

Revisiting Synthetic Human Trajectories: Imitative Generation and Benchmarks Beyond Datasaurus

Bangchao Deng
University of Macau
Macao SAR, China
yc37980@um.edu.mo

Xin Jing
University of Macau
Macao SAR, China
yc27431@um.edu.mo

Tianyue Yang
University of Macau
Macao SAR, China
mc35301@um.edu.mo

Bingqing Qu
BNU-HKBU United International
College, China
bingqingqu@uic.edu.cn

Dingqi Yang*
University of Macau
Macao SAR, China
dingqiyang@um.edu.mo

Philippe Cudre-Mauroux
University of Fribourg
Switzerland
philippe.cudre-mauroux@unifr.ch

ABSTRACT

Human trajectory data, which plays a crucial role in various applications such as crowd management and epidemic prevention, is challenging to obtain due to practical constraints and privacy concerns. In this context, synthetic human trajectory data is generated to simulate as close as possible to real-world human trajectories, often under summary statistics and distributional similarities. However, the complexity of human mobility patterns is oversimplified by these similarities (a.k.a. “Datasaurus”), resulting in intrinsic biases in both generative model design and benchmarks of the generated trajectories. Against this background, we propose MIRAGE, a huMan-Imitative tRAjectory GenERative model designed as a neural Temporal Point Process integrating an Exploration and Preferential Return model. It imitates the human decision-making process in trajectory generation, rather than fitting any specific statistical distributions as traditional methods do, thus avoiding the Datasaurus issue. Moreover, we also propose a comprehensive task-based evaluation protocol beyond Datasaurus to systematically benchmark trajectory generative models on four typical downstream tasks, integrating multiple techniques and evaluation metrics for each task, to comprehensively assess the ultimate utility of the generated trajectories. We conduct a thorough evaluation of MIRAGE on three real-world user trajectory datasets against a sizeable collection of baselines. Results show that compared to the best baselines, MIRAGE-generated trajectory data not only achieves the best statistical and distributional similarities with 59.0-67.7% improvement, but also yields the best performance in the task-based evaluation with 10.9-33.4% improvement. A series of ablation studies also validate the key design choices of MIRAGE. Our code and datasets are at <https://github.com/UM-Data-Intelligence-Lab/MIRAGE>.

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD’25, August 03–07, 2025, Toronto, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/XXXXXXX.XXXXXXX>

CCS CONCEPTS

• Information systems → Spatial-temporal systems.

KEYWORDS

Human Mobility, Trajectory, Generative Model, Mobility Simulation

ACM Reference Format:

Bangchao Deng, Xin Jing, Tianyue Yang, Bingqing Qu, Dingqi Yang, and Philippe Cudre-Mauroux. 2018. Revisiting Synthetic Human Trajectories: Imitative Generation and Benchmarks Beyond Datasaurus. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (KDD’25)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Human trajectory data is a key ingredient for a wide range of applications, including urban planning [67], traffic management [34], epidemic analysis [15], predictive policing [57], and crowd monitoring [7]. These applications heavily rely on the quality of human mobility models learnt from human trajectory data. However, acquiring large-scale real human trajectory data is often challenging due to practical constraints and privacy concerns. Therefore, synthetic human trajectory data has been widely used as an alternative, where generative models are learnt to generate artificial trajectories that closely resemble real-world human trajectories, making human trajectory data more readily available.

In the current literature, existing works on synthetic human trajectory generation mostly focus on the resemblance under summary statistics and distributional similarities [28]. For example, these works measure the resemblance in different aspects of mobility characteristics such as spatial distribution (e.g. G-Rank) [26], temporal distribution (e.g. stay duration) [38], OD flows (trips per OD pair) [63], and user mobility patterns (e.g. I-Rank and DailyLoc) [54], using divergence/distance metrics such as Kullback-Leibler divergence (KLD) [3], Jensen-Shannon divergence (JSD) [12, 37], earth mover’s distance (EMD) [2], Root Mean Squared Error (RMSE) [38]. These similarities often serve on one hand as part of the model design such as the model fitting objective [25], while on the other hand also as the benchmarks for evaluating trajectory generative models [23, 32, 37, 53]. However, while these similarities provide insights into the differences between real and generated data from various perspectives, they indeed oversimplify the complexity of human mobility patterns, resulting in intrinsic biases in

both generative model design and benchmarks of the generated trajectories. Specifically, datasets that are similar over a number of statistical properties may yield very different patterns, known as “Datasaurus” [5, 33]. In the context of trajectory data, it implies that statistically/distributionally similar trajectories may imply different mobility patterns, which thus leads to different performances in downstream tasks (as evidenced by our experiments in Section 5.3). Therefore, the resemblance under summary statistics and distributional similarities cannot fully reflect the ultimate utility of generated trajectories in supporting downstream tasks.

A few recent works started to consider the utility of downstream tasks, by either incorporating task-specific prior knowledge in the generative model design, such as constraining the next locations in a close neighborhood for traffic flow simulation [26], or benchmarking generative models on a heuristically designed downstream task with one specific technique to solve the task [12, 32, 65]. Nevertheless, these works are either limited by their task-specific design or lack a comprehensive view of utility benchmarks. In particular, heuristically designed downstream tasks may lead to unknown biases in the utility evaluation, as the performance of different techniques solving the same task often varies (as evidenced by our experiments in Appendix A) and heuristically choosing one technique as the benchmark is thus untrustworthy.

Against this background, we study in this paper the problem of task-agnostic human trajectory generation and its benchmarks beyond “Datasaurus”. Specifically, the task-agnostic generation should be independent of specific downstream tasks and imitate the human decision-making process in trajectory generation. Such human-imitative design needs to consider the unique characteristics of human trajectories. First, real-world human trajectories usually consist of sparse and irregularly observed presence events, where individuals voluntarily share their presence at semantic-enriched locations under their own preference, such as check-ins at Points of Interest (POIs¹) on social media [59], which differ from other periodic and regularly sampled mobility traces with only GPS coordinates such as taxi trajectories [6]. To handle such human trajectories, some works simply treat human trajectories as sequences without considering the temporal information [19, 23], while some others heuristically transform human trajectories to sequences of regularly observed events under a pre-defined and dataset-specific time interval (e.g., every 30 minutes [12] or 1-hour [3, 38]). However, these approaches are limited either in ignoring the important temporal information, or in incorporating a strong assumption of observing human trajectory in regular time intervals and heuristically interpolating missing and removing redundant events. Second, human decision-making on mobility choices has been widely recognized and evidenced to follow an Exploration and Preferential Return (EPR) model [25, 38, 44]. It presents individuals with a choice between two competing mechanisms: the exploration mechanism selecting previously unvisited locations and the preferential return mechanism to encourage returning to a previously visited location. Existing works mostly design mechanistic EPR-alike models reliant on simple statistic models under Markovian assumptions [3]. However, despite the advantage of

being interpretable by design, these statistic models have intrinsic limitations in accurately modeling the complex mobility patterns observed in real-world human trajectories.

To address the above issues, we propose MIRAGE, a huMan-Imitative tRAjectory GenErative model, which imitates the holistic human decision-making process in trajectory generation and does not explicitly fit any specific statistical distributions (while traditional EPR model does), thus avoiding the “Datasaurus” issue. Specifically, a human trajectory consisting of a sequence of stochastic presence events on continuous time is naturally a Temporal Point Process (TPP). Subsequently, we design MIRAGE as a neural TPP with intensity-free parameterization, benefiting from both the flexibility of the neural network encoding the trajectory history and the efficiency of the TPP in modeling continuous-time stochastic events [43]. To generate one event of a human trajectory, we first sample an event time from the neural TPP based on the encoded trajectory history, and then sample the activity category further conditioned on the sampled event time. Afterward, we design a neural EPR model, mimicking the human decision-making process to choose either exploring unvisited locations or returning to previously visited locations, conditioned on the sampled time and activity category. We also adopt a user variational autoencoder to imitate individual preferences in the trajectory generation process.

Meanwhile, we also propose a comprehensive task-based evaluation protocol to systematically benchmark trajectory generative models beyond summary statistics and distributional similarities (i.e. “Datasaurus”). The key idea is to evaluate whether the generated trajectories are similar to real trajectories in practically supporting different downstream tasks. To this end, we evaluate the performance of both real and generated mobility trajectories on a variety of typical downstream tasks using multiple techniques and evaluation metrics for each task, so as to average out the biases of individual techniques and metrics. We measure the paired performance discrepancy between the real and generated mobility trajectories using relative errors, such as Mean Absolute Percentage Error (MAPE) and Mean Squared Percentage Error (MSPE), which finally serve as benchmarks to assess the ultimate utility of the generated trajectories.

We summarize our contribution as follows:

- We reveal the limitations of existing human trajectory generative models in focusing on the summary statistics and distributional similarities between real and generated trajectories, which lead to intrinsic biases in both generative model design and benchmarks.
- We propose MIRAGE, a huMan-Imitative tRAjectory GenErative model designed as an intensity-free neural Temporal Point Process integrating a neural Exploration and Preferential Return model to imitate the human decision-making process in trajectory generation.
- We propose a comprehensive task-based evaluation protocol to systematically benchmark trajectory generative models on four typical downstream tasks, integrating multiple techniques and evaluation metrics for each task, to assess the ultimate utility of the generated trajectories.
- We conduct a thorough evaluation of MIRAGE on three real-world human trajectory datasets against a sizeable collection of state-of-the-art baselines. Results show that compared to the best

¹Given the context of human trajectories, we do not distinguish the two terms “locations” and “POIs” throughout this paper.

baselines, MIRAGE-generated trajectory data not only achieves the best statistical and distributional similarities with 59.0–67.7% improvement but also yields the best performance in the task-based evaluation with 10.9–33.4% improvement.

