

UC Davis

UC Davis Previously Published Works

Title

MobilityGPT: Enhanced Human Mobility Modeling With a GPT Model

Permalink

<https://escholarship.org/uc/item/6q5753r8>

Journal

IEEE Transactions on Intelligent Transportation Systems, 27(1)

ISSN

1524-9050

Authors

Haydari, Ammar

Chen, Dongjie

Lai, Zhengfeng

et al.

Publication Date

2026

DOI

10.1109/tits.2025.3626357

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

MobilityGPT: Enhanced Human Mobility Modeling with a GPT model

Ammar Haydari, Dongjie Chen, Zhengfeng Lai, Michael Zhang, Chen-Nee Chuah, *Fellow, IEEE*

Abstract—Generative models have shown promising results in capturing human mobility characteristics and generating synthetic trajectories. However, it remains challenging to ensure that the generated geospatial mobility data is semantically realistic, including consistent location sequences, and reflects real-world characteristics, such as constraining on geospatial limits. We reformat human mobility modeling as an autoregressive generation task to address these issues, leveraging the Generative Pre-trained Transformer (GPT) architecture. To ensure its controllable generation to alleviate the above challenges, we propose a geospatially-aware generative model, MobilityGPT. We propose a gravity-based sampling method to train a transformer for semantic sequence similarity. Then, we constrained the training process via a road connectivity matrix that provides the connectivity of sequences in trajectory generation, thereby keeping generated trajectories in geospatial limits. Lastly, we proposed to construct a preference dataset for fine-tuning MobilityGPT via Reinforcement Learning from Trajectory Feedback (RLTF) mechanism, which minimizes the travel distance between training and the synthetically generated trajectories. Experiments on real-world datasets demonstrate MobilityGPT’s superior performance over state-of-the-art methods in generating high-quality mobility trajectories that are closest to real data in terms of origin-destination similarity, trip length, travel radius, link, and gravity distributions. We release the source code and reference links to datasets at <https://github.com/ammahydr/MobilityGPT>.

Index Terms—Human Mobility Modeling, Spatial-Temporal Systems, Generative Models, Generative Pretrained Transformer

I. INTRODUCTION

The widespread integration of location-based services and smart GPS devices, such as smartphones and watches, has made continuous monitoring of human mobility both desirable and feasible. These technologies capture diverse and detailed human movement information, with mobility trajectories representing the finest granularity of individual-level mobility characteristics. Such trajectories are crucial in various applications, including mobility modeling, commercial business analysis, and disease spread control [1, 2, 3]. Understanding and modeling these mobility patterns has become increasingly critical as population growth and complex travel behaviors reshape modern transportation systems. The impacts of human mobility extend beyond transportation, influencing urban planning, environmental quality, public health, and energy consumption.

Ammar Haydari, Dongjie Chen and Chen-Nee Chuah are with the Department of Electrical and Computer Engineering, University of California, Davis, CA, 95616 USA (e-mail:ahaydari@ucdavis.edu; cdjchen@ucdavis.edu; chuah@ucdavis.edu).

Zhengfeng Lai is with Apple Inc. Cupertino, CA, 95014 USA (e-mail: lzhenfeng@ucdavis.edu).

Michael Zhang is with the Department of Civil and Environmental Engineering, University of California, Davis, CA 95616 USA (e-mail: hmzhang@ucdavis.edu).




Fig. 1. Generative models that do not incorporate mobility characteristics struggle to capture human mobility patterns accurately.

Recent global events, particularly the COVID-19 pandemic, have highlighted the dynamic nature of mobility patterns and their profound societal implications [4, 5]. This has intensified the need for sophisticated mobility modeling approaches that can capture, analyze, and predict human movement patterns at various spatial and temporal scales.

Despite increasing demand for human mobility trajectory datasets, numerous challenges hinder their access and distribution [6, 7]. First, these datasets are typically collected by private companies or government agencies, hence presenting privacy concerns due to the potential disclosure of an individual’s sensitive lifestyle patterns (e.g., home or office addresses and points of interest). Legal implications, such as those outlined in the California Consumer Privacy Act [8] and General Data Protection Regulation [9], highlight the importance of carefully handling spatiotemporal traces. Second, datasets owned by companies may expose proprietary business models and are often inaccessible for research purposes [10]. Lastly, publicly available datasets often lack diversity or quality, with gaps in data points and intrinsic noise that significantly reduce their utility. These limitations impede the progress of urban planning or transportation research [11]. Therefore, it is imperative to establish alternative trajectory data sources that are both high-quality and accessible for research purposes.

Current generative models designed for human mobility exhibit multiple drawbacks: they fail to capture sequential mobility characteristics, often lack the continuity of generated trajectories, and do not consistently adhere to geospatial constraints. Approaches based on the Generative Adversarial Network (GAN) [12] and the Variational Autoencoder (VAE) [13] structure data in a tabular format, which unfortunately fails to preserve the inherent correlations between locations.

Models like Long-Short-Term Memory (LSTM) and sequential GAN [14] struggle to maintain smooth continuity of locations and capture realistic human mobility characteristics. A further limitation of these methods lies in the elevated mismatches observed in geospatial location sequences. In short, a successful generative model for trajectory datasets should precisely grasp the intrinsic spatial-temporal behaviors.

GPT (Generative Pre-trained Transformer) is a language model architecture that learns to generate coherent sequences by predicting the next token in a sequence, based on the context of previous tokens. It is trained in an autoregressive manner and excels at modeling sequential dependencies. This approach offers a robust foundation for sequence generation tasks and can be refined through fine-tuning to cater to distinct objectives.

Trajectories and sentences share several similarities that make GPTs a promising approach for trajectory generation. Both consist of ordered sets of elements chosen from finite pools (road links and words, respectively). They exhibit semantic or spatiotemporal relationships, adhering to distinct rule systems such as language rules for sentences and geographical constraints for trajectories. Due to these parallels, the techniques developed for natural language processing can be adapted to model and generate realistic trajectories [15]. In this work, we adapt the GPT architecture to the domain of human mobility by treating trajectories as sequences of spatial tokens (e.g., road links). We found that the self-attention mechanism captures intricate dependencies in human mobility sequences. This positions such models as powerful tools for overcoming challenges in human trajectory generation. Besides, the gravity model holds significance in human mobility modeling as it provides a structured framework for estimating and understanding the flow of interactions between different locations [16]. Using gravity as part of human mobility modeling could bring further advantages to synthetic trajectory generation tasks.

We leverage the power of GPT architecture for synthetic trajectory generation. However, directly training a GPT on mobility data without geospatial insights would generate unrealistic sequences. To generate more realistic synthetic trajectories, we introduced several innovative methods, including (1) gravity-aware sampling, which incorporates gravity modeling of trajectory data into training updates, and (2) the use of road connectivity matrix (RCM) masking to eliminate disconnected location sequences from logits (see Figure 1). Pretraining GPT models on trajectory data captures the general sequence characteristics in an unstructured manner. We, furthermore, proposed an automated fine-tuning pipeline that improves the trajectory quality by leveraging transportation-specific metrics to evaluate and optimize the generated sequences without relying on human labeling. Our multi-objective framework, MobilityGPT, not only extends the application of GPT models to human mobility modeling but also collectively enhances the model’s capacity to capture intricate patterns and adhere to realistic geospatial constraints (e.g., trajectories consist of road segments that are indeed connected). MobilityGPT can generate contextually rich and semantically accurate trajectories that reflect realistic human mobility without requiring prohibitive computational complexity.

Contributions: In summary, this paper addresses concerns

like trajectory continuity, adherence to geospatial constraints, and computational complexity in existing models. The core of our contribution is MobilityGPT, a human mobility modeling framework that features geospatially aware pretraining and automated fine-tuning without requiring human feedback. We enhance trajectory generation by implementing gravity-aware sampling that trains our generative model using Origin-Destination (OD) pair-gravity values, while ensuring trajectory continuity through next sequence prediction conditioned on the road network through RCM. To improve trajectory quality, we developed a strategy to construct preference datasets for fine-tuning MobilityGPT via reinforcement learning, which significantly enhances travel length similarity in generated paths. Our comprehensive experimental analysis with real-world datasets demonstrates that these methods produce high-fidelity trajectories that maintain both the statistical properties and semantic characteristics essential to authentic human mobility patterns.

