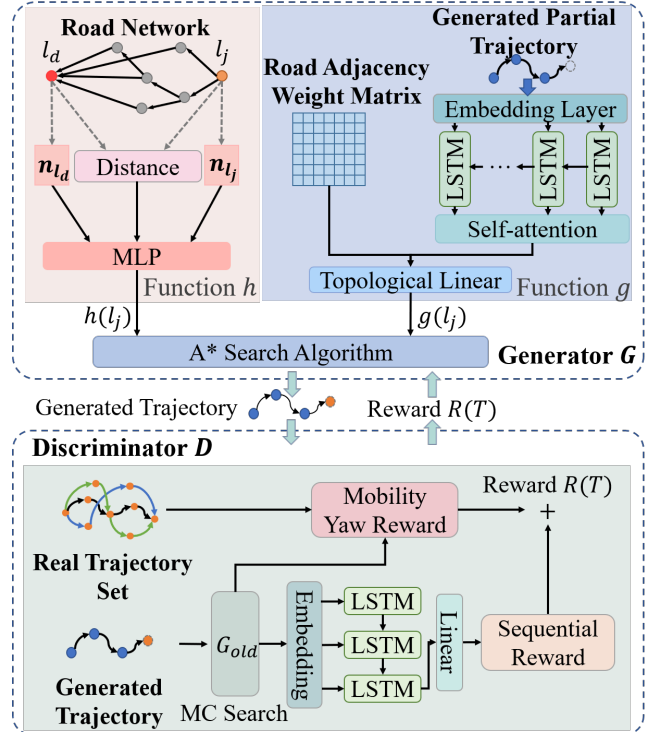


Figure 1: The overall generative adversarial framework of TS-TrajGen.

$$\hat{T} = \max \prod_{i=1}^n \pi(a_i|s_i) = \max_{P_\theta} \prod_{i=1}^n P_\theta(l_{i+1}|x_{1:i}, l_d). \quad (2)$$

The Proposed Framework

In this section, we first give a brief overview of our proposed framework and then describe each part of our framework in detail. We build a generator G to learn the human movement policy $\pi(a|s)$, and a discriminator D to generate the reward R to guide the optimization process of the generator, which is represented in Figure 1. Besides, we propose a novel two-stage generation process to generate trajectories more accurately and efficiently. Finally, we introduce the training process of our proposed framework.

Modeling Human Movement Policy by Generator

As stated in Definition 2, the key factor that determines the human movement policy $\pi(a|s)$ is the current state s , which is described by current partial trajectory $x_{1:i}$ and the travel destination l_d . To consider both the impact of the current partial trajectory and the impact of the travel destination, we adopt the hypothesis of A* algorithm to model the human movement policy $\pi(a|s)$. The A* algorithm is a heuristic search algorithm used extensively on the road network, which evaluates candidate road l_j exactly based on the above two impacts as

$$f(l_j) = g(l_j) + h(l_j), \quad (3)$$

where the function $g(l_j)$ evaluates the observed cost of the current partial trajectory, and the function $h(l_j)$ estimates the expected cost of arriving the destination l_d through the candidate road l_j .

Combining above two costs, A* algorithm learn the human movement policy as follows:

$$f(l_j) = -\log \pi(a|s) = -\log P(l_j|x_{1:i}, l_d). \quad (4)$$

With the function f , A* algorithm evaluates which candidate road is the best candidate to be searched, and then generates the heuristic best continuous trajectory.

However, there are two weak points for generating trajectories with the original A* algorithm. First, the original functions g and h are calculated based on the spherical distance between road segments, which makes it difficult to learn diverse human movement policy. Second, the spherical distance can not accurately estimate the expected cost. For example, although roads l_i and l_j are close in space, they may be one-way roads in two different directions on the road network, which makes the real expected cost between them greater than the spherical distance. Thus, we implement the functions g and h with the neural network to learn the human movement policy.

Function $g(\cdot)$. To model the observed cost, we combine the LSTM network with the self-attention mechanism to learn the current moving state and then predict the observed cost with a topologically constrained linear layer.

First, given a current partial trajectory $x_{1:i}$, we extract the spatial-temporal information through embedding each point $x_i = (l_i, t_i)$ into a dense vector \mathbf{x}_i as

$$\mathbf{x}_i = \text{Embed}(l_i) \parallel \text{Embed}(\text{encode}(t_i)), \quad (5)$$

where we encode the timestamp t_i by mapping the weekday timestamp to a group of 24×60 time slots (1 to 1440), and the weekend timestamp to another group time slots (1441 to 2880).