2 RELATED WORK

2.1 Mobility Trajectory Generation

Early works mostly model individual mobility with explicit physical meanings [16, 45] and generate synthetic trajectories under the distributions of key characteristics observed in real mobility patterns, such as trip lengths, start locations, or start times, etc. As a widely recognized mobility model, the Exploration and Preferential Return model (EPR) [44] unifies exploration and return mobility patterns by selecting new locations based on a random walk process for exploration and revisiting previously visited locations based on their frequency for preferential return. Subsequent studies further extend the EPR model by integrating sophisticated spatial or social information, such as mining the correlation between location capacity and social network sizes [1], incorporating a nested gravity model into the EPR model [39], linking mobility to social ties and studying the dynamics of spatial choices based on social behavior [48], and integrating the circadian propensity of human mobility and Markov-based models into the EPR model [25]. In addition, the EPR model is also shown to be universal in human behavior modeling in general, such as user activities in recommendation systems [36] and human behaviors in cyberspace [22]. However, these models, often reliant on heuristic statistical assumptions, have limitations in accurately modeling the complex mobility patterns observed in real-world trajectories.

Recently, deep learning generative models have been widely adopted for mobility trajectory generation. They can flexibly capture the complex spatiotemporal patterns encoded in real-world mobility trajectories without strong prior assumptions. For example, SeqGAN [62] is the pioneering work of sequence generation based on Generative Adversarial Networks (GAN); MoveSim [12] later incorporates information about physical distance, temporal periodicity, and historical transition matrix of location into a GAN framework; TrajGen [6] employs a CNN-based GAN to map mobility trajectories to images and to generate synthetic trajectory images, followed by a Seq2Seq model to output the synthetic trajectory; DeltaGAN [54] adopts a two-stage generative model to simulate human mobility trajectories, capturing fine-grained timestamps and effectively representing temporal irregularities; TS-TrajGen [26] combines the A* algorithm [17] with a GAN framework to generate continuous trajectories on urban road networks; SAVE [23] combines VAE and LSTM for mobility trajectory generation.

In addition, (neural) temporal point processes [8, 42, 64, 69] are also widely used to model the temporal dynamics of user behaviors. In the context of trajectory generation, VOLUNTEER [32] incorporates a two-layer VAE model with a temporal point process to capture the characteristics of human mobility; ActSTD [65] enhances the dynamic modeling of individual trajectories by utilizing neural ordinary equations in the continuous location domain; DSTPP [64] further models the complex spatiotemporal joint distributions using diffusion models. In this paper, beyond traditional generative deep learning models, we further design and integrate a neural

EPR model with neural TPPs to imitate the human decision-making process in trajectory generation.

2.2 Synthetic Mobility Trajectory Benchmarks

The benchmarks for synthetic mobility trajectories can be classified into two categories [28]. First, statistical and distributional similarities are the most widely used benchmarks, such as Kullback-Leibler divergence (KLD) [3], Jensen-Shannon divergence (JSD) [12, 37], earth mover’s distance (EMD) [2], Root Mean Squared Error (RMSE) [38], which are used to measure the similarities between real and generated trajectories in different aspects. For example, typical mobility statistics include the radius of gyration [12, 26, 65], the number of distinct locations visited per user per day [12, 54], I-Rank (frequency of visiting personal top locations) [12, 38], and the number of daily trips per user [38], trip lengths between consecutive trajectory points or between the origin and destination [12, 26, 53]; spatial distributions characterize the distribution of locations based on factors like visits per location (i.e. G-Rank) [12, 26, 37, 50, 53] or location popularity ranking [11]; temporal distributions characterize the number of trips per hour of the day [38], stay duration [12, 32, 37] and time intervals between check-ins [65]. However, while these similarity metrics provide insights into the differences between real and generated data from various perspectives, they cannot fully reflect the ultimate utility of generated trajectories in supporting downstream tasks.

Second, benchmarking on downstream tasks recently emerged as an evaluation scheme for synthetic mobility trajectories. These tasks include road map updating [6], next location prediction [26, 32], and spreading simulation [12, 53, 65]. However, these works often use a heuristically designed downstream task with one specific technique/algorithm to solve the task, which leads to unknown biases in the utility evaluation. As evidenced in our experiments in Appendix A, the performance of different techniques solving the same task varies; heuristically choosing the results of one technique as the benchmark is thus untrustworthy. Therefore, we propose a comprehensive task-based evaluation protocol to systematically benchmark synthetic mobility trajectories.

3 PRELIMINARIES

3.1 Problem Definition

Human Trajectory. A trajectory is defined as a time-ordered sequence $X = \{x_1, x_2, \dots, x_n\}$, where $x_i = (t_i, k_i, l_i)$ is a presence event defined as a tuple consisting of a timestamp t_i , and a semantic activity category k_i , and a location (POI) l_i .

Trajectory Generation. Given a real-world human trajectory dataset, the objective is to generate a new trajectory dataset while preserving the fidelity and utility of the original real-world dataset.

3.2 Neural Temporal Point Processes

3.2.1 Temporal Point Processes. A Temporal Point Process (TPP) is a stochastic process where its realization is a sequence of discrete events in time, represented as a sequence $\mathcal{T} = \{t_1, \dots, t_N\}$, which can be equivalently represented as a sequence of strictly positive inter-event times $\tau_{i+1} = (t_{i+1} - t_i) \in \mathbb{R}^+$. The conditional intensity $\lambda^*(t)$, which fully specifies the TPP, represents the instantaneous rate of arrival of new events at time t given the history of past




Figure 1: An overview of our MIRAGE with 1) a trajectory history encoder and 2) a probabilistic event decoder integrating a time decoder with intensity-free TPP, a category decoder, a location decoder with neural EPR, and a user VAE.

events $\mathcal{H}_t = \{t_j \in \mathcal{T} | t_j < t\}$, where $(*)$ is used as a shorthand for conditioning on the history. The conditional probability density function of time τ_{i+1} until the next event is computed by the integration of $\lambda^*(t)$ as follows:

$$p^*(\tau_{i+1}) = \lambda^*(t_i + \tau_{i+1}) \exp\left(-\int_0^{\tau_{i+1}} \lambda^*(t_i + s) ds\right) \quad (1)$$

3.2.2 TPP Parameterization. Traditional TPPs often specify a simple parametric intensity function capturing relatively simple patterns in event occurrences, such as the Hawkes process [18], which often leads to poor results because of its limited flexibility in modeling complex data. In this context, neural TPPs are developed to use neural networks to learn complex dependencies from the history of TPPs, and then approximate the intensity function by some parametric forms [43]. However, these intensity-based approaches cannot achieve flexibility, efficiency, and ease-to-use simultaneously [42]. Alternatively, an intensity-free learning model of neural TPPs is proposed to use a mixture of log-normal distributions to represent the conditional probability density function $p^*(\tau)$ directly, and thus bypass the cumbersome intensity function [42]. Due to its advantage of closed-form sampling and likelihood computation, the intensity-free approach has been widely adopted in learning neural TPPs [13, 54]. We also adopt this approach in our work.

4 MIRAGE

To imitate the human decision-making process in trajectory generation, we design MIRAGE as an intensity-free neural Temporal Point Process (TPP) integrating a neural Exploration and Preferential Return (EPR) model. In the following, we first present the overview of MIRAGE, followed by the details of individual components.

4.1 Overview

Figure 1 shows the overview of MIRAGE consisting of two components. To imitate the human decision-making process, a trajectory

history encoder learns from the sequences of events using a recurrent neural network. Afterward, based on the output hidden states of the trajectory history encoder, a probabilistic event decoder generates the probability distributions for the next event time, activity categories, and locations in a cascading manner, where the latter distribution depends on the samples drawn from the former distribution. Notably, the distribution of the next event time is modeled using an intensity-free TPP, while the distribution of the next location is modeled by a neural EPR model. In addition, we also use a Variational Autoencoder to model the distribution of individual preferences on trajectory (User VAE).

4.2 Trajectory History Encoder

Representing a human trajectory as a sequence of events $\{x_1, x_2, \dots, x_n\}$, the trajectory history encoder first encodes individual event $x_i = (t_i, k_i, l_i)$ to event embedding e_i , and then adopts a Gated Recurrent Unit (GRU) to encode the sequence of event embeddings. Specifically, for each event (t_i, k_i, l_i) , we first encode the timestamp t_i by converting it to the log inter-event time $\log(\tau_i) = \log(t_i - t_{i-1})$. This lossless conversion is to fit the formulation of the intensity-free TPP; however, the actual timestamp t_i is still useful in the decoding process, and we will discuss this point later. We also encode the category k_i and location l_i using two embedding layers E_k and E_l , respectively.

$$e_i^t = \log(\tau_i), \quad e_i^k = E^k k_i, \quad e_i^l = E^l l_i \quad (2)$$

The event embedding e_i is obtained by concatenating e_i^t, e_i^k and e_i^l :

$$e_i = [e_i^t; e_i^k; e_i^l] \quad (3)$$

After obtaining a sequence of event embeddings, we then encode it using a recurrent neural network of GRU:

$$h_i = g(h_{i-1}, e_i) \quad (4)$$

where g represents the recurrent updating function of GRU. The output hidden state h_i encodes all the trajectory history until t_i .