The remainder of the paper is organized as follows. We present related work in Section II, establish our preliminary definitions in Section III, and introduce MobilityGPT components in Section IV. After empirical evaluation in Section V, we discuss limitations in Section VI. Finally, Section VII concludes the paper.

II. RELATED WORK

Human mobility simulation presents a significant challenge in mobility data mining due to complex geospatial and temporal correlations in movement sequences. Existing methods aim to generate realistic trajectories that reflect human behavior. Many focus on movement prediction [17, 18] or trajectory reconstruction [19]. However, they often miss dynamic mobility patterns. They also struggle to generalize across large-scale datasets, limiting their use in comprehensive mobility modeling.

The methods for synthetic mobility trajectory generation can be broadly categorized into knowledge-driven *model-based* and data-driven *model-free* approaches. The latter predominantly leveraging machine learning techniques. *Model-based* methods rely on theoretical frameworks and domain-specific hypotheses, such as gravity models and spatial interaction theory, to simulate mobility patterns. Examples include microscopic traffic simulators like SUMO [20] and VISSIM [21]. These methods, however, often struggle to capture the complexity and variability of real-world human mobility, particularly at finer granularities.

There exists another important distinction in the trajectory literature between point-based and point-of-interest-based generative models. Point-based location sequences are dense and contain richer human movement information at high granularity. Recently, TrajGPT [22] proposed a visitation-based sequence generation model that generates sparse location sequences based on visited locations. However, this method falls short in dense sequence prediction tasks. MobilityGPT takes a fundamentally different approach from these works, addressing the limitations of sparse sequence models by focusing on continuous, high-resolution trajectory generation that captures the complete dynamics of human movement patterns.

Early data-driven *model-free* approaches, employing techniques such as gravity models, decision trees, and Markovian methods such as Mobility Markov Chains (MMC) [23, 24, 25], face challenges in capturing sequential transitions between locations. In contrast, generative machine learning models have gained traction for their ability to learn from data without external input parameters [26]. Such generative models have been explored for human mobility modeling using various data representations, including grids in tabular format [27], image-like trajectory modeling [28], and sequential grid format [29]. However, these models exhibit limitations in generating trajectories that accurately capture geospatial complexity.

One promising generative model direction relies on GPS sampling. Due to mismatches, LSTM-based trajectory generation models cannot provide continuous motion of human mobility in diverse geospatial areas [30]. A similar continuous GPS trajectory generation method with a U-Net neural network using Diffusion models is proposed in [31]. Sampling GPS locations without geospatial constraints diminishes the quality of synthetic trajectories. They rarely incorporate geospatial constraints and road network topologies into their training pipeline, resulting in physically impossible trajectories that ignore natural barriers and transportation infrastructure limitations when applied to new environments.

Another generative direction in ML for mobility trajectories employs two-stage training methods [32, 19, 33]. A recent GAN-based model that relies on road links with the A^* path-finding algorithm between regions using two-stage GANs has shown promise in identifying diverse paths for specified OD pairs [19]. It falls short in generating truly synthetic trajectories, as the generation process requires initiating with an OD pair from testing trajectories. Another two-stage approach generates mobility trajectories using dual VAEs conditioned on user-level mobility characteristics learned in the second VAE model [33]. This approach is limited by its requirement for learning user-driver behaviors across multiple trips, necessitating substantial data per individual. In contrast, our MobilityGPT model learns generalist human mobility patterns without requiring personalization, making it more adaptable and efficient with available data.

Finally, sequential GAN methods have also been employed for mobility modeling with different variants [14, 29]. These earlier models generate data in sequences but struggle to capture geospatial complexity and add significant computational complexity, as discussed later in our ablation studies VI. Unlike previous methods, MobilityGPT leverages the autoregressive GPT model for trajectory generation. This enables a comprehensive exploration of spatial and temporal distributions with high-quality outputs by incorporating road links into sequence modeling.

The GPT models, developed by [34], transformed the field of Natural Language Processing (NLP) with the transformer’s capabilities. GPT, with its pre-training on diverse language data, demonstrated remarkable proficiency in understanding and generating coherent sequences to support various NLP tasks. The transformer-based GPT concept has been applied to diverse applications, from vision [35], music [36] to network data [37]. In addition, researchers have made great progress

on various fine-tuning models for enhancing the GPT model capabilities [38, 39, 40]. However, these fine-tuning models mainly rely on human feedback that evaluates the quality of the generated sequences. This study proposes a geospatially-aware GPT model, MobilityGPT, with an innovative fine-tuning approach for generating synthetic mobility trajectories.

III. PRELIMINARIES

A. Discretization with Map-Matching

Let $D(V, E)$ represent the road network as a weighted digraph, where the set of nodes V corresponds to the intersection of the road, the set of edges E to the roads and the weights representing the metrics of the link, such as the length of the link or traffic volume. There are two sets of trajectories: GPS trajectories and link trajectories.

1) GPS trajectories: A sequence of GPS coordinates with l number of samples $\mathbf{x} \in \mathcal{G} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_l\}$ forms a GPS trajectory that reflects the continuous motion of the object.

2) Link trajectories: Given the s number of vehicles on the road network, each vehicle travels between origins and destinations using an ordered link path generating a user travel path known as a *micro-graph* $\Phi \in D$. Every link trajectory has n number of links $\phi \in \Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ and $\Phi \subset E$. The set of trajectories is the corpus of all link trajectories with s users $\Phi \in \Lambda = \{\Phi_1, \Phi_2, \dots, \Phi_s\}$.

3) Map-matching: GPS coordinates provide an approximation of a device’s location by utilizing satellite broadcast information. Nonetheless, these coordinates may not consistently reflect the precise travel path, as they are susceptible to intrinsic and environmental errors, including satellite geometry, signal blockage, or tree cover [41]. Consequently, the recorded GPS locations might not align with a specific link on the road network. To address this, map-matching is employed to generate an ordered set of road network links that accurately describes the user’s trajectory, taking into account the road network $D(V, E)$ and GPS points [42].

Addressing location sequences in the continuous domain poses challenges. Consequently, conventional human mobility modeling methods often resort to the discretized grid domain. However, relying on a grid-like discrete representation does not eliminate the risk of geospatial mismatching. For instance, random sampling from a sparse road network grid may lead to nonsensical locations in generated trajectories. In response to this limitation, our MobilityGPT model adopts a strategy of discretizing locations into road segments, enhancing the quality of generated trajectories.

B. Problem Definition

In an urban setting, a generative model is tasked with learning the intrinsic mobility characteristics from a provided set of link trajectories, $\Phi \in \Lambda = \{\Phi_1, \Phi_2, \dots, \Phi_s\}$. The objective is to model human mobility using a sequential learning approach and generate synthetic trajectories that resemble training data across key utility metrics.

The sequential nature of mobility trajectories encapsulates human routing behavior, where the successive road segments influence subsequent link choice. The generative

model aims to learn the conditional probability distribution $P(\phi_{k+1}|\phi_k, \phi_{k-1}, \dots, \phi_{k-m})$, capturing the sequential dependencies that characterize human mobility patterns. The generation process involves producing synthetic trajectories that maintain both network connectivity constraints and realistic routing behaviors observed in the training data.

IV. MOBILITY MODELING WITH GENERATIVE TRANSFORMERS

This section describes the components of the proposed MobilityGPT for generating synthetic high-fidelity trajectories. MobilityGPT achieves this through a multi-objective training pipeline. Figure 2 illustrates the framework of our MobilityGPT mechanism, which can be divided into two parts: pretraining for learning sequence characteristics and fine-tuning for enhancing the quality of generated trajectories. The pretraining stage (Step 1 in Figure 2) learns the mobility modeling on the road-link sequences using the gravity of training data and RCM. The fine-tuning stage (Steps 2 and 3 in Figure 2) aims to improve the similarity of synthetic trajectories to the training trajectory samples in terms of travel length.

A. Generative Transformers

Mobility trajectories are the sequential records of locations or movements over time, capturing the spatiotemporal patterns of an entity's mobility. Their characteristics are crucial in human mobility modeling, providing insights into individual or collective movement behaviors.