After embedding, we transform the current partial trajectory into a sequence of dense vectors $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_i\}$. Then, we apply the LSTM network (Hochreiter and Schmidhuber 1997) to model the sequential trajectory moving state $\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_i\}$. We further leverage the dot-product attention mechanism (Vaswani et al. 2017) to enhance the latent moving state $\tilde{\mathbf{h}}_i$ in the form of

$$\tilde{\mathbf{h}}_i = \sum_{k=1}^i \text{att}(\mathbf{h}_i, \mathbf{h}_k) \cdot \mathbf{h}_k, \quad (6)$$

where the function $\text{att}(\cdot, \cdot)$ evaluates the correlation weight between the current moving state \mathbf{h}_i and the historical moving state \mathbf{h}_k through a dot-production method.

Finally, after capturing the enhanced moving state, we construct a topologically constrained linear layer to predict the observed conditional probability $P(l_j|x_{1:i})$. In the topologically constrained linear layer, we introduce a topological adjacency weight adj_{l_i, l_j} into the linear layer, which is 1 when the candidate road l_j is adjacent to the current road l_i and 0 otherwise. With the topological adjacency weight, we set the conditional probability $P(l_j|x_{1:i})$ of the non-adjacent

candidate road l_j to zero, which ensures the continuity of the generated trajectory. The observed conditional probability is formulated as

$$P(l_j | x_{1:i}) = \frac{\exp(\mathbf{w}_{l_j} \cdot \tilde{\mathbf{h}}_i \cdot \text{adj}_{l_i, l_j})}{\sum_{l_k} \exp(\mathbf{w}_{l_k} \cdot \tilde{\mathbf{h}}_i \cdot \text{adj}_{l_i, l_k})}, \quad (7)$$

where the parameter vector \mathbf{w}_{l_j} is the l_j -th row of the weight matrix of the linear layer.

The observed cost $g(l_j)$ can be calculated as the negative log of the conditional probability as

$$g(l_j) = -\log P(l_j | x_{1:i}) \quad (8)$$

Function $h(\cdot)$. To model the expected cost, we assume that a person usually needs to know the relative position of two roads in the road network and the distance between them when estimating the expect cost of getting from a candidate road to the destination. From this perspective, we use a graph attention network to extract the relative position information from the road network and calculate the spherical distance between two road segments. Based on above information, we finally use a multilayer perceptron network to estimate the expected cost h .

First, we build a graph attention network (Veličković et al. 2018) to learn the structural road representation containing the relative position information. The update of the graph attention network can be given as

$$\mathbf{N}^{(z+1)} = \text{GAT}(\mathbf{N}^z, \mathbf{A}), \quad (9)$$

where $\mathbf{N}^z \in \mathbb{R}^{|L| \times d_s}$ is the matrix of road representation at the z -th iteration, the l_i -th row $\mathbf{n}_{l_i} \in \mathbb{R}^{d_s}$ corresponds to the representation of road segment $l_i \in L$, and \mathbf{A} is the road adjacency matrix.

For initialization, we construct the road contextual embedding vectors \mathbf{v}_{l_j} and set $\mathbf{n}_{l_i}^0 = \mathbf{v}_{l_i}$. Given a road segment l_i , we consider 6 kinds of road attributes, namely segment length, segment width, max speed limitation, lane number, road type, and longitude and latitude. We perform normalization on continuous attributes and one-hot encoding on discrete attributes, and concatenate the encoded attributes as the contextual embedding vector \mathbf{v}_{l_i} .

After obtaining the road representation, we use a multilayer perceptron network to predict the expected conditional probability from a candidate road l_j to the destination l_d as

$$P(l_j | l_d) = \text{MLP}(\mathbf{n}_{l_j}, \mathbf{n}_{l_d}, d_{l_j, l_d}). \quad (10)$$

where the \mathbf{n}_{l_j} and \mathbf{n}_{l_d} is the road representation of road l_j and l_d , and the d_{l_j, l_d} is the spherical distance between them.

In the end, the expected cost $h(l_j)$ is calculated as the negative log of the conditional probability $P(l_j|l_d)$ as

$$h(l_j) = -\log P(l_j | l_d). \quad (11)$$

Remark. Compared with the original A* algorithm, our function f is capable to model the diverse human movement policy with the introduction of temporal information and the sequential moving state. Besides, combining the structural road representation, the effectiveness of the expected cost $h(l_j)$ is improved.

Enhancing Generator with Discriminator

In general, the generative adversarial network learned by a min-max game as follows

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x} \sim p_d} [\log D_{\phi}(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim G_{\theta}} [\log(1 - D_{\phi}(\mathbf{x}))], \quad (12)$$

where $\mathbf{x} \sim p_d$ denotes the sample from the real data distribution, $\mathbf{x} \sim G_{\theta}$ denotes the sample generated by the generator, and D_{ϕ} is the ϕ -parameterized discriminator.