4.3 Probabilistic Event Decoder

Based on the output hidden state h_i , the probabilistic event decoder generates the probability distribution for the next inter-event time $p^*(\tau_{i+1})$, activity category $p^*(k_{i+1} | \tau_{i+1})$, and location $p^*(l_{i+1} | \tau_{i+1}, k_{i+1})$ in a cascading manner, where the latter distribution depends on the samples drawn from the former distributions. Note that here $(*)$ denotes the conditioning on the trajectory history, while the conditional probability emphasizes on conditioning on drawn samples.

4.3.1 Time decoder with intensity-free TPP. Following the intensity-free TPP [42], we use a mixture of log-normal distributions to characterize the conditional probability density function of inter-event time $p(\tau)$.

$$p(\tau | \omega, \mu, \sigma) = \sum_{m=1}^M \omega_m \frac{1}{\tau \sigma_m \sqrt{2\pi}} \exp\left(-\frac{(\log \tau - \mu_m)^2}{2\sigma_m^2}\right) \quad (5)$$

where M is the number of components of the mixture, ω_m denotes the weight of each component, μ_m and σ_m are the logarithmic mean and standard deviation of each component.

Moreover, beyond the intensity-free TPP, we have two further design considerations to better accommodate the human trajectories. First, the original intensity-free TPP is defined on inter-event times only (as mentioned in the trajectory history encoder), which ignores the actual timestamp t_i that could play a crucial role in human mobility modeling. For example, human trajectories exhibit diurnal rhythms, where few events are observed during the night on weekdays; in other words, if an event is observed at t_i at midnight, the next event will probably be in the morning of the next day. This implies that $p^*(\tau_{i+1})$ should be skewed to the sleep duration here. Therefore, we transform the timestamp t_i into an hour-in-week timestamp embedding t_i^{emb} where $1 \leq i \leq 168$ (i.e., one of the 168 hours in a week), and t_i^{emb} is used to further condition $p^*(\tau_{i+1})$. The hour-in-week granularity is selected to capture both daily and weekly dynamics as suggested by [59]. Second, human trajectories are intrinsically generated under individual preferences [32]. For example, user check-ins at POIs on social networks have been widely used for user preference modeling [61]. Therefore, we also define an individual preference embedding u^{emb} for each trajectory to condition $p^*(\tau_{i+1})$.

Then, we concatenate the hidden state of trajectory history h_i , the hour-in-week timestamp embedding t_i^{emb} , and the individual preference embedding u^{emb} as the overall time context for the conditional probability density function of inter-event time $p^*(\tau_{i+1})$:

$$c_i^\tau = [h_i; t_i^{emb}; u^{emb}] \quad (6)$$

We then compute the parameters of the distribution $p^*(\tau_{i+1})$ using an affine function of c_i^τ :

$$\begin{cases} \omega_i = \text{softmax}(V_\omega c_i + b_\omega) \\ \mu_i = V_\mu c_i + b_\mu \\ \sigma_i = \exp(V_\sigma c_i + b_\sigma) \end{cases} \quad (7)$$

where the softmax and exp transformations are used to enforce the constraints on the distribution parameters ($\sum_m \omega_m = 1$, $\omega_m \geq 0$, and $\sigma_m > 0$), and $\{V_\omega, b_\omega, V_\mu, b_\mu, V_\sigma, b_\sigma\}$ are learnable parameters.

Finally, we sample a next inter-event time τ_{i+1} according to the log-normal mixture $p^*(\tau_{i+1}) = p(\tau_{i+1} | \omega_i, \mu_i, \sigma_i)$, and also obtain the corresponding next event time $t_{i+1} = t_i + \tau_{i+1}$.

4.3.2 Category decoder. After obtaining the next inter-event time τ_{i+1} , we sample an activity category. In this step, we use a slightly different (category) context vector c_i^k by replacing the timestamp embedding t_i^{emb} in the time context vector c_i^τ with the embedding of the sampled next event time t_{i+1}^{emb} , as follows:

$$c_i^k = [h_i; t_{i+1}^{emb}; u^{emb}] \quad (8)$$

This is because the next activity category k_{i+1} is highly correlated with its corresponding event time t_{i+1} ; for example, the activity category in the noon time should probably be “food” (one of the nine categories on Foursquare [59]). Then, we compute a categorical distribution conditioned on the category context vector c_i^k :

$$p^*(k_{i+1} | \tau_{i+1}) = \text{softmax}(\text{MLP}_\phi(c_i^k)) \quad (9)$$

where MLP_ϕ refers to a multi-layer perception with parameters ϕ and softmax transforms its output to a categorical probability distribution. We finally sample a next event category k_{i+1} from $p^*(k_{i+1} | \tau_{i+1})$. Note that as the next event time $t_{i+1} = t_i + \tau_{i+1}$ can

be computed in closed-form from the sampled inter-event time τ_{i+1} , we keep using τ_{i+1} as the condition of $p^*(k_{i+1} | \tau_{i+1})$ for emphasizing the sampling dependences.

4.3.3 Location decoder with neural EPR. After obtaining the next event time t_{i+1} and category k_{i+1} , we sample the next location by designing a neural EPR model. Specifically, our location decoder follows a two-step design. First, we sample a binary decision on two competing modes: exploration (visiting a new location that is not in the trajectory history) or return (visiting a previously visited location in the trajectory history). Second, based on the sampled mode, we then sample a location from the corresponding candidate locations of the mode. We present the detailed design below.

Exploration/Return mode sampling. In this step, we learn to sample one mode based on an extended context vector c_i^l , by further adding the embedding of the sampled next event category k_{i+1}^{emb} to the context vector of category encoder c_i^k :

$$c_i^l = [h_i; t_{i+1}^{emb}; k_{i+1}^{emb}; u^{emb}] \quad (10)$$

Subsequently, we feed this context vector to an MLP followed by a softmax function, to output a distribution over the two modes of *explore* and *return*:

$$p^*(z | \tau_{i+1}, k_{i+1}) = \text{softmax}(\text{MLP}_\theta(c_i^l)) \quad (11)$$

Finally, we sample one mode z according to $p^*(z | \tau_{i+1}, k_{i+1})$. Note that for the first event in a trajectory sequence, the exploration mode is selected, because of the empty trajectory history.

Location sampling in the exploration mode. In the exploration mode $z = \text{explore}$, we compute a categorical distribution over all previously unvisited locations, which is obtained by feeding the location context vector c_i^l to an MLP followed by a softmax function:

$$p^*(l_{i+1} | \tau_{i+1}, k_{i+1}, \text{explore}) = \text{softmax}(\text{MLP}_\xi(c_i^l)), \quad (12)$$

where $l \notin \{x_j | x_j \in X, j < i + 1\}$

From this, we can sample a previously unvisited location l .

Location sampling in the return mode. In the return mode $z = \text{return}$, we consider the temporal periodicity of human behaviors, where individuals tend to revisit previously visited locations under a regular temporal distance, such as returning home daily (with a temporal distance of 24 hours) [44]. Therefore, the return mode is designed to imitate such revisiting behaviors. Specifically, instead of imposing specific constraints or predefined temporal periodicity as [9, 56], we use temporal distance embeddings to learn the probability distribution of returning to previously visited locations. We define the temporal distance under an hour granularity between the $(i + 1)$ -th and j -th events as $\Delta_{i+1,j}^t = |t_{i+1} - t_j| \in \mathbb{Z}_{\geq 0}$. Subsequently, for the next event time t_{i+1} , we obtain its temporal distances to all previous events as follows:

$$\Delta^t = [\Delta_{i+1,1}^t, \Delta_{i+1,2}^t, \dots, \Delta_{i+1,i}^t] \quad (13)$$

We then define a learnable embedding for each hourly time distance, which is fed to an MLP followed by softmax to output the categorical distribution of returning to the locations of the previous events:

$$p^*(l_{i+1} | \tau_{i+1}, \text{return}) = \text{softmax}(\text{MLP}_\eta([\Delta_{i+1,1}^{emb}, \Delta_{i+1,2}^{emb}, \dots, \Delta_{i+1,i}^{emb}])),$$

where $l \in \{x_j | x_j \in X, j < i + 1\}$

(14)

Table 1: Dataset statistics

	TKY	IST	NYC
#Sequences	8,890	14,380	17,682
#POIs	7,590	16,809	15,723
#Categories	9	9	9

Note that here $p^*(l_{i+1}|\tau_{i+1}, \text{return})$ is not conditioned on the sampled next event category k_{i+1} because the revisiting behavior depends only on the next event time. We then sample a previously visited location l according to $p^*(l_{i+1}|\tau_{i+1}, \text{return})$. Our experiments also reveal that this learnt probability indeed exhibits a clear daily return pattern over different time distances (see Figure 4).

In summary, our location decoder with neural EPR samples a next location l_{i+1} according to:

$$p^*(l_{i+1}|\tau_{i+1}, k_{i+1}) = p^*(\text{explore}|\tau_{i+1}, k_{i+1}) \cdot p^*(l_{i+1}|\tau_{i+1}, k_{i+1}, \text{explore}) + p^*(\text{return}|\tau_{i+1}, k_{i+1}) \cdot p^*(l_{i+1}|\tau_{i+1}, \text{return}) \quad (15)$$

4.3.4 User VAE. Additionally, we also train a Variational Autoencoder (VAE) to learn to generate novel individual preference embeddings on trajectories, to avoid exposing the learnt user preference embeddings from the real dataset in the generation process. In the sampling stage, the trained user VAE generates individual preference embedding from random noises.