Recent breakthroughs in NLP have proven that transformer architectures are remarkably effective in processing word sequences [43]. Trajectory sequences, representing the spatiotemporal movements of individuals or objects, share similarities with sentences across four key dimensions: sequential dependencies, spatial relationships, contextual embeddings, and variable-length sequences. First, trajectory sequences exhibit sequential dependencies similar to sentences, where the order of locations matters. Additionally, trajectories involve spatial relationships between locations, just as words in a sentence convey semantic relationships. Furthermore, trajectory sequences, like sentences, can have variable lengths. Finally, trajectories benefit from contextual embeddings that consider the entire sequence, similar to the contextual understanding of words in sentences.

Given these strong similarities with trajectories and word sequences, two key elements of transformers, the self-attention mechanism, and autoregressive generation, empower MobilityGPT to handle sequential data efficiently. The self-attention mechanism captures the long-range dependencies and weighs the importance of each location in the context of the entire trajectory, enabling the modeling of sequential dependencies. Besides, autoregressive sequence generation involves predicting one element at a time based on the context of preceding elements, contributing to the model's ability to generate coherent and contextually relevant trajectories.

The self-attention mechanism computes a weighted sum of values V based on attention scores A assigned to each element in the input sequence:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) \cdot V \quad (1)$$

where Q , K , and V represent the query, key, and value matrices. The softmax operation normalizes the attention scores, and d_k is the dimensionality of the key vectors. The weighted sum is then used as the output of the self-attention mechanism.

In an autoregressive nature of transformers, the probability of predicting the next element y_t in the sequence given the context $Y_{<t}$ is computed using the chain rule of probability:

$$P(y_t|Y_{<t}) = P(y_1) \cdot P(y_2|y_1) \cdot \dots \cdot P(y_t|y_{t-1}) \quad (2)$$

By formulating human mobility trajectories as a sequence generation task with road links from $D(V, E)$, transformer-based models can be applied to generate realistic and contextually relevant synthetic trajectories given the real trajectories. The tokenization of road links, combined with the self-attention mechanism, enables the model to capture spatial dependencies and generate coherent and meaningful sequences of road links. Next, we explain our unique tokenizer modeling.

B. Tokenizer Modeling

Tokenization, the breakdown of sequences into smaller units or tokens, is a key step in preparing raw data for transformer models. The influence of tokenization extends to trajectory sequences, where tokens represent locations or spatiotemporal points.

In MobilityGPT, tokens symbolize locations in terms of road links, and the contextual embeddings derived from tokenization contribute to generating embeddings that consider the holistic context of the entire trajectory.

To structure the input data into a format suitable for autoregressive modeling, each trajectory is treated as a discrete sequence of location-based tokens. For each trajectory, we incorporate an "end of trajectory" token, denoted as $\langle \text{EOT} \rangle$, which is appended to each trajectory during training. Given a trajectory sequence $\phi \in \Phi = \{\phi_1, \phi_2, \dots, \phi_n\}$ in terms of road links, where ϕ_i denotes the i -th token in the trajectory. The tokenization process is defined as follows:

$$\Phi_{\text{tokenized}} = \{\phi_1, \phi_2, \dots, \phi_n, \langle \text{EOT} \rangle\} \quad (3)$$

This inclusion of $\langle \text{EOT} \rangle$ serves a dual purpose. First, it acts as a sentinel token to signify the termination of a trajectory. Formally, if $\phi_i = \langle \text{EOT} \rangle$, then trajectory $i + 1$ begins. This structural sequence tokenization enables the model to distinguish between successive locations, thus avoiding positional leakage across trajectories during batch processing. This prevents the model from interpreting the start of a new trajectory in a batch as a continuation of the previous one, maintaining clear trajectory boundaries during both training and generation. Second, the proposed tokenization supports the generation of diverse and randomized trajectory sequences by clearly marking sequence endpoints, which facilitates autoregressive sampling without conflating distinct OD pairs. As $\langle \text{EOT} \rangle$ is part of




Fig. 2. Overview. The proposed MobilityGPT framework employs Generative Pretrained Transformer (GPT) architectures and self-adaptive reinforcement learning (RL) fine-tuning methods within a human mobility-aware training pipeline for synthetic trajectory generation. The process involves three steps: 1) Pretraining the GPT model with a gravity model and road connectivity matrix, 2) Constructing a trajectory-driven reward model, and 3) Fine-tuning the MobilityGPT pre-trained model using RL policy optimization methods with the trained reward model providing valuable feedback.

the model’s vocabulary, MobilityGPT assigns it a probability at every decoding step. This allows trajectories to terminate at variable positions, enabling the generation of diverse and randomized geospatial sequences while preserving structural coherence with real-world mobility patterns.

Moreover, this design abstracts away from raw GPS coordinates, allowing the model to learn high-level movement patterns through a discrete vocabulary of meaningful road link identifiers. This abstraction mitigates the impact of noise from GPS sparsity or irregular sampling rates and grounds the representation in topological connectivity rather than continuous space.

Utilizing this tokenization strategy is integral to our model’s ability to capture spatial patterns, differentiate between trajectories, and generate realistic sequences with varied ODs. This formulation contributes to the robustness and versatility of our trajectory generation approach.

C. Map Understanding with Gravity Model and Road Connectivity Matrix

Transformer-based models capture sequential information within long data sequences. However, human mobility modeling presents distinct challenges, necessitating a comprehensive understanding of geographical attributes and constraints. For instance, the sequence of links does not contain hidden attributes of human mobility, such as regional mobility flows (traveling from eastbound of city to west), traffic rules (must right-turn, no U-turn), and trip length. To address these limitations, we integrate two structured priors from classical mobility modeling—the gravity model and the RCM—into the training process of the transformer.

While both models are well-established in the transportation literature [44, 45], our work recontextualizes them as inductive biases within a deep generative model. Specifically, the gravity model is repurposed as a data sampling prior to expose the

transformer to representative and balanced regional flows, while the RCM is injected into the final decision layer of the model to constrain predictions to physically feasible road transitions. These strategies collectively guide the transformer to learn geospatial structure without requiring explicit auxiliary data such as maps or network graphs.

1) Gravity model. The area is discretized into regions $r \in \mathcal{R}$, with $\text{RegionW}(r_x)$ calculated as the number of trajectories starting or ending in the region r_x . The gravity value $\text{Gravity}(r_x, r_y)$ quantifies traffic flow between origin r_x and destination r_y regions as:

$$\text{Gravity}(r_x, r_y) = \frac{\text{RegionW}(r_x) \cdot \text{RegionW}(r_y)}{d^2(r_x, r_y)} \quad (4)$$

where $d^2(r_x, r_y)$ is the squared Euclidean distance between regions. This classical formulation has historically been used to model inter-regional travel demand. In our case, we adopt it as a training-time sampling prior to ensure that the model observes a diverse and flow-consistent distribution of OD pairs. During training, as shown in Figure 2, a gravity-aware sampler selects trajectory sequences $\Phi = \phi_1, \phi_2, \dots, \phi_n$ based on $\text{Gravity}(r_x, r_y)$, where $\phi_1 \in r_x$ and $\phi_n \in r_y$. This strategy reduces sampling bias toward overrepresented regions and encourages the model to internalize global mobility trends.

2) Road connectivity matrix. The clarity of road connectivity is often diminished when using GPS trajectories compared to traditional 2D maps. This is because, in some instances, roads that appear connected on a map might not be so in reality due to physical barriers or design differences, such as one-way streets and pedestrian zones. To incorporate real-world structural constraints into sequence generation, we propose the RCM, a sparse binary matrix that reflects the physical adjacency of road links.