In the min-max game, the goal of the discriminator D_{ϕ} is to distinguish the input sample, and the generator G_{θ} optimizes itself based on the discriminator’s output signal, i.e., the reward R as

$$D_{\phi}(\mathbf{x}) = R(\mathbf{x}). \quad (13)$$

In our framework, the input sample \mathbf{x} is a trajectory T , which is a time-series with spatial information. Thus, we distinguish the input trajectory in terms of the time-series similarity and the spatial similarity, and generate corresponding reward R as

$$R(T) = R_s(T) + R_m(T), \quad (14)$$

where R_s denotes the sequential reward from the time-series aspect, where we evaluate the similarity based on the hidden sequential transition pattern extracted by an LSTM network. R_m denotes the mobility yaw reward from the spatial aspect, where we evaluate the similarity according to the trajectory’s mobility yaw distance from the real trajectories.

Sequential Reward. To extract the sequential transition pattern, we build a sequential discriminator D_s following the same steps in the function g : we first embed the input trajectory, and then obtain the hidden sequential transition pattern \mathbf{h} based on an LSTM network. Then, we use a linear layer to predict the probability that the trajectory is real, which is the sequential reward R_s as

$$R_s(T) = \begin{cases} D_s(T) & \text{if } T \text{ is finished;} \\ \frac{1}{N} \sum_{n=1}^N D_s(\text{MC}(T)) & \text{else,} \end{cases} \quad (15)$$

where we further leverage the Monte Carlo search to evaluate the average sequential reward for the intermediate step in the trajectory.

This is because the trajectory from the generator is sequentially generated, which requires the discriminator to provide rewards for each step of the trajectory. However, current discriminator can not provide proper reward for the intermediate step. For the intermediate step i , the current trajectory $x_{1:i}$ is still unfinished. This means that the reward, based on current sub-trajectory $x_{1:i}$, only considers the sequential transition pattern of the current sub-trajectory but ignores the future outcome. Thus, following the previous work (Yu et al. 2017), we apply Monte Carlo search to evaluate the intermediate step. To do this, we first maintain an older version generator G_{old} , which is the last step version of the current generator. Then, we use the G_{old} to complete the current sub-trajectory by repeating Monte Carlo search N times. The complete trajectories are then fed into the D_s to generate sequential reward R_s for the intermediate step.

Mobility Yaw Reward. The mobility yaw distance of a input trajectory T is defined as the minimum distance to the real trajectories set S_{od} in the form of

$$dis = \min\{\text{DTW}(T, T') \mid T' \in S_{od}\}, \quad (16)$$

where we leverage the widely-used trajectory distance metric DTW (Keogh 2002) to measure the mobility yaw distance, and the real trajectories set S_{od} contains the real trajectories shared the same OD with the input trajectory T .

Furthermore, we calculate the mobility yaw distance for each step of the trajectory T by the Monte Carlo search. The mobility yaw reward is defined as the change in yaw distance $\Delta dis_i = dis_i - dis_{i-1}$ as

$$R_m(T) = \text{norm}(\Delta dis_i), \quad (17)$$

where we further apply a normalization to dimensionless the change of mobility yaw distance.

The change in yaw distance Δdis_i indicates how much the current step’s mobility yaw distance has decreased compared to the previous step. The more the mobility yaw distance is decreased, the more correct the decision of the current step is. Correspondingly, the reward for the current step should be larger.

Two-Stage Generation Process

The stochastic generation process has the problem of error accumulation, where trajectory is randomly generated according to the probability given by the generator as

$$P_{\theta}(\hat{T}) = \prod_{i=1}^n \pi(a_i | s_i) = \prod_{i=1}^n P_{\theta}(l_{i+1} | x_{1:i}, l_d). \quad (18)$$

However, once the generator predicts wrong, the stochastic generation process continues to generate under the wrong premise state. Especially, when generating long trajectories, the probability of the generator making an error increases as the number of predictions made by the generator increases. This makes it difficult to reach the destination and reduces the quality of the generated trajectory.

To solve the above problem, we propose a novel two-stage generation process based on A* search. We first construct the structural regions above the road network. Then, we generate the regional trajectory in the first stage. In the second stage, we generate the continuous trajectory in the guidance of the regional trajectory. With the A* search, our generation process can roll back to correct the error, when the generator predicts incorrectly. Because, in the A* search, all possible trajectories currently searched are maintained, and the trajectory with the least cost (highest probability) is selected for searching each time. In addition, with the two-stage generation, not only the space and time complexity of the generation process is reduced, but the generation of long trajectories is also avoided.