4.4 Model Training and Trajectory Generation

4.4.1 Model Training. The training objective of MIRAGE is to minimize the Negative Log Likelihood (NLL) of all sequences of events in a human trajectory dataset $\mathcal{D}$.

$$\mathcal{L}(\mathcal{D}) = - \sum_{X \in \mathcal{D}} \sum_{x_i \in X} (\ln p^*(\tau_i) + \ln p^*(k_i|\tau_i) + \ln p^*(l_i|\tau_i, k_i)) \quad (16)$$

In the training process, the above conditional probability distribution is conditioned on the real data without sampling. After training this objective function, we then train the user VAE on the learnt individual preference embeddings. The complexity of our model is discussed in Appendix B.

4.4.2 Trajectory Generation. Our trajectory generation process is conducted via a series of sampling steps without any real data as input. Specifically, the user VAE first generates an individual preference embedding. Afterward, the trajectory is generated iteratively based on the encoded trajectory history as follows: 1) the time decoder samples an inter-event time τ_i according to $p^*(\tau_i)$; 2) the category decoder samples an activity category k_i conditioned on τ_i according to $p^*(k_i|\tau_i)$; and 3) the location decoder samples a location l_i conditioned on τ_i and k_i according to $p^*(l_i|\tau_i, k_i)$. Note that the initial hidden state h_0 is learnt during training, and the generation process terminates until a desired time length.

5 EXPERIMENTS

5.1 Experimental Settings

5.1.1 Dataset. We conduct extensive experiments on three user trajectory datasets collected from a location-based social network Foursquare [59, 60], in three respective cities Tokyo (TKY), Istanbul (IST), and New York City (NYC). Table 1 shows the dataset statistics.

5.1.2 Baselines. We consider the following state-of-the-art baselines of three categories: statistical models **Semi-Markov** [29] and **TimeGeo** [25]; neural TPP **RMTTP** [10], **ERTPP** [52], **THP** [69], **ActSTD** [65] and **LogNormMix** [42]; deep learning generative models **LSTM** [24], **SeqGAN** [62], **MoveSim** [12], **VOLUNTEER** [32] and **DiffTraj** [68]. The baseline details are in Appendix C.

5.1.3 Statistical and Distributional Similarity Metrics. We adopt five popular metrics [12, 37] to evaluate the resemblance between real and generated trajectories in different aspects. **Distance** measures the distance between successive locations in a trajectory. **Radius** of gyration is calculated as the root mean squared distance of all locations from the central one in a trajectory. **Interval** is computed as time intervals between successive events in a trajectory. **DailyLoc** computes the unique locations visited by users. **Category** computes the overall distribution of the POI categories. We use the Jensen-Shannon divergence (JSD) [14] as the similarity metric between the distributions of real and generated trajectories.

5.2 Task-Based Evaluation Protocol

We introduce our proposed task-based evaluation protocol to comprehensively benchmark trajectory generative models. Specifically, the ultimate utility of human trajectory generation is to support different downstream tasks in practice; subsequently, the benchmark objective is to evaluate whether the generated trajectories are similar to real trajectories when being used to conduct different downstream tasks. In the current literature, existing works all use heuristically designed downstream tasks with one specific technique to solve a task [6, 12, 26, 32, 65], and thus lack a comprehensive view of utility benchmarks. In particular, heuristically designed downstream tasks may lead to unknown biases in the utility evaluation, as the performance of different techniques solving the same task often varies (as evidenced by our experiments in Appendix A), and heuristically regarding the results of one technique as the benchmark is thus untrustworthy.

Our proposed evaluation protocol implements four typical tasks: location recommendation, next location prediction, semantic location labeling, and epidemic simulation, which model user trajectory data in four aspects, user preferences on locations, sequential mobility patterns, collective traffic patterns, and spatiotemporal contact patterns, respectively. For each task, we choose multiple state-of-the-art techniques to conduct experiments and report the results on multiple metrics, to average out the biases of individual techniques and metrics. We then measure the paired performance discrepancy between the real and generated trajectories using Mean Absolute Percentage Error (**MAPE**) and Mean Squared Percentage Error (**MSPE**), which serve as final benchmarks to assess the ultimate utility of the generated trajectories. In the following, we first present our dataset settings, followed by the details of each task.

5.2.1 Dataset Settings. Unlike some previous works [26, 32, 65] that use the generated trajectories to augment the real trajectories (combining generated and real data in certain proportions) and then perform the downstream tasks, we put one step forward to directly perform the tasks on the generated trajectories only, without exposing any real trajectories. Our dataset setting is more strict which completely avoids the leakage of real human trajectories. In

Table 2: Performance on Statistical and Distributional Similarity Metrics

Method	TKY					IST					NYC				
	Distance	Radius	Interval	DailyLoc	Category	Distance	Radius	Interval	DailyLoc	Category	Distance	Radius	Interval	DailyLoc	Category
Semi-Markov	0.6931	0.5840	0.1196	0.6931	0.0791	0.6931	0.5545	0.2041	0.6931	0.0295	0.6931	0.5097	0.2072	0.6930	0.0135
Time Geo	0.5307	0.5712	0.0467	0.4441	0.0140	0.5984	0.6339	0.0345	0.3644	0.0124	0.4437	0.5673	0.0387	0.1982	0.0144
RMTTP	0.1988	0.1724	0.4166	0.5778	0.0113	0.1864	0.2008	0.4151	0.5650	0.0322	0.1745	0.2288	0.4422	0.3728	0.0141
ERTPP	0.6807	0.4021	0.1249	0.6636	0.0187	0.6699	0.3601	0.2110	0.6762	0.0236	0.6783	0.4395	0.2127	0.6824	0.0169
THP	0.3604	0.0858	0.2290	0.6183	0.0815	0.3476	0.4975	0.1122	0.2414	0.0340	0.4239	0.5109	0.0738	0.6235	0.0544
ActSTD	0.2470	0.1921	0.0272	0.1931	0.0231	0.1710	0.1225	0.0098	0.1325	0.0231	0.1884	0.1861	0.0152	0.1360	0.0131
LogNormMix	0.3898	0.2585	0.0186	0.2114	0.0056	0.2567	0.1787	0.0299	0.1547	0.0152	0.2172	0.1751	0.0352	0.1120	0.0205
LSTM	0.2745	0.1983	0.0919	0.1581	0.0143	0.3640	0.2516	0.0758	0.1960	0.0671	0.2876	0.2425	0.0476	0.2501	0.0133
SeqGAN	0.3006	0.2616	0.0507	0.2406	0.0326	0.1548	0.1734	0.0584	0.1279	0.0125	0.3023	0.2897	0.1206	0.3102	0.0078
MoveSim	0.3623	0.3018	0.0619	0.2761	0.0136	0.4549	0.3176	0.1537	0.4076	0.0527	0.3454	0.3700	0.0967	0.2293	0.0527
VOLUNTEER	0.5098	0.3208	0.0273	0.2630	0.0184	0.2498	0.2623	0.0415	0.1031	0.0171	0.3167	0.2845	0.0443	0.1456	0.0201
DiffTraj	0.2209	0.2653	0.1388	0.2684	0.0769	0.1974	0.0713	0.2137	0.1513	0.0237	0.1589	0.1039	0.1949	0.2361	0.0074
MIRAGE	0.1295	0.0330	0.0037	0.0229	0.0038	0.0714	0.0375	0.0033	0.0100	0.0045	0.0622	0.0415	0.0050	0.0185	0.0024

Table 3: Performance in the Task-Based Evaluation on MAPE

Method	TKY				IST				NYC			
	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim
Semi-Markov	0.8335	0.9919	0.4784	0.6406	0.9165	0.9969	0.5179	0.5652	0.8410	0.9951	0.4472	2.5354
Time Geo	0.7487	1.7867	0.2725	0.8531	3.4468	3.4048	0.3962	0.8838	1.0636	3.0667	0.2691	0.8122
RMTTP	0.8694	0.8616	0.2888	0.9402	0.7237	0.8180	0.4412	0.8599	0.8879	0.9406	0.3786	0.8801
ERTPP	0.4367	0.8181	0.1878	0.5592	0.3605	0.7854	0.3596	0.7666	0.6810	0.9156	0.5158	3.4278
THP	1.4991	2.2344	0.3985	0.3496	0.5154	0.8491	0.4963	0.9836	3.2324	3.7909	0.3496	0.6705
ActSTD	0.5341	0.6174	0.3738	0.2366	0.4632	0.5734	0.4454	0.4095	0.6648	0.5953	0.4158	0.8923
LogNormMix	0.4549	0.7148	0.1822	0.1480	0.4374	0.8592	0.3862	0.1620	0.7031	0.8702	0.4146	0.7216
LSTM	0.4816	0.7230	0.1631	0.1490	0.4025	0.6398	0.4016	0.2745	0.6884	0.5246	0.2889	0.6809
SeqGAN	0.4460	0.4514	0.2909	0.2041	0.8808	0.6971	0.4172	0.2216	1.0619	0.5271	0.5199	2.3157
MoveSim	0.7371	0.3347	0.2791	0.1790	5.0527	0.5815	0.4978	0.5957	1.1737	0.7824	0.3368	2.2911
VOLUNTEER	0.7277	0.5392	0.3050	0.4577	0.4796	0.5417	0.3750	0.2299	0.6695	0.8459	0.4831	1.6571
DiffTraj	0.7649	1.0815	0.4555	0.6752	0.7345	1.1090	0.4840	0.8263	0.6622	1.9322	0.4059	0.7631
MIRAGE	0.3485	0.3173	0.1433	0.1390	0.2947	0.3046	0.2198	0.2170	0.4855	0.2299	0.0961	0.7609

our experiments, for each real trajectory dataset, we generate a synthetic dataset having the same number of trajectories as the real dataset using each trajectory generative model. We then perform downstream tasks on both real and generated datasets separately, to evaluate whether the two datasets encode the same amount of information for supporting different downstream tasks.