Let Φ denote the set of training trajectories in terms of road




Fig. 3. The scheme of fine-tuning dataset generation for Reinforcement Learning from Trajectory Feedback.

links. The RCM is defined as:

$$RC_{\phi_x, \phi_y} = \begin{cases} 1 & \text{if } (\phi_x, \phi_y) \text{ is a consecutive pair in } \Phi, \\ & \text{or} \\ & \text{if } \phi_x \text{ or } \phi_y \text{ is in the last row/column,} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

RCM effectively represents the connectivity between different road links. If two links, ϕ_x and ϕ_y , are directly connected in the real world, their corresponding element in RCM is set to 1; otherwise, it is set to 0. This structure reflects only plausible transitions observed in the data. To enforce these structural constraints consistently, we apply the RCM to the transformer's output logits during both training and generation, ensuring that the model only learns and predicts physically plausible road transitions. Additionally, one row and one column, both filled with 1s, are added to the matrix to facilitate the generation of $\langle \text{EOT} \rangle$. Specifically, the last linear layer of the transformer, denoted as F , is adjusted by the RCM prior to the softmax activation function:

$$F' = F \cdot RC \quad (6)$$

$$\text{Softmax}(F') = \frac{e^{F'}}{\sum e^{F'}} \quad (7)$$

this integration is visualized in Figure 2, showing how the RCM modifies the output of the transformer's last linear layer. By multiplying F with the RCM, we ensure that only the connections between actually connected roads are considered. This process filters out any inaccuracies in road connectivity that might have been introduced by GPS data, thus enabling the model to have a more accurate understanding of the actual road network.

D. Reinforcement learning from trajectory feedback (RLTF)

Fine-tuning is essential for sequence generative models [46]. Pre-trained transformer architectures provide a foundation, but fine-tuning customizes models to specific objectives, such as similarity of trajectories, enabling them to learn complex patterns and nuances. This process is particularly vital in applications like human mobility, where capturing intricate movement patterns is crucial. Building on this, we propose a novel approach to construct a trajectory-aware preference

dataset that combines established fine-tuning strategies with a new preference feedback dataset - all without requiring human labeling. This enables the model to self-improve based on its output feedback, minimizing human intervention and addressing scalability and bias challenges [47].

Reinforcement learning has proven to be a powerful approach for capturing the long-term dependencies and spatiotemporal characteristics inherent in human mobility modeling, owing to the complexity of this task [48]. The proposed approach, incorporating our novel preference dataset with reinforcement learning, termed Reinforcement Learning from Trajectory Feedback (RLTF), in an automated feedback process, involves several key steps. First, a reward dataset is created by pairing prompts (partial trajectories) with compilations (complete trajectories generated by the model) and assigning scores based on the similarity of the compilation lengths to the original trajectory lengths (see Figure 3). This reward dataset is then utilized to train a reward model U through supervised learning. Subsequently, the trained reward model U is employed to fine-tune the MobilityGPT model through a Proximal Policy Optimization (PPO) policy π , effectively leveraging reinforcement learning without requiring human feedback.

The reward model U is a crucial component of our RLTF framework. We adopted the same lightweight and efficient model architecture as MobilityGPT itself for the reward model. This design choice ensures computational efficiency and a seamless integration with the main model, allowing the reward signal to be accurately modeled without introducing significant overhead. By using a shared architecture, we can effectively capture the complexities of the reward function while keeping the overall system resource-friendly.

Preference dataset construction for RLTF. Instead of relying on human feedback in fine-tuning of MobilityGPT, we propose using quantitative metrics to construct the preference dataset. Given a partial trajectory prompt, we generate two complete trajectories and assess their similarity to the actual trajectory by assigning a trajectory similarity score based on the relative lengths of the generated and reference trajectories from the training set (See Step 2 in Figure 2). This score measures how well the generated trajectories match the characteristics, particularly length, of the training data, serving as a proxy for



Fig. 4. Example of a preference dataset, showing a sample input alongside two completed trajectory pairs. Candidate 1’s trajectory demonstrates greater similarity to the Baseline trajectory compared to Candidate 2.

human preference without requiring human labeling.

The process, as described in Figure 3, begins with a full trajectory $\Phi = \phi_1, \phi_2, \dots, \phi_n$ of length $L(\Phi)$. The first m road-links $\Phi'_0 = \phi_1, \phi_2, \dots, \phi_m$ serve as a prompt for MobilityGPT to predict two possible continuations: $\Phi'_1 = \phi_{m+1}, \dots, \phi_k$ and $\Phi'_2 = \phi_{m+1}, \dots, \phi_z$. These are concatenated with Φ'_0 to form two complete trajectories, evaluated using $\Gamma_1 = |L(\Phi) - L(\Phi'_0 + \Phi'_1)|$ and $\Gamma_2 = |L(\Phi) - L(\Phi'_0 + \Phi'_2)|$. The trajectory with lower Γ is labeled Φ_{chosen} , and the other Φ_{rejected} , forming a preference dataset for RLTF. By leveraging this proximity-based metric, we eliminate the need for human labeling, streamlining the dataset creation process and potentially mitigating biases introduced by human subjectivity. MobilityGPT can be generalized to different fine-tuning strategies with the proposed preference dataset. In Figure 4, we exemplified one step of the preference dataset creation schema with an actual sample trajectory using the Porto dataset [19].

Reward model training. A Reward (preference) Model U is trained using the proposed preference dataset to predict the probability that a given trajectory is chosen over a rejected alternative (See Step 2 in Figure 2). The loss function is formulated as follows:

$$-\log(\sigma(U(\Phi_{\text{chosen}}) - U(\Phi_{\text{rejected}}))) \quad (8)$$

The goal of the loss function is to maximize the probability that the chosen trajectory has a higher score, thereby encouraging the model to assign a larger difference between Φ_{chosen} and Φ_{rejected} . As the training progresses, the model becomes more capable of recognizing the characteristics of trajectories that align with the preferred trip length.

Fine-tuning with reinforcement learning. Given the pre-trained MobilityGPT model (Step 1) and supervised trained reward model (Step 2), we train an RL policy that optimizes the MobilityGPT model for generating higher-quality synthetic trajectories as depicted in Step 3 of Figure 2. Since policy learning relies on the reward that policy receives from the environment, the reward model provides feedback for RL policy optimization. We followed a general policy optimization method as in [49]. The RL agent learns a fine-tuning policy

π , which is optimized with the well-known Proximal Policy Optimization (PPO) method [50]. The reward of policy π consists of two terms: reward model logits given the prompt Φ'_p and compilation Φ'_c pair received from the MobilityGPT as $U(\Phi'_p, \Phi'_c)$ and policy shift constraint in terms of the (Kullback–Leibler) KL divergence given the baseline policy π . The full reward can be represented as follows:

$$R(\Phi'_p, \Phi'_c) = U(\Phi'_p, \Phi'_c) - \beta \log \left(\frac{\pi^{\text{PPO}}_{\theta'}(\Phi'_c | \Phi'_p)}{\pi^{\text{base}}_{\theta'}(\Phi'_c | \Phi'_p)} \right). \quad (9)$$

It is worth emphasizing that the neural network architecture of the reward model U can be simplified to enhance efficiency and interpretability. Moreover, there is the option to incorporate an online fine-tuning stage, where the policy π and the reward model U are periodically re-trained. However, in our experimental assessments, we did not delve into the performance of the retraining process. This aspect and the formal exploration of privacy-preserving ML models are left for future research endeavors.

V. EXPERIMENTS

MobilityGPT’s effectiveness is validated through extensive experiments on real-world datasets against recent benchmarks. We performed all the experiments on a single NVIDIA TITAN RTX GPU with 24GB memory. The source code of MobilityGPT is publicly available. The preprint version of this paper is available on arXiv at [51].¹

A. Experimental setup

We have used two datasets of GPS trajectories from taxi drivers in Porto, Portugal, and Beijing, China. The Porto dataset has a sampling period of 15 seconds, whereas the Beijing dataset has a sampling period of 1 minute. Both datasets were preprocessed and map-matched following the protocol described by [19]. Specifically, the authors filtered out abnormal trajectories by removing those shorter than five road links or containing loops. After preprocessing, the Porto dataset comprises 695,085 unique trips across 11,095 road links, while the Beijing dataset includes 956,070 trajectories spanning 40,306 road links.

We chose these two datasets for two primary reasons. First, they are widely recognized and frequently used as benchmark datasets in the field of human mobility modeling, particularly in recent works such as TS-TrajGen [19] and DiffTraj [31]. This established use allows for a direct and fair comparison of our model’s performance against state-of-the-art benchmarks. Second, the datasets represent distinct geographic regions and mobility patterns. By validating our model on data from different continents with varying urban structures, we aim to demonstrate its robustness and ability to accurately model diverse human mobility characteristics, irrespective of geographical location or a specific urban context.

¹<https://github.com/ammahydr/MobilityGPT>

B. Benchmark models

We performed experiments on original implementation hyperparameters for all benchmark models using their publicly available source codes. We experimented with similar structured generative models with our proposed MobilityGPT model, including discrete and continuous structure formats.