Building Structural Region. We leverage the multilevel graph partitioning algorithm KaFFPa (Sanders and Schulz 2013) to construct structural regions on the top of the road network. It partitions the graph into k blocks under the constraint that the number of edges between blocks is minimized and the maximum block size does not exceed $1 + \epsilon$

Algorithm 1: Process of Mapping the Regional Trajectory

Input: The regional trajectory $T_R = \{r_1, r_2, \dots, r_n\}$ and the origin l_s and destination l_d .

Output: The discontinuous road segment trajectory T .

- 1: Initialize the T as $\{l_s\}$.
 - 2: **for** $i = 1$ to $n - 1$ **do**
 - 3: Randomly choose a boundary road segment l from the boundary road segments set $S_{r_i, r_{i+1}}$ based on the historical visit frequency.
 - 4: Append the boundary road segment l to the T .
 - 5: **end for**
 - 6: Append the l_d to the T .
 - 7: Return T .
-

times the average block size. With the KaFFPa algorithm, we can obtain relatively independent structural regions of similar size, and the mapping relationship between road and region. Formally, we describe the mapping relationship between road and region as the mapping matrix $M \in \mathbb{R}^{|L| \times |R|}$, where each element is defined as

$$m_{l_i, r_j} = \begin{cases} 1, & l_i \in r_j; \\ 0, & \text{other,} \end{cases} \quad (19)$$

where the l_j is the i -th road segment in road segment set L , and the r_j is the j -th structural region in region set R .

Based on the mapping matrix M , the region-level data can be obtained through mapping from the road-level data. Then, we use the same proposed adversarial generative framework to learn the region-level human movement policy. See the appendix for more details.

Regional Trajectory Generation. In the first stage, we aim to generate the regional trajectory. At the beginning, we sample a tuple of origin and destination (l_s, l_d) and the start timestamp t_s from the historical OD matrix, which can be counted from the given real-world trajectory dataset. Then, we map the road-level origin and destination to the region-level origin and destination (r_s, r_d) according to the road-to-region mapping relationship, where we map the l_s to r_s only if the $m_{l_s, r_s} = 1$. With the region-level origin and destination and the start timestamp as input, we generate regional trajectory based on the A* search and the region-level generator.

Road-level Continuous Trajectory Generation. In the second stage, we generate the continuous trajectory in the guidance of the regional trajectory. In practice, we construct the boundary road segments set S_{r_i, r_j} between each region r_i and region r_j , and count the visiting frequency of each boundary road segment in advance. Then, the regional trajectory can be mapped to the road-level, according to the Algorithm 1.

After obtaining the discontinuous road segment trajectory T , we use the road-level generator to complete the trajectory between each pair (l_i, l_{i+1}) in the T , which is the final continuous trajectory. In this way, we only need to generate in a small structural region instead of the huge road network.

Data Statistics	BJ-Taxi	Porto-Taxi
Number of Taxis	15,642	435
Number of Roads	40,306	11,095
Number of Intersections	16,927	5,184
Number of Trajectories	956,070	695,085

Table 1: Basic statistics of two real-world datasets.

Training Process

At the beginning of the training, we use the trajectory next-location prediction task to pre-train the function g and h separately, as they model human movement policy from different perspectives. Then, we pre-train the sequential discriminator based on the generated trajectories and the real trajectories. After the pre-training, we follow the REINFORCE algorithm (Williams 1992) to generate the policy gradient by receiving the reward $R(T)$ from the discriminator D as

$$\nabla_{\theta} = \nabla_{\theta} \mathbb{E}_{p_{\theta}(T)}[R(T)] = \mathbb{E}_{p_{\theta}(T)}[R(T) \nabla_{\theta} \log p_{\theta}(T)], \quad (20)$$

where θ is the parameter of the generator G , T is the generated trajectory, the reward $R(T)$ is the sum of sequential reward R_s and mobility yaw reward R_m from the discriminator D . Based on the policy gradient ∇_{θ} , parameter θ is updated by $\theta \leftarrow \theta + \alpha \nabla_{\theta}$, where α is the learning rate.

Experiment

We first evaluate TS-TrajGen in terms of macro-similarity and micro-similarity between the synthetic trajectories and the real trajectories. Next, we conduct two case studies: the trajectory next-location prediction and the traffic control simulation, to assess the utility of synthetic trajectories.