5.2.2 Location Recommendation Task. (LocRec) suggests new (previously unvisited) locations for users by modeling *user preferences on locations* from trajectory data [61]. Following the default setting of [66], we transform a trajectory dataset into a set of (user, location, visit_count) triplets and then split them into training/valid/test datasets under a ratio of 8:1:1. To discount the impact of the specific techniques and metrics, we consider five popular recommendation algorithms, i.e., BPR [40], DMF [55], LightGCN [20], MultiVAE [31], and NeuMF [21] (details in Appendix D), and report their performance on Mean Reciprocal Rank@N (MRR@N), Normalized Discounted Cumulative Gain@N (NDCG@N), hit@N (where N = 5 and 10). Finally, we compare the paired performance discrepancy between the real and generated trajectories (i.e., 30 paired results from five algorithms and six metrics each) using MAPE and MSPE.

5.2.3 Next Location Prediction Task. (NexLoc) forecasts a user's next location in the future by learning the *sequential mobility patterns* from historical user trajectories [35]. Specifically, for one

dataset, we chronologically split each trajectory into training/valid/test trajectories under a ratio of 8:1:1. We also consider five sequence prediction algorithms, i.e., FPMC [41], BERT4Rec [46], Caser [47], SRGNN [51], and SASRec [27] (details in Appendix E), and report their performance on MRR@N, NDCG@N, hit@N (where N = 5 and 10). Finally, we compare the paired performance discrepancy between the real and generated trajectories using MAPE and MSPE.

5.2.4 Semantic Location Labeling Task. (SemLoc) assigns a semantic label (i.e., activity category) to a location based on the *collective traffic pattern* of the location, extracted from users' trajectory data [58]. Specifically, for each location, we extract its weekly temporal traffic pattern with an hour granularity, resulting in a feature vector of size 168 where each entry represents the empirical probability of all users' visits to this location. For each trajectory dataset, we then split all POIs into training/valid/test datasets under a ratio of 8:1:1. As a classification problem in nature, we consider five typical classification algorithms, i.e., Decision Tree, Naive Bayes, K-Nearest Neighbors, Logistic Regression, and Support Vector Machine, and report their performance on Accuracy, F1-Micro, and F1-Macro scores. Finally, we compare the performance discrepancy between the real and generated trajectories using MAPE and MSPE.

5.2.5 Epidemic Simulation Task. (EpiSim) simulates the epidemic spreading over a contact network characterizing the *spatiotemporal*

Table 4: Ablation Study on Statistical and Distributional Similarity Metrics

Method	TKY					IST					NYC				
	Distance	Radius	Interval	DailyLoc	Category	Distance	Radius	Interval	DailyLoc	Category	Distance	Radius	Interval	DailyLoc	Category
MIRAGE-noTPP	0.1673	0.0448	0.1436	0.0465	0.0043	0.1550	0.1449	0.2742	0.0551	0.0080	0.2275	0.0987	0.2453	0.1817	0.0105
MIRAGE-noEPR	0.1468	0.0752	0.0183	0.1276	0.0108	0.0720	0.0429	0.0336	0.0221	0.0063	0.0920	0.0565	0.0374	0.0955	0.0044
MIRAGE-noIMI	0.1368	0.0334	0.0051	0.0238	0.0101	0.0749	0.0356	0.0047	0.0342	0.0084	0.0660	0.0416	0.0081	0.0194	0.0096
MIRAGE-TD	0.1339	0.0429	0.0037	0.0289	0.0048	0.0933	0.0368	0.0036	0.0149	0.0059	0.0692	0.0413	0.0064	0.0206	0.0033
MIRAGE	0.1295	0.0330	0.0037	0.0229	0.0038	0.0714	0.0375	0.0033	0.0100	0.0045	0.0622	0.0415	0.0050	0.0185	0.0024

Table 5: Ablation study in the Task-Based Evaluation on MAPE

Method	TKY				IST				NYC			
	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim
MIRAGE-noTPP	0.3688	0.3293	0.1600	0.5806	0.4458	0.4380	0.2323	0.3768	0.6654	0.5663	0.1153	1.4793
MIRAGE-noEPR	0.3725	0.4025	0.1483	0.1409	0.3620	0.5077	0.2427	0.2342	0.6243	0.5243	0.1045	0.9413
MIRAGE-noIMI	0.3479	0.3458	0.1820	0.1457	0.5026	0.3399	0.3283	0.2148	0.5437	0.2807	0.2808	0.8046
MIRAGE-TD	0.3557	0.3732	0.2336	0.1574	0.3526	0.3625	0.3232	0.2490	0.5003	0.2381	0.2517	0.8387
MIRAGE	0.3485	0.3173	0.1433	0.1390	0.2947	0.3046	0.2198	0.2170	0.4855	0.2299	0.0961	0.7609

contact patterns of user trajectories [49]. The contact network is extracted from a trajectory dataset as a dynamic graph of users, and a directed edge from a user p to a user q represents that p 's visit precedes q 's visit to the same location during a day, indicating a potential chance of epidemic spreading from p to q . Following the setting of recent works [12, 30, 65], we adopt the Susceptible–Exposed–Infected–Recovered (SEIR) model. We consider the simulation of COVID-19 using the parameters suggested by [12, 65] and influenza using the parameters suggested by [4] (details in Appendix F). We randomly select 50 individuals as exposed individuals, simulate the spread of the epidemic, and report the daily counts of Exposed, Infectious, and Recovered individuals as three metrics. We report the average results of 10 repeated simulations to discount the impact of the random selection of initially exposed individuals. Finally, we compare the performance discrepancy between the real and generated trajectories using MAPE and MSPE.

5.3 Statistical & Distributional Similarity

Table 2 shows the performance comparison of MIRAGE and baselines on the three datasets. We observe that MIRAGE consistently achieves the best performance with the lowest JSD, yielding 59.0%, 64.4%, and 67.7% improvement (reduction on JSD) over the best-performing baselines on TKY, IST, and NYC datasets, respectively. However, even though we selected four similarity metrics covering three different aspects of spatial (Distance and Radius), temporal (Interval), and user (DailyLoc) distributions, these metrics cannot fully reflect the ultimate utility of generated trajectories in supporting downstream tasks. For example, the consistent superiority of MIRAGE over baselines on these similarity metrics is still biased, because MIRAGE indeed underperforms some baselines in a few cases in our task-based evaluation, as we discuss below.

5.4 Task-Based Evaluation Performance

Table 3 shows the performance in our proposed task-based evaluation on MAPE (similar results on MSPE, shown in Appendix G). Each entry in this Table is the MAPE averaging over all metrics of all algorithms solving a task on a dataset, to average out the biases of individual techniques and metrics. For example, the MAPE of MIRAGE in the LocRec task on the TKY dataset is 0.3485, which

is computed from 30 paired (real and generated) results from five algorithms and six metrics used in the LocRec task.

We observe that MIRAGE achieves the best performance with the smallest MAPE in most cases. In general, compared to the best-performing baselines, our MIRAGE achieves 10.9%, 16.7%, and 33.4% improvement on TKY, IST, and NYC datasets, respectively. In addition, we also see that in the EpiSim task, MIRAGE achieves second and even fourth places, on IST and NYC datasets, respectively, which departs from the consistent superiority of MIRAGE over all baselines on the statistical and distributional similarity metrics.

5.5 Ablation Study

We conduct an ablation study on our proposed method MIRAGE, considering the following four variants. **MIRAGE-noTPP** is a variant of MIRAGE without the neural TPP, where we define a regression task to predict the next inter-event from the time context vector c_i^t . **MIRAGE-noEPR** is a variant of MIRAGE without neural EPR, where we directly sample a location from a categorical distribution conditioned on the location context vector c_i^l . **MIRAGE-noIMI** is a variant of MIRAGE uses t_i^{emb} and k_i^{emb} (rather than t_{i+1}^{emb} and k_{i+1}^{emb}) for sampling the next category and next POI. **MIRAGE-TD** is a variant of MIRAGE that uses temporal distances (rather than timestamps) in the exploration mode. Tables 4 and 5 show the results on similarities and task-based evaluation on MAPE, respectively (similar results on MSPE in Appendix G).

We observe that MIRAGE consistently outperforms MIRAGE-noTPP on both similarities and task-based evaluation by 63.8% and 30.0% (on average over tasks and datasets), respectively, showing the effectiveness of neural TPPs modeling the event stochasticity of human trajectories. Second, MIRAGE outperforms MIRAGE-noEPR in most cases, with an improvement of 50.3% and 17.8% on similarities and task-based evaluation, respectively, which verifies the usefulness of our neural EPR model. Third, MIRAGE outperforms MIRAGE-noIMI in general, by 24.8% and 18.1% on similarities and task-based evaluation, respectively, since the next location should depend on the next category and time, rather than the current ones. Finally, MIRAGE outperforms MIRAGE-TD in most cases, by 15.2% and 18.5% on similarities and task-based evaluation, respectively.