- **Random Walk:** This baseline randomly selects a starting point on the road network and performs a random walk guided by visitation counts of the road segments (locations). It operates solely on the road graph structure without learning from the data distribution.
- **Mobility Markov Chain (MMC)** [25]: This model treats road segments as discrete states and constructs a transition probability matrix to capture first-order dependencies between segments. The next segment is sampled based on the learned transition probabilities.
- **SeqGAN** [14]: SeqGAN is a state-of-the-art sequential generative model that trains trajectories in discrete sequences with RL policy gradient methods using Monte Carlo sampling.
- **LSTM-TrajGAN** [30]: This is a tabular LSTM model by learning the normalized GPS location divergences with respect to the center of the dataset.
- **TS-TrajGen** [19]: TS-TrajGen is a method based on GANs that incorporates a modified A^* path search algorithm. This adjustment ensures spatial continuity in the generated data. Additionally, the method includes topological constraints to align the generated trajectories with existing road networks, ensuring road matching. TS-TrajGen requires fixed OD pairs as inputs. We use OD pairs from the real training set as inputs during inference.
- **DiffTraj** [31]: is the most recent and promising generative model for mobility trajectories. DiffTraj employs a diffusion model in the U-Net structure for generating GPS trajectories.

C. Evaluation metrics

To evaluate the performance of the generative models, we require robust comparison metrics that capture a broad spectrum of mobility characteristics and effectively quantify the quality of the generated trajectories relative to real ones. The selected metrics assess whether the synthetic trajectories preserve both local transition patterns—i.e., movements between adjacent road segments—and global route structures that reflect realistic and rational human mobility behavior in urban environments. In this study, we adopted four widely used evaluation metrics, along with our proposed gravity and connectivity utility metrics, as introduced in prior works [52, 53]. These metrics have become standard tools for assessing the utility of synthetic mobility data generation methods [54, 31]. Given that MobilityGPT introduces randomness and can alter trajectory patterns, direct one-to-one trajectory comparisons are not applicable in such generative settings [55].

Query Error: The query error, mainly used for evaluating data synthesis algorithms, is a popular metric for synthetic data generation models. We use spatial counting queries in the form

of “the number of trajectories passing through a certain road-link”. Given road network $D(V, E)$, we uniformly sampled 500 road links from the road network. Then, the normalized absolute difference between the number of real and synthetic trajectories passing through each of these links is computed by the following:

$$QE(f(\Sigma)) = \frac{|f(\Psi) - f(\Sigma)|}{\max\{f(\Psi), s\}}, \quad (10)$$

where $f(\Psi)$ and $f(\Sigma)$ stand for the query outcome from the original and synthetic trajectories, respectively, and s is sanity bound for mitigating the effect of the extremely small selective queries. We specified the sanity bound s as 1% of the users.

Apart from the query error, the main utility metrics inspect the distributional similarities for synthetic trajectories with respect to real trajectories that train the generative model. Jensen-Shannon divergence (JSD) is a well-known similarity metric mainly used for measuring the similarity of two probability distributions [56]. We employ JSD for several human mobility characteristics.

We can look at overall distributions of the generated trajectories Λ_{gen} and real trajectories Λ_{real} , and JSD can provide a summary statistic for a pair of distributions. These are the main distributions we extract from the Λ_{gen} and Λ_{real} :

- **OD:** This metric evaluates the origin-destination similarity between two datasets for the overall characteristics preserved in terms of OD links.
- **Trip Length:** This metric considers the travel distances of trajectories by using the total length of link trajectories.
- **Radius:** In human mobility patterns, the spatial range of daily movements is an important metric. We examined the user’s radius of gyration within the controlled road network area.
- **Gravity:** The gravity summary metric examines the impact of the gravity sampling model that was introduced for MobilityGPT at high-level mobility patterns with the same JSD method.

In addition to divergence metrics, we introduce a new metric that measures link trajectory connectivity as a percentage. The connectivity metric quantifies the percentage of trajectories that are fully connected from origin to destination, owing to the RCM that samples the next road link conditioned on the connected road links, thereby capturing the underlying road network topology and connectivity patterns. We employed the connectivity metric in our ablation studies to assess the impact of various components on link trajectory connectivity.

D. Model Details

MobilityGPT utilizes a minimal version of the standard decoder-only transformer architecture [57], built upon the minimal GPT implementation from [58]. Table I summarizes the important model and training parameters employed in this work. Since the Porto and Beijing datasets exhibit different trajectory lengths and characteristics, we utilize distinct block sizes for training and maximum trajectory lengths for synthetic trajectory generation. However, we employ the same set of hyperparameters across both datasets for all training stages, including the reward model and PPO training.

TABLE I
MODEL AND TRAINING PARAMETERS UTILIZED IN MOBILITYGPT

Parameter	Value
N. layers	6
N. heads	4
N embeddings	64
Batch size	64
Learning rate	1e-5
Training size	80%
Test size	20%
Training steps	3000
Block size (Porto)	300
Block size (Beijing)	60
Max. Traj. length (Porto)	278
Max. Traj. length (Beijing)	60

E. Comparison with benchmark generative methods

Table II summarizes the performance of MobilityGPT and benchmark models. We sampled 5000 random trajectories from the test set and generated same number of synthetic trajectories. The results show the significant advantage of MobilityGPT over existing state-of-the-art models for trajectory generation. Across both Porto-Taxi and BJ-Taxi, MobilityGPT consistently outperforms all baseline models, achieving the best or second-best performance on multiple utility metrics.

The results in Table II demonstrate that MobilityGPT consistently outperforms baseline models across both the Porto and Beijing datasets and all utility metrics. Rule-based methods such as Random Walk and MMC show relatively high divergence scores, indicating poor alignment with real-world distributions. While deep generative models like SeqGAN and DiffTraj achieve competitive results on some metrics, they lack consistency across all aspects of trajectory modeling. In contrast, MobilityGPT achieves the lowest JS divergence across all categories on both datasets—for instance, reducing OD distribution divergence to 0.114 (Porto) and 0.099 (Beijing), outperforming the next best scores of 0.116 and 0.105, respectively. This performance reflects the combined impact of MobilityGPT’s autoregressive transformer architecture and its gravity-based sampling strategy, which better captures global mobility patterns. By grounding generation in structural road connectivity and learned spatio-temporal dynamics, MobilityGPT shows strong potential for applications in transportation, urban planning, and mobility analytics.

The box plots in Figure 5 illustrate the performance of MobilityGPT and several other models across two key metrics: the number of points per trajectory and inter-point distance. Analysis of the Number of Points per Trajectory reveals that MobilityGPT’s performance is consistently strong. On the Porto dataset (Figure 5(a)), its median of 185.00 is very close to the Baseline’s 176.00, demonstrating that MobilityGPT effectively replicates the trajectory length characteristics of this dataset. Similarly, for the Beijing dataset (Figure 5(b)), MobilityGPT’s median of 84.00 is remarkably close to the Baseline’s 89.00. This indicates that MobilityGPT successfully replicates the fundamental characteristic of trajectory length from the original datasets, a crucial aspect of generating realistic, synthetic mobility data. In contrast, models like TS-TrajGen, MMC, and Random Walk show significantly different distributions,

suggesting they generate trajectories that are either excessively long or follow a different statistical pattern than the baseline.

Regarding Inter-point Distance, MobilityGPT also demonstrates a high degree of fidelity to the baseline data. For the Porto dataset (Figure 5(c)), the results are highly aligned, with MobilityGPT’s mean of 20.77 meters and median of 10.01 meters nearly identical to the Baseline’s 20.59 meters and 10.06 meters, respectively. On the Beijing dataset (Figure 5(d)), MobilityGPT’s mean of 61.98 meters and median of 26.18 meters are highly comparable to the Baseline’s mean of 61.20 meters and median of 25.21 meters. These results underscore that MobilityGPT not only generates trajectories of the correct length but also accurately captures the realistic step distances between consecutive points. This is a significant finding, as models like DiffTraj and SeqGAN exhibit a much wider variance or different median values, suggesting they may struggle to accurately represent the typical movement patterns found in real-world data. Overall, the consistent alignment of MobilityGPT’s metrics with the Baseline across both datasets confirms its effectiveness in generating high-fidelity, synthetic trajectory data.