Experiment Setting

Dataset. We use two real-world mobility dataset to measure the performance of our proposed framework. The *BJ-Taxi* dataset contains real GPS trajectory data of Beijing taxis for one week from November 1 to 7, 2015 within Beijing’s Fourth Ring Roads, which is sampled every minute (Jiang et al. 2023). The *Porto-Taxi* dataset is originally released for a Kaggle trajectory prediction competition with a sample period of 15 seconds. For the two datasets, we collect corresponding road network information from the open street map, and then perform map matching (Yang and Gidofalvi 2018) to obtain real continuous trajectories. To eliminate abnormal trajectories, we remove trajectories with lengths less than 5 and trajectories with loops. The statistics of the two datasets after preprocessing are shown in Table 1. For each trajectory dataset, we randomly split the whole dataset into three parts in the ratio of 6:2:2: a training set, a validation set, and a test set.

Comparative Benchmarks. We compare TS-TrajGen with five baselines to validate the performance of our proposed framework: SeqGAN (Yu et al. 2017), SVAE (Huang et al. 2019), MoveSim (Feng et al. 2020), TSG (Wang et al. 2021b), and TrajGen (Cao and Li 2021). For more information about the baselines, please refer to the appendix.

Performance Criteria. From the macro perspective, we mainly consider the overall distribution of the trajectory dataset. Following the common practice in previous works (Cao and Li 2021; Feng et al. 2020), we quantitatively assess the quality of the generated data by computing the similarity of 4 important mobility patterns between the generated and real data: the travel distance, the radius of gyration, the location distribution and the OD flow. To obtain quantitative results, we use Jensen-Shannon Divergence to calculate the similarity on the four metrics. Refer to the appendix for more information on the metrics.

From the microscopic perspective, we focus on measuring the similarity between the real trajectories and the generated trajectories with the same OD. In practice, we use three widely used trajectory distance metrics to measure the micro similarity: Hausdorff distance (Xie, Li, and Phillips 2017), DTW (Keogh 2002), and EDR (Chen, Özsu, and Oria 2005). We randomly select 5,000 trajectories from the test set, and make the generative model generate the corresponding trajectory. Finally, the average distance between the generated trajectories and the real trajectories is taken as the evaluation result.

Evaluation Results

The evaluation results of macro-similarity and micro-similarity are shown in Table 2. In each row, the best result is highlighted in boldface and the second best is underlined>. From the statistics, we can first conclude that our proposed TS-TrajGen significantly outperforms all baselines in terms of both macro-similarity metrics and micro-similarity metrics.

From the perspective of macro-similarity, compared with the second best model, TS-TrajGen improves the *Distance*, *Radius*, and *Location Frequency* metrics by 97.57%, 98.94% and 70.09% respectively in *BJ-Taxi*. As for the *OD Flow* metric, since we take both the origin and destination as input, TS-TrajGen achieve to completely restore the real OD flow state of the city. Among the baselines, TrajGen performs best because the trajectories it generates are derived from images, which enables it to generate continuous trajectories as the image resolution increases. While other sequential generative baselines fail to model the continuous mobility pattern. Besides, to better illustrate the detailed difference, we leverage heat map to visualize the *location distribution* of the top three generated data in Beijing, as shown in Figure 2.

As for the micro-similarity, our proposed TS-TrajGen achieves to simulate the individual trajectory, and shows clear advantages over others, which is mainly because other baseline models ignore the impact of destination on the human mobility behavior.

Case Study

Case Study 1: Trajectory Next-Location Prediction.

Trajectory next-location prediction task aims to predict the location of mobility trajectory through mining the human mobility patterns. Therefore, we can leverage this task to examine whether there are real mobility patterns in the generated data based on *BJ-Taxi* dataset.

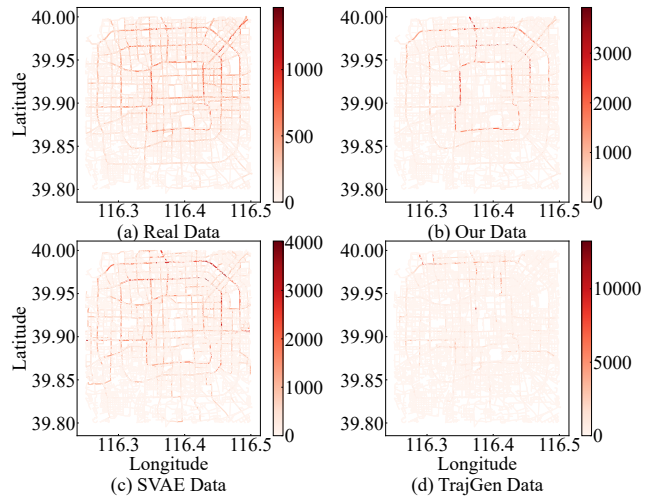


Figure 2: Visualized *location distribution* of different trajectory data.