Figure 2: Returning Probability over a Week

Because timestamps offer more direct information for location sampling in the exploration mode compared to temporal distances. Our ablation studies systematically validate our key design choices.

In addition, to further show the utility of the EPR model, we plot the empirical returning probability of users over time, which is defined as the probability of a user returning to a location a certain period (temporal distance) after the user’s first presence at the location [16, 56]. Figure 2 shows the plots of both real and generated data on NYC. We see that the real trajectory exhibits strong periodicity, which can be well imitated by MIRAGE but not by MIRAGE-noEPR and baselines. The superiority of MIRAGE is also validated by the return probability over sequence lengths. Figure 3(a) shows the return probability of the whole trajectories of real and generated data. We selected top baselines SeqGAN, VOLUNTEER, and MoveSim for comparison. Figures 3(b), 3(c), and 3(d) show the detailed return probability w.r.t. the (percentage) length of trajectories of three datasets (100% corresponds to the whole trajectory). We see that our MIRAGE is more similar to the real data in terms of return trend across different (percentage) lengths. Additionally, Figure 4 shows the learnt return probability $p^*(l_{i+1}|\tau_{i+1}, return)$ over different time distances (hour granularities in a week) from each dataset. We see a clear daily return pattern, implying that our model can effectively capture the periodicity encoded in human trajectories.

6 CONCLUSIONS AND FUTURE WORKS

In this paper, by revisiting the existing human trajectory generative models, we identify their limitations in focusing on the summary




Figure 3: Returning Probability over Sequence Length




Figure 4: Return probability over different time distances

statistics and distributional similarities between real and generated trajectories, which could lead to intrinsic biases in both generative model design and benchmarks of the generated trajectories. Against this background, we propose MIRAGE, a huMAN-Imitative tRAjectory GenErative model designed as an intensity-free neural Temporal Point Process integrating a neural Exploration and Preferential Return model to imitate the human decision-making process in trajectory generation. Meanwhile, we also propose a comprehensive task-based evaluation protocol to systematically benchmark trajectory generative models on four typical downstream tasks, integrating multiple techniques and evaluation metrics for each task, to assess the ultimate utility of the generated trajectories. The evaluation results show that MIRAGE-generated trajectory data not only achieves the best statistical and distributional similarities with 59.0-67.7% improvement but also yields the best performance in the task-based evaluation with 10.9-33.4% improvement.

Our future work will investigate the relationship between task performance variations and generative model design.

ACKNOWLEDGMENTS

This project has received funding from the University of Macau (MYRG2022-00048-IOTSC), Taihu Innovation scheme, and the Science and Technology Development Fund, Macau SAR (0047/2022/A1, 001/2024/SKL). This work was performed in part at SICCC which is supported by SKL-IOTSC, University of Macau.

REFERENCES

- [1] Laura Alessandretti, Piotr Sapiezynski, Vedran Sekara, Sune Lehmann, and Andrea Baronchelli. 2018. Evidence for a conserved quantity in human mobility. *Nature human behaviour* 2, 7 (2018), 485–491.
- [2] Cuauhtemoc Anda, Sergio A Ordonez Medina, and Kay W Axhausen. 2021. Synthesizing digital twin travellers: Individual travel demand from aggregated mobile phone data. *Transportation Research Part C: Emerging Technologies* 128 (2021), 103118.
- [3] Vincent Bindschadler and Reza Shokri. 2016. Synthesizing plausible privacy-preserving location traces. In *S&P. IEEE*, 546–563.
- [4] Fred Brauer, Carlos Castillo-Chavez, Zhilan Feng, Fred Brauer, Carlos Castillo-Chavez, and Zhilan Feng. 2019. Models for Influenza. *Mathematical Models in Epidemiology* (2019), 311–350.
- [5] Alberto Cairo. 2016. Download the Datasaurus: Never trust summary statistics alone; always visualize your data. <http://www.thefunctionalart.com/2016/08/download-datasaurus-never-trust-summary.html>
- [6] Chu Cao and Mo Li. 2021. Generating mobility trajectories with retained Data Utility. In *KDD*. 2610–2620.
- [7] Chunhua Chen, Yuxin Yang, Hao Yuan, Longbiao Chen, Leye Wang, Bingqing Qu, and Dingqi Yang. 2024. Animating the Crowd Mirage: A WiFi-Positioning-Based Crowd Mobility Digital Twin for Smart Campuses. *IMWUT* 8, 4 (2024), 1–32.
- [8] Ricky TQ Chen, Brandon Amos, and Maximilian Nickel. 2020. Neural spatio-temporal point processes. *arXiv:2011.04583* (2020).
- [9] Bangchao Deng, Dingqi Yang, Bingqing Qu, Benjamin Fankhauser, and Philippe Cudré-Mauroux. 2023. Robust Location Prediction over Sparse Spatiotemporal Trajectory Data: Flashback to the Right Moment! *ACM Transactions on Intelligent Systems and Technology* 14, 5 (2023), 1–24.
- [10] Nan Du, Hanjun Dai, Rakshit Trivedi, Utkarsh Upadhyay, Manuel Gomez-Rodriguez, and Le Song. 2016. Recurrent marked temporal point processes: Embedding event history to vector. In *KDD*. 1555–1564.
- [11] Yuntao Du, Yujia Hu, Zhikun Zhang, Ziquan Fang, Lu Chen, Baihua Zheng, and Yunjun Gao. 2023. LDPTrace: Locally Differentially Private Trajectory Synthesis. *arXiv:2302.06180* (2023).
- [12] Jie Feng, Zeyu Yang, Fengli Xu, Haisu Yu, Mudan Wang, and Yong Li. 2020. Learning to simulate human mobility. In *KDD*. 3426–3433.
- [13] Leo Feng, Frederick Tung, Hossein Hajimirsadeghi, Yoshua Bengio, and Mohamed Osama Ahmed. 2023. Memory Efficient Neural Processes via Constant Memory Attention Block. (2023). *arXiv:2305.14567 [cs.LG]*
- [14] Bent Fuglede and Flemming Topsøe. 2004. Jensen-Shannon divergence and Hilbert space embedding. In *ISIT. IEEE*, 31.
- [15] Song Gao, Jinneng Rao, Yuhao Kang, Yunlei Liang, and Jake Kruse. 2020. Mapping county-level mobility pattern changes in the United States in response to COVID-19. *SIGSpatial Special* 12, 1 (2020), 16–26.
- [16] Marta C Gonzalez, Cesar A Hidalgo, and Albert-László Barabási. 2008. Understanding individual human mobility patterns. *nature* 453, 7196 (2008), 779–782.
- [17] Peter E Hart, Nils J Nilsson, and Bertram Raphael. 1968. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics* 4, 2 (1968), 100–107.
- [18] Alan G Hawkes. 1971. Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58, 1 (1971), 83–90.
- [19] Xi He, Graham Cormode, Ashwin Machanavajjhala, Cecilia Procopiuc, and Divesh Srivastava. 2015. DPT: differentially private trajectory synthesis using hierarchical reference systems. *Vldb* 8, 11 (2015), 1154–1165.
- [20] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. LightGCN: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [21] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. 173–182.
- [22] Tianran Hu, Yinglong Xia, and Jiebo Luo. 2019. To return or to explore: Modelling human mobility and dynamics in cyberspace. In *WWW*. 705–716.
- [23] Dou Huang, Xuan Song, Zipei Fan, Renhe Jiang, Ryosuke Shibasaki, Yu Zhang, Haizhong Wang, and Yugo Kato. 2019. A variational autoencoder based generative model of urban human mobility. In *MIPR. IEEE*, 425–430.
- [24] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv:1508.01991* (2015).
- [25] Shan Jiang, Yingxiang Yang, Siddharth Gupta, Daniele Veneziano, Shounak Athavale, and Marta C González. 2016. The TimeGeo modeling framework for urban mobility without travel surveys. *PNAS* 113, 37 (2016), E5370–E5378.
- [26] Wenjun Jiang, Wayne Xin Zhao, Jingyuan Wang, and Jiawei Jiang. 2023. Continuous Trajectory Generation Based on Two-Stage GAN. *arXiv:2301.07103* (2023).
- [27] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *ICDM. IEEE*, 197–206.
- [28] Alexandra Kapp, Julia Hansmeyer, and Helena Mihaljević. 2023. Generative Models for Synthetic Urban Mobility Data: A Systematic Literature Review. *Comput. Surveys* 56, 4 (2023), 1–37.
- [29] Vo S Korolyuk, SM Brodi, and AF Turbin. 1975. Semi-Markov processes and their applications. *Journal of Soviet Mathematics* 4, 3 (1975), 244–280.
- [30] Shengjie Lai, Nick W Ruktanonchai, Liangcai Zhou, Olivia Prosper, Wei Luo, Jessica R Floyd, Amy Wesolowski, Mauricio Santillana, Chi Zhang, Xiangjun Du, et al. 2020. Effect of non-pharmaceutical interventions to contain COVID-19 in China. *nature* 585, 7825 (2020), 410–413.
- [31] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *WWW*. 689–698.
- [32] Qingyue Long, Huandong Wang, Tong Li, Lisi Huang, Kun Wang, Qiong Wu, Guangyu Li, Yanping Liang, Li Yu, and Yong Li. 2023. Practical synthetic human trajectories generation based on variational point processes. In *KDD*. 4561–4571.
- [33] Justin Matejka and George Fitzmaurice. 2017. Same stats, different graphs: generating datasets with varied appearance and identical statistics through simulated annealing. In *CHI*. 1290–1294.
- [34] Diala Naboulsi, Marco Fiore, Stephane Ribot, and Razvan Stanica. 2016. Large-Scale Mobile Traffic Analysis: A Survey. *IEEE Communications Surveys and Tutorials* 18, 1 (2016), 124–161.
- [35] Anastasios Noulas, Salvatore Scellato, Neal Lathia, and Cecilia Mascolo. 2012. Mining user mobility features for next place prediction in location-based services. In *ICDM. IEEE*, 1038–1043.
- [36] Kensuke Onuma, Hanghang Tong, and Christos Faloutsos. 2009. Tangent: a novel ‘surprise me’ recommendation algorithm. In *KDD*. 657–666.
- [37] Kun Ouyang, Reza Shokri, David S Rosenblum, and Wenzhuo Yang. 2018. A non-parametric generative model for human trajectories. In *IJCAI*, Vol. 18. 3812–3817.
- [38] Luca Pappalardo and Filippo Simini. 2018. Data-driven generation of spatio-temporal routines in human mobility. *Data Mining and Knowledge Discovery* 32, 3 (2018), 787–829.
- [39] Luca Pappalardo, Filippo Simini, Salvatore Rinzivillo, Dino Pedreschi, Fosca Giannotti, and Albert-László Barabási. 2015. Returners and explorers dichotomy in human mobility. *Nature communications* 6, 1 (2015), 8166.
- [40] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv:1205.2618* (2012).
- [41] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized markov chains for next-basket recommendation. In *WWW*. 811–820.
- [42] Oleksandr Shchur, Marin Biloš, and Stephan Günnemann. 2019. Intensity-free learning of temporal point processes. *arXiv:1909.12127* (2019).
- [43] Oleksandr Shchur, Ali Caner Türkmen, Tim Januschowski, and Stephan Günnemann. 2021. Neural temporal point processes: A review. *arXiv:2104.03528* (2021).
- [44] Chaoming Song, Tal Koren, Pu Wang, and Albert-László Barabási. 2010. Modelling the scaling properties of human mobility. *Nature physics* 6, 10 (2010), 818–823.
- [45] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. 2010. Limits of predictability in human mobility. *Science* 327, 5968 (2010), 1018–1021.
- [46] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *CIKM*. 1441–1450.
- [47] Jiaxi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *WSDM*. 565–573.
- [48] Jameson L Toole, Carlos Herrera-Yañe, Christian M Schneider, and Marta C González. 2015. Coupling human mobility and social ties. *Journal of The Royal Society Interface* 12, 105 (2015), 20141128.
- [49] Rajat Verma, Takahiro Yabe, and Satish V Ukkusuri. 2021. Spatiotemporal contact density explains the disparity of COVID-19 spread in urban neighborhoods. *Scientific Reports* 11, 1 (2021), 10952.
- [50] Xingrui Wang, Xinyu Liu, Ziteng Lu, and Hanfang Yang. 2021. Large scale GPS trajectory generation using map based on two stage GAN. *Journal of Data Science* 19, 1 (2021), 126–141.
- [51] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. In *AAAI*, Vol. 33. 346–353.
- [52] Shuai Xiao, Junchi Yan, Xiaokang Yang, Hongyuan Zha, and Stephen Chu. 2017. Modeling the intensity function of point process via recurrent neural networks. In *AAAI*, Vol. 31.
- [53] Gang Xiong, Zhishuai Li, Meihua Zhao, Yu Zhang, Qinghai Miao, Yisheng Lv, and Fei-Yue Wang. 2023. TrajSGAN: A Semantic-Guiding Adversarial Network for Urban Trajectory Generation. *IEEE Transactions on Computational Social Systems* (2023).
- [54] Nan Xu, Loc Trinh, Sirisha Rambhatla, Zhen Zeng, Jiahao Chen, Samuel Assefa, and Yan Liu. 2021. Simulating continuous-time human mobility trajectories. In *ICLR*. 1–9.
- [55] Hong-Jian Xue, Xinyu Dai, Jianbing Zhang, Shujian Huang, and Jiajun Chen. 2017. Deep matrix factorization models for recommender systems. In *IJCAI*, Vol. 17. Melbourne, Australia, 3203–3209.
- [56] Dingqi Yang, Benjamin Fankhauser, Paolo Rosso, and Philippe Cudré-Mauroux. 2020. Location Prediction over Sparse User Mobility Traces Using RNNs: Flashback in Hidden States! In *IJCAI*. 2184–2190.