F. Ablation study


As MobilityGPT incorporates three key enhancements to the standard GPT model for capturing human mobility patterns — gravity model, RCM, and RLTF — we conducted ablation experiments to evaluate their individual and combined contributions. Results, presented in Table IIIa and IIIb, show that removing the gravity model and RCM leads to noticeable performance declines. The gravity model plays a critical role in modeling trip lengths, offering substantial improvements even when used in isolation. Without the RCM, MobilityGPT struggles to maintain spatial connectivity in the generated trajectories. By conditioning the next-link prediction on actual road connections, RCM effectively enforces geospatial constraints. Furthermore, fine-tuning MobilityGPT with RLTF improves synthetic trajectory quality—particularly trip length—since the preference dataset is optimized for that objective. These results highlight how each component contributes uniquely to different aspects of mobility modeling, and together, they deliver comprehensive improvements across all evaluation metrics.

We compare our proposed fine-tuning approach, RLTF, with two alternative methods: Supervised Fine-Tuning (SFT) and Direct Preference Optimization (DPO)[59]. Their performance is summarized in Table IV. The results show that RLTF is more effective in capturing the complexities of human mobility and outperforms the other fine-tuning strategies. While all methods perform similarly in terms of OD prediction, RLTF achieves significantly lower error in modeling trip length. This improvement suggests that RLTF better captures the nuanced relationship between OD pairs and their associated trip lengths. Its superior performance can be attributed to its use of reward signals that explicitly incorporate spatial distribution and trip length objectives. In contrast, SFT and DPO primarily optimize for input-output matching, which may overlook these deeper patterns in mobility behavior.


TABLE II

COMPARISON OF THE UTILITY METRICS WITH BENCHMARK STUDIES. BOLD AND UNDERLINED RESULTS SHOW THE BEST AND SECOND-BEST RESULTS, RESPECTIVELY. FOR ALL METRICS, LOWER VALUES INDICATE BETTER PERFORMANCE. WHILE QUERY ERROR RELIES ON NORMALIZED ERROR RATES FOR MOST VISITED PLACES, OTHER METRICS USE THE JENSEN-SHANNON DISTRIBUTIONAL DIVERGENCE STATISTIC.


Methods	Porto-Taxi					BJ-Taxi				
	Query E.	OD	Trip L.	Radius	Gravity	Query E.	OD	Trip L.	Radius	Gravity
Random Walk	1.729	0.375	0.739	0.393	0.307	0.105	0.167	0.473	0.187	0.133
MMC	1.409	0.210	0.742	0.393	0.307	0.212	0.167	0.550	0.465	0.291
SeqGAN	0.138	0.116	0.220	0.191	0.326	0.063	0.296	0.162	0.275	0.120
LSTM-TrajGAN	0.234	0.254	0.435	0.425	0.277	0.078	0.182	0.398	0.365	0.219
TS-TrajGen	0.309	0.120	0.161	<u>0.115</u>	<u>0.248</u>	0.277	<u>0.105</u>	0.570	0.485	0.336
DiffTraj	0.267	0.177	<u>0.126</u>	0.123	0.252	0.074	0.146	0.158	<u>0.123</u>	<u>0.114</u>
MobilityGPT	<u>0.144</u>	0.114	0.124	0.107	0.225	<u>0.066</u>	0.099	0.123	0.102	0.105




(a) Distribution of number of points per trajectory for the Porto-Taxi dataset.



(b) Distribution of number of points per trajectory for the BJ-Taxi dataset.



(c) Distribution of inter-point distance for the Porto-Taxi dataset.



(d) Distribution of inter-point distance for the BJ-Taxi dataset.

Fig. 5. Comparison of different trajectory generation models on the Beijing and Porto datasets. The plots show the distribution of the number of points per trajectory and the inter-point distance.

G. Effect of temperature on MobilityGPT inferencing

The temperature parameter in LLMs, including MobilityGPT, influences the randomness in generating outputs. At lower temperatures, the model behavior is more deterministic, often selecting the most likely next token. As the temperature

increases, the model introduces greater randomness, leading to a diverse array of outputs.

Interestingly, as illustrated in Figure 7, MobilityGPT demonstrates a clear trend: its performance tends to increase with higher temperature settings. This observation suggests that




Fig. 6. Visual representation of the geospatial density of different methods. MobilityGPT exhibits the highest similarity to the original trajectories.

TABLE III

ABLATION STUDY ON PORTO (A) AND BEIJING (B) DATASETS EXAMINING THE IMPACT OF DIFFERENT COMPONENTS: RCM, GRAVITY MODEL, AND RLTF. LOWER VALUES ARE BETTER FOR OD AND TRIP LENGTH, WHILE HIGHER VALUES ARE BETTER FOR CONNECTIVITY.

RCM	Gravity Model	RLTF	OD	Trip L.	Connectivity
✗	✗	✗	0.123	0.154	0.87
✗	✓	✗	0.122	0.146	0.83
✓	✗	✗	0.120	0.147	1.0
✓	✓	✗	0.114	0.140	1.0
✓	✓	✓	0.114	0.124	1.0

(a)

RCM	Gravity Model	RLTF	OD	Trip L.	Connectivity
✗	✗	✗	0.102	0.201	0.330
✗	✓	✗	0.103	0.151	0.367
✓	✗	✗	0.093	0.145	1.0
✓	✓	✗	0.099	0.131	1.0
✓	✓	✓	0.099	0.123	1.0

(b)

TABLE IV

ABLATION STUDY ON DIFFERENT FINE-TUNING STRATEGIES FOR MOBILITYGPT. THE METRICS USE THE JENSEN-SHANNON DISTRIBUTIONAL DIVERGENCE STATISTIC.

SFT	DPO	RLTF	OD	Trip Length
✓			0.116	0.171
	✓		0.115	0.163
		✓	0.114	0.124

mobility data generation can benefit from the randomness of GPT-like models.

H. Tokenizer Modeling

In the development of the MobilityGPT tokenizer model, we explored the use of specific tokens to mark the start (<BOT>) and end (<EOT>) of trajectories. Our goal was to assess how these tokens affect the model's performance in human mobility modeling, particularly for sequential trajectory data. Our findings reveal a key insight: using a single token to indicate a trajectory boundary—whether <BOT> or <EOT>—significantly




Fig. 7. Performance of MobilityGPT under different temperature settings. All metrics use the Jensen-Shannon distributional divergence statistic.

improves performance compared to using separate start and end tokens. This result indicates that emphasizing trajectory boundaries with a single flag is more beneficial than marking both beginning and end with separate tokens for our mobility modeling task. The effectiveness of a single-token approach highlights the importance of task-specific tokenization strategies in sequential data generation with transformer models. An ablation study, detailed in Table V, compares single- and dual-token approaches using various metrics. This study demonstrates that a single-token strategy outperforms the dual-token method in the MobilityGPT workflow, underscoring the advantages of customized and simplified tokenization.

TABLE V

ABLATION STUDY ON THE EFFECT OF SINGLE-TOKEN VS DUAL-TOKEN STRATEGIES. ALL METRICS USE THE JENSEN-SHANNON DIVERGENCE.

<BOT>	<EOT>	OD	Trip Length	Radius	Gravity
✓	✓	0.110	0.151	0.128	0.235
✗	✓	0.114	0.124	0.107	0.225

I. Computational Demand

Evaluating computational demand is essential to ensure the practical feasibility, scalability, and environmental sustainability of generative models. To ensure fairness, all benchmark models

and MobilityGPT were trained on the same hardware setup, NVIDIA TITAN RTX GPU with 24GB memory. MobilityGPT achieves an effective balance between model complexity and efficiency—it is more compact than some of the larger transformer-based models while being more sophisticated than lightweight models like LSTM-based approaches. Regarding training time, MobilityGPT is significantly faster than the more complex baselines, DiffTraj and TS-TrajGen, yet it maintains far superior performance compared to the simplest models.