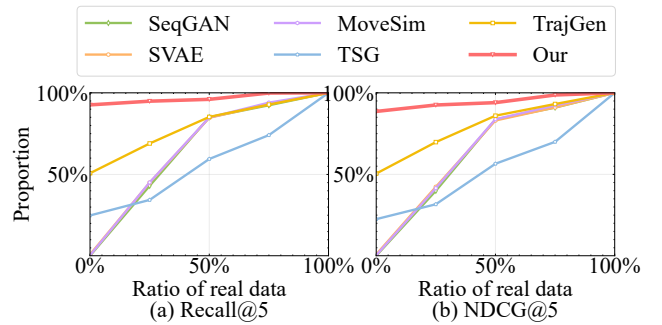


Figure 3: The proportion of performance of the model trained on mixed data over the performance of the model trained on real data.

In this case study, we choose DeepMove (Feng et al. 2018) implemented by LibCity (Wang et al. 2021a) as the trajectory next-location prediction model. We mix the generated data with the real data in various ratios and train the DeepMove model using the mixed data. Then, we test the model performance on real data and calculate the performance difference from the model trained entirely on real data.

Figure 3 shows the performance difference between the model trained on the mixed data and the model trained on the real data, e.g., only the real data is used for training when the x value is 100%. We can conclude the utility of our generated data exceeds the data generated by other models. When only trained with our generated data, the DeepMove achieves 92.61% and 88.59% performance of the DeepMove trained on real data on the *Recall@5* and *NDCG@5* metrics respectively. However, among the baselines, the best proportions of performance are 50.74% and 50.47% respectively.

Case Study 2: Traffic Control Simulation.

Traffic simulation is an important application of trajectory generation. Thus, we choose a case of Beijing Sanyuan Bridge traffic

	BJ-Taxi						Porto-Taxi					
	SeqGAN	SVAE	MoveSim	TSG	TrajGen	Our	SeqGAN	SVAE	MoveSim	TSG	TrajGen	Our
Distance	0.495	0.303	0.537	0.609	<u>0.225</u>	0.005	0.350	0.361	0.473	0.509	<u>0.136</u>	0.008
Radius	0.472	0.236	0.407	<u>0.095</u>	0.138	0.002	0.303	0.187	0.202	0.128	<u>0.065</u>	0.002
Loc Dist	<u>0.112</u>	0.153	0.406	0.536	0.255	0.076	0.143	<u>0.134</u>	0.316	0.417	0.365	0.120
OD Flow	<u>0.652</u>	0.683	0.636	0.651	<u>0.632</u>	0.000	0.337	<u>0.627</u>	<u>0.250</u>	0.553	0.552	0.000
Hausdorff	13.85	<u>10.88</u>	13.41	11.07	11.65	0.76	4.19	<u>3.36</u>	4.17	3.88	4.37	0.65
DTW	<u>199.51</u>	430.88	240.16	2446.58	427.75	13.71	<u>82.55</u>	133.15	101.18	704.39	198.90	18.12
EDR	0.999	0.998	<u>0.945</u>	0.999	0.997	0.452	0.995	0.979	<u>0.953</u>	0.997	0.993	0.558

Table 2: Performance comparison of our model and baselines on two real-world datasets, where lower results are better. The first four lines are macro-similarity metrics, and the last three lines are micro-similarity metrics.



(a) Normal situation (b) Traffic control (c) Simulation result

Figure 4: Visualized vehicle trajectories near Sanyuan Bridge 4a) when there is no traffic control, 4b) when there is traffic control, and 4c) which is the simulation result.

control to verify that our framework is competent for the human mobility simulation task.

On November 13, 2015, the Sanyuan Bridge on the Third Ring Road in Beijing was imposed a one-day traffic control due to the construction reason. We randomly visualized 15 trajectories without and with traffic control near the Sanyuan Bridge, as shown in Figure 4a and Figure 4b, where the green lines are the Beijing road network, the blue lines are the trajectories of different vehicles, and the area in the green box is the Sanyuan Bridge.

In this case study, we set the roads under construction as unreachable and then simulate trajectories nearby Sanyuan Bridge. The part of simulation result is visualized in Figure 4c. In addition, we quantitatively calculate the micro similarity between real trajectories and generated trajectories. However, since none of the above baseline models can handle the traffic control situation, we chose two route planning models as baseline models in this case study, i.e., NASR(Wang et al. 2019) and A* search(E., J., and Bertram 1968). The evaluation result is shown in Table 3. From the test result, our TS-TrajGen is closer to the real traffic control situation than existing route planning method NASR, and achieves significant progress on the first two metrics.

	Hausdorff	DTW	EDR
A* search	2.7599	49.3791	0.7160
NASR	<u>2.0393</u>	<u>37.1353</u>	<u>0.5658</u>
Our	0.9790	16.3914	0.5451

Table 3: The micro similarity between real trajectories and simulated trajectories.