Table 6: MAPE Performance of individual methods solving the NexLoc task on the IST dataset

	FPMC	BERT4Rec	Caser	SRGNN	SASRec
Semi-Markov	0.9968	0.9965	0.9960	0.9983	0.9971
Time Geo	3.5045	3.7099	3.5759	3.3284	2.9051
RMTTP	0.8838	0.7807	0.7849	0.8156	0.8250
ERTPP	0.7783	0.7618	0.7726	0.8007	0.8136
THP	0.9598	0.7937	0.7955	0.8205	0.8758
ActSTD	0.8454	0.4274	0.5147	0.4729	0.6069
LogNormMix	0.8442	0.8523	0.8689	0.8686	0.8621
LSTM	0.6218	0.5986	0.6191	0.6707	0.6887
SeqGAN	0.5952	0.7406	1.7957	0.1857	0.1685
MoveSim	0.5114	0.5820	0.5215	0.7367	0.5557
VOLUNTEER	0.5279	0.4924	0.5147	0.5719	0.6019
DiffTraj	1.3272	1.2091	0.9671	1.0018	1.0399
MIRAGE	0.3111	0.2912	0.2921	0.3195	0.3089

- [57] Dingqi Yang, Terence Heaney, Alberto Tonon, Leye Wang, and Philippe Cudré-Mauroux. 2018. CrimeTelescope: crime hotspot prediction based on urban and social media data fusion. *World Wide Web* 21 (2018), 1323–1347.
- [58] Dingqi Yang, Bin Li, and Philippe Cudré-Mauroux. 2016. POIsketch: semantic place labeling over user activity streams. In *IJCAI*. 2697–2703.
- [59] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudré-Mauroux. 2019. Revisiting user mobility and social relationships in lbsns: a hypergraph embedding approach. In *WWW*. 2147–2157.
- [60] Dingqi Yang, Bingqing Qu, Jie Yang, and Philippe Cudré-Mauroux. 2020. Lbsn2vec++: Heterogeneous hypergraph embedding for location-based social networks. *IEEE Transactions on Knowledge and Data Engineering* 34, 4 (2020), 1843–1855.
- [61] Mao Ye, Peifeng Yin, and Wang-Chien Lee. 2010. Location recommendation for location-based social networks. In *SIGSPATIAL*. 458–461.
- [62] Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. Seqgan: Sequence generative adversarial nets with policy gradient. In *AAAI*, Vol. 31.
- [63] Jing Yuan, Yu Zheng, Xing Xie, and Guangzhong Sun. 2011. Driving with knowledge from the physical world. In *KDD*. 316–324.
- [64] Yuan Yuan, Jingtao Ding, Chenyang Shao, Depeng Jin, and Yong Li. 2023. Spatio-temporal Diffusion Point Processes. *arXiv:2305.12403* (2023).
- [65] Yuan Yuan, Jingtao Ding, Huandong Wang, Depeng Jin, and Yong Li. 2022. Activity trajectory generation via modeling spatiotemporal dynamics. In *KDD*. 4752–4762.
- [66] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Yushuo Chen, Xingyu Pan, Kaiyuan Li, Yujie Lu, Hui Wang, Changxin Tian, Yingqian Min, Zhichao Feng, Xinyan Fan, Xu Chen, Pengfei Wang, Wendi Ji, Yaliang Li, Xiaoling Wang, and Ji-Rong Wen. 2021. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. In *CIKM*. ACM, 4653–4664.
- [67] Yuren Zhou, Billy Pik Lik Lau, Chau Yuen, Bige Tuncer, and Erik Wilhelm. 2018. Understanding Urban Human Mobility through Crowdsensed Data. *IEEE Communications Magazine* 56, 11 (2018), 52–59.
- [68] Yuanshao Zhu, Yongchao Ye, Shiyao Zhang, Xiangyu Zhao, and James Yu. 2024. DiffTraj: Generating gps trajectory with diffusion probabilistic model. *NIPS* 36 (2024).
- [69] Simiao Zuo, Haoming Jiang, Zichong Li, Tuo Zhao, and Hongyuan Zha. 2020. Transformer hawkes process. In *ICML*. PMLR, 11692–11702.

A BIASES IN DOWNSTREAM TASKS

We show the performance variation of different methods solving the same downstream task in our task-based evaluation on MAPE in Table 6. Taking MoveSim and VOLUNTEER as examples, MoveSim is better than VOLUNTEER when benchmarking using FPMC or SASRec in this task, while we have the opposite results if benchmarking using BERT4Rec, Caser, or SRGNN.