TABLE VI
COMPUTATIONAL PERFORMANCE OF DIFFERENT MODELS VS
MOBILITYGPT

Model	N_param		Total running time	
	BJ	Porto	BJ	Porto
SeqGAN	2.6M	800	~ 12hr	~ 5hr
LSTM-TrajGAN	100k	100k	~ 30mins	~ 1hr
TS-TrajGen	23.2M	7.6M	> 200hr	> 200hr
DiffTraj	15.7M	15.7M	~ 48hr	~ 60hr
MobilityGPT	5.5M	5.5M	~ 5hr	~ 5hr

Overall, MobilityGPT strikes a strong balance between model capacity and computational efficiency. It achieves state-of-the-art performance without the heavy resource requirements of larger models and avoids the accuracy trade-offs seen in lightweight alternatives. This balance makes MobilityGPT a practical and scalable solution for real-world applications where both predictive performance and computational feasibility are critical.

VI. DISCUSSION

While MobilityGPT achieves state-of-the-art results in generating realistic trajectories that respect road network constraints, there are notable limitations. First, the current model does not provide formal privacy guarantees for the synthetic trajectories, which limits its immediate applicability in privacy-sensitive domains such as ride-sharing or urban planning. Addressing this will require integrating differential privacy mechanisms or other formal privacy-preserving techniques into the generative process.

Second, MobilityGPT generalizes well across diverse urban environments (e.g., Porto and Beijing) when trained and evaluated within the same city. However, its ability to transfer across regions remains limited. A core reason for this is the region-specific nature of the tokenizer, where tokens correspond to individual road segments. Since urban areas vary significantly in road network topology and density, a tokenizer trained on one city cannot be directly applied to another. This restricts the model’s out-of-the-box generalizability and highlights the need for future work on region-agnostic tokenization or transferable representations across cities.

Lastly, while our current RLTF strategy with trajectory preferences derived from a similarity-based baseline enables scalable and automated optimization, it does not incorporate direct human feedback. Although human-in-the-loop signals could enhance adaptability to evolving mobility trends, collecting

such feedback at scale poses practical challenges. We consider this a promising direction for future research, particularly for applications where user alignment and behavioral nuance are critical.

VII. CONCLUSION

This paper presents a GPT-based generative method for modeling human mobility characteristics. The proposed mechanism considers a multi-objective learning method with pre-training and fine-tuning stages. MobilityGPT pre-training leverages the understanding of human mobility using the strong self-attention structures by conditioning the trajectory sampling with the gravity model and the sequence sampling with the RCM. To further enhance the MobilityGPT capabilities for generating similar trajectories to real trajectories, we propose a fine-tuning strategy that does not require a human assessment. The novel fine-tuning method trains a reward using the trajectory trip length and employs the reward model as an input to the RL policy gradient method. The experiments validate the effectiveness of the proposed MobilityGPT approach for generating realistic mobility trajectories from real trajectory samples by comparing with state-of-the-art benchmark models.

Future directions include enhancing the model’s cross-regional generalization capabilities, incorporating robust privacy guarantees and fine-tuning strategies with human feedback, enabling broader applicability in real-world deployments.

REFERENCES

- [1] S. Uppoor, O. Trullols-Cruces, M. Fiore, and J. M. Barcelo-Ordinas, “Generation and analysis of a large-scale urban vehicular mobility dataset,” *IEEE Transactions on Mobile Computing*, vol. 13, no. 5, pp. 1061–1075, 2013.
- [2] S. Wang, Z. Bao, J. S. Culpepper, and G. Cong, “A survey on trajectory data management, analytics, and learning,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 2, pp. 1–36, 2021.
- [3] H. Zhou, K. Ma, S. Liang, X. Li, and X. Qu, “A unified longitudinal trajectory dataset for automated vehicle,” *Scientific Data*, vol. 11, no. 1, p. 1123, 2024.
- [4] J. Zhang, B. Feng, Y. Wu, P. Xu, R. Ke, and N. Dong, “The effect of human mobility and control measures on traffic safety during COVID-19 pandemic,” *PLoS one*, vol. 16, no. 3, p. e0243263, 2021.
- [5] T. Hu, S. Wang, B. She, M. Zhang, X. Huang, Y. Cui, J. Khuri, Y. Hu, X. Fu, X. Wang *et al.*, “Human mobility data in the COVID-19 pandemic: characteristics, applications, and challenges,” *International Journal of Digital Earth*, vol. 14, no. 9, pp. 1126–1147, 2021.
- [6] X. Kong, Q. Chen, M. Hou, H. Wang, and F. Xia, “Mobility trajectory generation: a survey,” *Artificial Intelligence Review*, vol. 56, no. Suppl 3, pp. 3057–3098, 2023.
- [7] D. Song, R. Tharmarasa, T. Kirubarajan, and X. N. Fernando, “Multi-vehicle tracking with road maps and car-following models,” *IEEE transactions on intelligent transportation systems*, vol. 19, no. 5, pp. 1375–1386, 2017.

- [8] R. Bonta, "California consumer privacy act (CCPA)," *Retrieved from State of California Department of Justice: <https://oag.ca.gov/privacy/ccpa>*, 2022.
- [9] European Parliament and Council, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," *Official Journal of the European Union*, Regulation L 119/1, 2016, accessed: April 09, 2025. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:02016R0679-20160504>
- [10] J. Wang, X. Kong, F. Xia, and L. Sun, "Urban human mobility: Data-driven modeling and prediction," *ACM SIGKDD explorations newsletter*, vol. 21, no. 1, pp. 1–19, 2019.
- [11] M. Luca, G. Barlacchi, B. Lepri, and L. Pappalardo, "A survey on deep learning for human mobility," *ACM Computing Surveys (CSUR)*, vol. 55, no. 1, pp. 1–44, 2021.
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [13] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [14] L. Yu, W. Zhang, J. Wang, and Y. Yu, "Seqgan: Sequence generative adversarial nets with policy gradient," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 31, no. 1, 2017.
- [15] M. Musleh, M. F. Mokbel, and S. Abbar, "Let's speak trajectories," in *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, 2022, pp. 1–4.
- [16] G. K. Zipf, "The P 1 P 2/D hypothesis: on the intercity movement of persons," *American sociological review*, vol. 11, no. 6, pp. 677–686, 1946.
- [17] J. J.-C. Ying, W.-C. Lee, T.-C. Weng, and V. S. Tseng, "Semantic trajectory mining for location prediction," in *Proceedings of the 19th ACM SIGSPATIAL international conference on advances in geographic information systems*, 2011, pp. 34–43.
- [18] Q. Long, Y. Yuan, and Y. Li, "A Universal Model for Human Mobility Prediction," in *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, ser. KDD '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 894–905. [Online]. Available: <https://doi.org/10.1145/3690624.3709236>
- [19] W. Jiang, W. X. Zhao, J. Wang, and J. Jiang, "Continuous trajectory generation based on two-stage GAN," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 4, 2023, pp. 4374–4382.
- [20] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using SUMO," in *2018 21st international conference on intelligent transportation systems (ITSC)*. Ieee, 2018, pp. 2575–2582.
- [21] M. Fellendorf and P. Vortisch, "Microscopic traffic flow simulator vissim," *Fundamentals of traffic simulation*, pp. 63–93, 2010.
- [22] S.-L. Hsu, E. Tung, J. Krumm, C. Shahabi, and K. Shafique, "Trajgpt: Controlled synthetic trajectory generation using a multitask transformer-based spatiotemporal model," in *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems*, 2024, pp. 362–371.
- [23] L. Pappalardo, S. Rinzivillo, and F. Simini, "Human mobility modelling: exploration and preferential return meet the gravity model," *Procedia Computer Science*, vol. 83, pp. 934–939, 2016.
- [24] S. Jiang, Y. Yang, S. Gupta, D. Veneziano, S. Athavale, and M. C. González, "The TimeGeo modeling framework for urban mobility without travel surveys," *Proceedings of the National Academy of Sciences*, vol. 113, no. 37, pp. E5370–E5378, 2016.
- [25] S. Gams, M.-O. Killijian, and M. N. del Prado Cortez, "Next place prediction using mobility markov chains," in *Proceedings of the first workshop on measurement, privacy, and mobility*, 2012, pp. 1–6.
- [26] M. Baratchi, N. Meratnia, P. J. Havinga, A. K. Skidmore, and B. A. Toxopeus, "A hierarchical hidden semi-markov model for modeling mobility data," in *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*, 2014, pp. 401–412.
- [27] K. Ouyang, R. Shokri, D. S. Rosenblum, and W. Yang, "A non-parametric generative model for human trajectories," in *IJCAI*, vol. 18, 2018, pp. 3812–3817.
- [28] C. Cao and M. Li, "Generating mobility trajectories with retained data utility," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 2610–2620.
- [29] J. Feng, Z. Yang, F. Xu, H. Yu, M. Wang, and Y. Li, "Learning to simulate human mobility," in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 3426–3433.
- [30] J. Rao, S. Gao, Y. Kang, and Q. Huang, "LSTM-TrajGAN: A deep learning approach to trajectory privacy protection," *arXiv preprint arXiv:2006.10521*, 2020.
- [31] Y. Zhu, Y. Ye, S. Zhang, X. Zhao, and J. Yu, "Difftraj: Generating GPS trajectory with diffusion probabilistic model," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [32] X. Wang, X. Liu, Z. Lu, and H. Yang, "Large scale GPS trajectory generation using map based on two stage GAN," *Journal of Data Science*, vol. 19, no. 1, pp. 126–141, 2021.
- [33] Q. Long, H. Wang, T. Li, L. Huang, K. Wang, Q. Wu, G. Li, Y. Liang, L. Yu, and Y. Li, "Practical synthetic human trajectories generation based on variational point processes," in *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023, pp. 4561–4571.
- [34] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever


- et al.*, “Improving language understanding by generative pre-training,” 2018.
- [35] M. Chen, A. Radford, R. Child, J. Wu, H. Jun, D. Luan, and I. Sutskever, “Generative pretraining from pixels,” in *International conference on machine learning*. PMLR, 2020, pp. 1691–1703.
- [36] B. Banar and S. Colton, “A systematic evaluation of GPT-2-based music generation,” in *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*. Springer, 2022, pp. 19–35.
- [37] D. K. Kholgh and P. Kostakos, “PAC-GPT: A novel approach to generating synthetic network traffic with GPT-3,” *IEEE Access*, 2023.
- [38] Z. Zhang, X. Han, Z. Liu, X. Jiang, M. Sun, and Q. Liu, “ERNIE: Enhanced language representation with informative entities,” *arXiv preprint arXiv:1905.07129*, 2019.
- [39] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, “Language models are few-shot learners,” *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [40] N. Lambert, L. Castricato, L. von Werra, and A. Havrilla, “Illustrating Reinforcement Learning from Human Feedback (RLHF),” *Hugging Face Blog*, 2022, <https://huggingface.co/blog/rlhf>.
- [41] M. Berber, A. Ustun, and M. Yetkin, “Comparison of accuracy of GPS techniques,” *Measurement*, vol. 45, no. 7, pp. 1742–1746, 2012.
- [42] M. A. Quddus, W. Y. Ochieng, and R. B. Noland, “Current map-matching algorithms for transport applications: State-of-the art and future research directions,” *Transportation research part c: Emerging technologies*, vol. 15, no. 5, pp. 312–328, 2007.
- [43] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong *et al.*, “A survey of large language models,” *arXiv preprint arXiv:2303.18223*, 2023.
- [44] O. Cabanas-Tirapu, L. Danús, E. Moro, M. Sales-Pardo, and R. Guimerà, “Human mobility is well described by closed-form gravity-like models learned automatically from data,” *Nature Communications*, vol. 16, no. 1, p. 1336, 2025.
- [45] Y. Wang, F. Bu, X. Lv, Z. Hou, L. Bu, F. Meng, and Z. Wang, “Attention-based message passing and dynamic graph convolution for spatiotemporal data imputation,” *Scientific Reports*, vol. 13, no. 1, p. 6887, 2023.
- [46] S. Zhang, L. Dong, X. Li, S. Zhang, X. Sun, S. Wang, J. Li, R. Hu, T. Zhang, F. Wu *et al.*, “Instruction tuning for large language models: A survey,” *arXiv preprint arXiv:2308.10792*, 2023.
- [47] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, A. Ray *et al.*, “Training language models to follow instructions with human feedback,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 27 730–27 744, 2022.
- [48] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg *et al.*, “A generalist agent,” *arXiv preprint arXiv:2205.06175*, 2022.
- [49] D. M. Ziegler, N. Stiennon, J. Wu, T. B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving, “Fine-tuning language models from human preferences,” *arXiv preprint arXiv:1909.08593*, 2019.
- [50] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [51] A. Haydari, D. Chen, Z. Lai, M. Zhang, and C.-N. Chuah, “MobilityGPT: Enhanced human mobility modeling with a GPT model,” *arXiv preprint arXiv:2402.03264*, 2024.
- [52] S. Hasan, C. M. Schneider, S. V. Ukkusuri, and M. C. González, “Spatiotemporal patterns of urban human mobility,” *Journal of Statistical Physics*, vol. 151, pp. 304–318, 2013.
- [53] M. E. Gursoy, L. Liu, S. Truex, L. Yu, and W. Wei, “Utility-aware synthesis of differentially private and attack-resilient location traces,” in *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, 2018, pp. 196–211.
- [54] A. Haydari, C.-N. Chuah, M. Zhang, J. Macfarlane, and S. Peisert, “Differentially Private Map Matching for Mobility Trajectories,” in *Proceedings of the 38th Annual Computer Security Applications Conference*, 2022, pp. 293–303.
- [55] A. Kapp, J. Hansmeyer, and H. Mihaljević, “Generative models for synthetic urban mobility data: A systematic literature review,” *ACM Computing Surveys*, vol. 56, no. 4, pp. 1–37, 2023.
- [56] J. Lin, “Divergence measures based on the shannon entropy,” *IEEE Transactions on Information theory*, vol. 37, no. 1, pp. 145–151, 1991.
- [57] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” 2019.
- [58] A. Karpathy, “mingpt,” <https://github.com/karpathy/minGPT>, accessed: May 17th 2024.
- [59] R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, and C. Finn, “Direct preference optimization: Your language model is secretly a reward model,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.



Ammar Haydari Received the B.Sc. degree in Electronic Engineering from Uludag University, Bursa, Turkey, in 2014, M.S. degree in Electrical Engineering from the University of South Florida, Tampa, FL, in 2019, and Ph.D. degree from the Department of Electrical and Computer Engineering at the University of California, Davis. His research interests include intelligent transportation systems, deep reinforcement learning, cybersecurity, and machine learning.



Dongjie Chen received his Ph.D. degree in Electrical and Computer Engineering from the University of California, Davis. His research focuses on Deep Generative Models, specifically Generative Adversarial Networks (GANs), Large Language Models (LLMs), and Vision-Language Models (VLMs). He also explores the applications of these models in various domains, particularly in healthcare and transportation.



Zhengfeng Lai received his B.S. in Information Engineering from Zhejiang University in 2019, and his M.S. degree in Electrical and Computer Engineering from the University of California, Davis in 2021. He is currently pursuing his Ph.D. degree at the University of California, Davis. He has served as a reviewer for the IEEE Trans. of Image Processing, ICML, ICLR, NeurIPS, ICCV, CVPR, and ECCV.



Michael Zhang is currently a professor in the Civil and Environmental Engineering Department at University of California Davis. His research is in traffic operations and control, transportation network analysis and intelligent transportation systems. Professor Zhang received his BS degree in Civil Engineering from Tongji University, and MS and PhD degrees in Engineering from University of California Irvine. He is an Area Editor of the journal Network and Spatial Economics, and an Associate Editor of Transportation Research, Part B, Methodological, and

Transportation Science.



Chen-Nee Chuah received her B.S. in Electrical Engineering from Rutgers University, and her M. S. and Ph.D. in Electrical Engineering and Computer Sciences from the University of California, Berkeley. She is currently the Child Family Professor in Engineering in Electrical and Computer Engineering at the University of California, Davis. Her research interests lie in applying data analytics and AI/ML techniques to complex networked systems and applications, including cyberinfrastructure, online social platforms, healthcare, and intelligent transportation systems.

Chuah is a Fellow of the AAAS and IEEE, and an ACM Distinguished Scientist. She currently serves as an Associate Editor for the ACM Transactions on Computing for Healthcare and has previously served on the editorial board for IEEE/ACM Transactions on Networking, IEEE Transactions on Mobile Computing, and IEEE Internet of Things Journal.