Related Work

Existing works of trajectory generation can be classified into two groups: model-based methods and model-free methods. Based on assumptions about the human mobility regularity, model-based methods model the individual mobility by limited parameters. However, these methods with simplified assumptions can not model the complex mobility in reality.

Recently, the researchers turned to apply model-free methods like VAE and GAN to generate the mobility trajectory. SAVE (Huang et al. 2019) is a combination of VAE and Seq2Seq model, which combines the ability of VAE to learn the underlying distribution of real data and the ability of LSTM to process sequential data. MoveSim (Feng et al. 2020) is a model-free GAN framework that integrates the domain knowledge of human mobility regularity. TrajGen (Cao and Li 2021) maps a trajectory into an image and utilizes a CNN-based GAN to generate the virtual trajectory image. Then, TrajGen applies a Seq2Seq model to infer the real trajectory sequence from the location sequence extracted from the virtual image. Although these methods do not rely on the simplified human mobility assumptions, but instead directly learned the underlying distribution of the data. However, these methods ignore the problem of the continuity of the generated trajectory, which is crucial in practical human mobility simulation applications. Hence, we propose a novel two-stage generative adversarial framework to solve this problem.

Conclusion

In this paper, we propose a novel two-stage generative adversarial framework to solve the continuous trajectory generation problem, which combines the advantages of human mobility prior knowledge with the model-free learning paradigm. Although many previous works have investigated the trajectory generation problem, we are the first to directly generate continuous trajectories on the road network. The extensive experiments demonstrate that our framework outperforms five state-of-the-art baselines significantly, including aspects of macro-similarity, micro-similarity and data utility. Besides, through a traffic control simulation case study, we prove that our framework can be directly applied to real-world city traffic simulations. As future work, we will further explore other potential factors of human mobility and extend the simulation to various applications.

Acknowledgments

This work was supported by the National Key R&D Program of China (Grant No. 2019YFB2102100). Prof. Wang's work was supported by the National Natural Science Foundation of China (No. 82161148011, 72222022, 72171013), the Fundamental Research Funds for the Central Universities (YWF-22-L-838) and the DiDi Gaia Collaborative Research Funds. Prof. Zhao's work was supported by the National Natural Science Foundation of China (No. 6222215).

References

- Cao, C.; and Li, M. 2021. Generating Mobility Trajectories with Retained Data Utility. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '21, 2610–2620. New York, NY, USA: Association for Computing Machinery. ISBN 9781450383325.
- Chaoming, S.; Tal, K.; Pu, W.; and Albert-László, B. 2010. Modelling the scaling properties of human mobility. *Nature physics*, 6(10): 818–823.
- Chen, L.; Özsü, M. T.; and Oria, V. 2005. Robust and Fast Similarity Search for Moving Object Trajectories. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, SIGMOD '05, 491–502. New York, NY, USA: Association for Computing Machinery. ISBN 1595930604.
- Dan, Y.; Qing, Y.; and Liran, M. 2018. GANs Based Density Distribution Privacy-Preservation on Mobility Data. *Sec. and Commun. Netw.*, 2018.
- E., H. P.; J., N. N.; and Bertram, R. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.
- Feng, J.; Li, Y.; Zhang, C.; Sun, F.; Meng, F.; Guo, A.; and Jin, D. 2018. DeepMove: Predicting Human Mobility with Attentional Recurrent Networks. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, 1459–1468. Republic and Canton of Geneva, CHE: International World Wide Web Conferences Steering Committee. ISBN 9781450356398.
- Feng, J.; Yang, Z.; Xu, F.; Yu, H.; Wang, M.; and Li, Y. 2020. Learning to Simulate Human Mobility. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '20, 3426–3433. New York, NY, USA: Association for Computing Machinery. ISBN 9781450379984.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.*, 9(8): 1735–1780.
- Huang, D.; Song, X.; Fan, Z.; Jiang, R.; Shibasaki, R.; Zhang, Y.; Wang, H.; and Kato, Y. 2019. A Variational Autoencoder Based Generative Model of Urban Human Mobility. In *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, 425–430.
- Ji, J.; Wang, J.; Jiang, Z.; Jiang, J.; and Zhang, H. 2022a. STDEN: Towards Physics-Guided Neural Networks for Traffic Flow Prediction. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, 4048–4056. AAAI Press.
- Ji, J.; Wang, J.; Wu, J.; Han, B.; Zhang, J.; and Zheng, Y. 2022b. Precision CityShield Against Hazardous Chemicals Threats via Location Mining and Self-Supervised Learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3072–3080.
- Jiang, J.; Pan, D.; Ren, H.; Jiang, X.; Li, C.; and Wang, J. 2023. Self-supervised Trajectory Representation Learning with Temporal Regularities and Travel Semantics. In *2023 IEEE 39th international conference on data engineering (ICDE)*. IEEE.
- Jiang, S.; Yang, Y.; Gupta, S.; Veneziano, D.; Athavale, S.; and González, M. C. 2016. The TimeGeo modeling framework for urban mobility without travel surveys. *Proceedings of the National Academy of Sciences*, 113(37): E5370–E5378.
- Keogh, E. 2002. Chapter 36 - Exact Indexing of Dynamic Time Warping. In Bernstein, P. A.; Ioannidis, Y. E.; Ramakrishnan, R.; and Papadias, D., eds., *VLDB '02: Proceedings of the 28th International Conference on Very Large Databases*, 406–417. San Francisco: Morgan Kaufmann. ISBN 978-1-55860-869-6.
- Ouyang, K.; Shokri, R.; Rosenblum, D. S.; and Yang, W. 2018. A Non-Parametric Generative Model for Human Trajectories. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, 3812–3817. International Joint Conferences on Artificial Intelligence Organization.
- Pellungrini, R.; Pappalardo, L.; Pratesi, F.; and Monreale, A. 2017. A Data Mining Approach to Assess Privacy Risk in Human Mobility Data. *ACM Trans. Intell. Syst. Technol.*, 9(3).
- Ren, H.; Wang, J.; and Zhao, W. X. 2022. Generative Adversarial Networks Enhanced Pre-training for Insufficient Electronic Health Records Modeling. In Zhang, A.; and Rangwala, H., eds., *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18, 2022*, 3810–3818. ACM.
- Sanders, P.; and Schulz, C. 2013. Think Locally, Act Globally: Highly Balanced Graph Partitioning. In *Proceedings of the 12th International Symposium on Experimental Algorithms (SEA'13)*, volume 7933 of LNCS, 164–175. Springer.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- Wang, J.; Gu, Q.; Wu, J.; Liu, G.; and Xiong, Z. 2016. Traffic Speed Prediction and Congestion Source Exploration: A Deep Learning Method. In Bonchi, F.; Domingo-Ferrer, J.;