B THE COMPLEXITY OF MIRAGE

Let n denote the sequence length, d is the GRU hidden state size, and z is the latent size of user VAE. Our time decoder with intensity-free TPP and category decoder are linear layers, taking $O(nd)$ time complexity and $O(d)$ space complexity. Similarly, the location decoder

with two modes with time complexity $O(2nd)$ and space complexity $O(2d)$. The GRU backbone has time and space complexities of $O(nd^2)$ and $O(d^2)$, respectively. For the user VAE, we consider both input and hidden size to be d for simplicity, and both time and space complexities are $O(d^2)$ as usually $d^2 \gg dz$. In summary, MIRAGE has time complexity $O(nd^2)$ and space complexity $O(d^2)$.

C BASELINES

Semi-Markov [29] model uses exponential distribution with Gamma priors for time intervals and incorporates a Dirichlet prior to construct the transition matrix for Bayesian inference. **TimeGeo** [25] designs a statistical EPR model integrated with temporal information including weekly home-based tour number, dwell rate, and burst rate to further characterize temporal choices on human mobility. **RMTTP** [10] uses Recurrent Neural Networks to jointly model the time and mark (location ID) dependency over history information. **ERTPP** [52] adopts distinct RNNs to independently model the timing of the next event and its associated mark. **THP** [69] incorporates the self-attention mechanism with the Hawkes Process to capture long-term dependencies in event sequence data. **LSTM** [24] learns the sequence patterns of human trajectory to predict both the location and time of the next event in a human trajectory. **SeqGAN** [62] introduces reinforcement learning into the GAN model to solve the sequence generation problem. **MoveSim** [12] is a GAN-based framework that integrates physical regularities and prior knowledge of human mobility in trajectory generation. **ActSTD** [65] enhances the dynamic modeling of individual trajectories by utilizing neural ordinary equations. **LogNormMix** [42] defines an intensity-free neural TPP modeling the conditional probability density distribution of trajectory events as a log-normal mixture. We extend it as marked TPP by sampling the time and location (i.e., mark) of the next event in a trajectory. **VOLUNTEER** [32] incorporates a two-layer VAE model with a temporal point process to capture the characteristics of human mobility from both group and individual views. **DiffTraj** [68] uses a diffusion probabilistic model for continuous location generation. We choose the nearest POI to each generated GPS coordinate as the generated POI.

D ALGORITHMS FOR THE RECLOC TASK

BPR [40] is a recommendation approach designed for implicit feedback. It operates by minimizing a pairwise ranking loss to learn user preferences on times effectively. **DMF** [55] incorporates a matrix factorization model with neural network architecture to the representations of users and locations. **LightGCN** [20] leverages neighborhood aggregation technique to learn user and location representations on a user-location interaction graph. **MultiVAE** [31] extends Variational Autoencoders (VAEs) to collaborative filtering for implicit feedback. **NeuMF** [21] proposes a neural network-based collaborative filtering technique and modeling user–location interaction function with non-linearities.

E ALGORITHMS FOR THE NEXLOC TASK

FPMC [41] uses matrix factorization techniques to estimate the personalized transition matrix of POIs in user mobility trajectories. **BERT4Rec** [46] uses bidirectional self-attention network

Table 7: Performance in the Task-Based Evaluation on MSPE

Method	TKY				IST				NYC			
	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim
Semi-Markov	0.7117	0.9839	0.2929	0.4125	0.8427	0.9939	0.2940	0.3378	0.7455	0.9903	0.2402	7.0165
Time Geo	0.6558	3.4390	0.0800	0.7302	14.1917	12.0970	0.2351	0.7822	1.2609	10.0716	0.1004	0.6650
RMTTP	0.7664	0.7472	0.1040	0.8849	0.5395	0.6712	0.2150	0.7432	0.8020	0.8854	0.1785	0.7773
ERTPP	0.2322	0.6696	0.0405	0.3146	0.1995	0.6179	0.2209	0.5942	0.5054	0.8387	0.5833	13.0105
THP	5.1891	5.6151	0.1768	0.1322	0.3651	0.7251	0.2737	0.9674	35.1497	16.2586	0.1697	0.4959
ActSTD	0.2921	0.4954	0.1528	0.1044	0.3556	0.3508	0.2045	0.2343	0.7544	0.6861	0.1763	0.9249
LogNormMix	0.2440	0.5112	0.0381	0.0299	0.2089	0.7387	0.1781	0.0473	0.5628	0.7577	0.3148	0.5571
LSTM	0.2653	0.5229	0.0355	0.0355	0.1955	0.4116	0.1820	0.1619	0.5268	0.2871	0.1221	0.4960
SeqGAN	0.2204	0.2290	0.0986	0.0453	1.2233	0.8872	0.1936	0.1986	4.8884	0.3213	0.5830	5.9233
MoveSim	1.2413	0.1249	0.0870	0.0353	30.6287	0.3623	0.2802	0.3817	5.2941	0.6145	0.2053	5.7852
VOLUNTEER	1.2227	0.2921	0.0972	0.2110	0.3791	0.2987	0.2433	0.1426	0.4571	0.7171	0.6019	2.9920
DiffTraj	0.5900	1.2704	0.2280	0.7155	0.5444	1.3134	0.2570	1.6888	0.4514	4.0397	0.1984	0.8163
MIRAGE	0.1416	0.1026	0.0222	0.0289	0.1792	0.0941	0.1235	0.0770	0.5062	0.0592	0.0147	0.8948

Table 8: Ablation study in the Task-Based Evaluation on MSPE

Method	TKY				IST				NYC			
	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim	LocRec	NexLoc	SemLoc	EpiSim
MIRAGE-noTPP	0.1583	0.1100	0.0397	0.4335	0.2176	0.2063	0.1631	0.2274	0.5299	0.3324	0.0184	3.8921
MIRAGE-noEPR	0.1517	0.1893	0.0288	0.0339	0.1879	0.2618	0.1869	0.0862	0.7906	0.2815	0.0173	0.9989
MIRAGE-noIMI	0.2250	0.2355	0.0421	0.0364	0.2670	0.1294	0.1252	0.0729	0.6697	0.0810	0.0982	0.9606
MIRAGE-TD	0.2352	0.2295	0.0653	0.0344	0.1903	0.1385	0.1264	0.0997	0.8839	0.0649	0.1032	1.1433
MIRAGE	0.1416	0.1026	0.0222	0.0289	0.1792	0.0941	0.1235	0.0770	0.5062	0.0592	0.0147	0.8948

Table 9: Parameters for COVID-19 simulation (d : day)

Parameters	c	T	T_i	T_f	R_0	β	α	γ
Value	0.2	5.8d	5.2d	11d	2.2	R_0/T	$1/T_i$	$1/T_f$

models human behavior sequences by employing a Cloze task approach. **Caser** [47] employs hierarchical and vertical CNNs to capture union-level sequential patterns and skip behaviors, enabling sequence-aware recommendation. **SRGNN** [51] designs a session-based recommendation model by utilizing the basic RNNs to predict the next location based on historical trajectories. **SASRec** [27] employs a multi-head attention mechanism to make predictions based on the historical trajectory of users.

F EPIDEMIC SIMULATION SETTINGS

We simulate the COVID-19 spreading with the SEIR model following recent works [12, 30, 65]. In particular, the simulation involves eight fundamental parameters: the close contact ratio (c), transmission period (T), incubation period (T_i), infection period (T_f), reproduction rate (R_0), transmission probability (β), infectious rate (α), and recovery rate (γ). Table 9 shows the parameter values.

During the simulation, we assume that infected or exposed individuals contact with s susceptible individuals connected by edges in the contact network each day. The probability of two people with an edge in the contact network becoming close contact is c . The transmission probability β is calculated by dividing the basic reproduction rate R_0 by the average duration (5.8 days) from onset to first medical visit and isolation. The infectious rate from the exposed state α is estimated as the reciprocal of the incubation period, which is the average time exposed but not infectious (5.2 days in [30]). The daily probability of transitioning to the removed state from infectious γ , is computed based on the average infection

Table 10: Parameters for influenza simulation

Parameters	c	β	α	γ
Value	0.2	0.402	0.526	0.244

period (11 days in [30]). The the infection process are as follows:

$$\frac{dS}{dt} = -sc\beta, \frac{dE}{dt} = sc\beta - \alpha E, \frac{dI}{dt} = \alpha E - \gamma I, \frac{dR}{dt} = \gamma I \quad (17)$$

where S represents the number of susceptible individuals, E is the number of exposed individuals, I stands for the number of infectious individuals, and R denotes the number of removed individuals. We randomly select 50 individuals as exposed and label their status accordingly. Using the differential equations 17, we simulate the spread of the epidemic and record the daily counts of exposed, infectious, and removed individuals.

Similarly to the COVID-19 simulation, we simulate influenza spreading using the parameters suggested by [4], where the transmission probability β infectious rate α and removed rate γ are shown in Table 10.

G PERFORMANCE ON MSPE

Table 7 shows performance comparison in the task-based evaluation on MSPE. Similar to the results on MAPE, we observe that MIRAGE achieves the best performance in most cases. Compared to the best-performing baselines, our MIRAGE achieves performance improvement with 17.6 % on average. Table 8 shows the ablation study results in the task-based evaluation on MSPE. Similar to the results on MAPE, we also observe that MIRAGE outperforms its variants MIRAGE-noTPP, MIRAGE-noEPR, MIRAGE-noIMI, and MIRAGE-TD by 41.8%, 28.7%, 30.0%, and 33.3%, respectively, further validating our key design choices.