Baeza-Yates, R.; Zhou, Z.; and Wu, X., eds., *IEEE 16th International Conference on Data Mining, ICDM 2016, December 12-15, 2016, Barcelona, Spain*, 499–508. IEEE Computer Society.

Wang, J.; He, X.; Wang, Z.; Wu, J.; Yuan, N. J.; Xie, X.; and Xiong, Z. 2018. CD-CNN: A Partially Supervised Cross-Domain Deep Learning Model for Urban Resident Recognition. In McIlraith, S. A.; and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, 192–199. AAAI Press.

Wang, J.; Jiang, J.; Jiang, W.; Li, C.; and Zhao, W. X. 2021a. LibCity: An Open Library for Traffic Prediction. In *Proceedings of the 29th International Conference on Advances in Geographic Information Systems, SIGSPATIAL '21*, 145–148. New York, NY, USA: Association for Computing Machinery. ISBN 9781450386647.

Wang, J.; Wang, X.; and Wu, J. 2018. Inferring Metapopulation Propagation Network for Intra-City Epidemic Control and Prevention. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '18*, 830–838. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355520.

Wang, J.; Wu, J.; Wang, Z.; Gao, F.; and Xiong, Z. 2020. Understanding Urban Dynamics via Context-Aware Tensor Factorization with Neighboring Regularization. *IEEE Trans. Knowl. Data Eng.*, 32(11): 2269–2283.

Wang, J.; Wu, N.; Zhao, W. X.; Peng, F.; and Lin, X. 2019. Empowering A* Search Algorithms with Neural Networks for Personalized Route Recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '19*, 539–547. New York, NY, USA: Association for Computing Machinery. ISBN 9781450362016.

Wang, X.; Liu, X.; Lu, Z.; and Yang, H. 2021b. Large Scale GPS Trajectory Generation Using Map Based on Two Stage GAN. *Journal of Data Science*, 19(1): 126–141.

Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3): 229–256.

Xie, D.; Li, F.; and Phillips, J. M. 2017. Distributed Trajectory Similarity Search. *Proc. VLDB Endow.*, 10(11): 1478–1489.

Yang, C.; and Gidofalvi, G. 2018. Fast map matching, an algorithm integrating hidden Markov model with precomputation. *International Journal of Geographical Information Science*, 32(3): 547–570.

Yu, L.; Zhang, W.; Wang, J.; and Yu, Y. 2017. SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17*, 2852–2858. AAAI Press.