Privacy-Aware Adversarial Network in Human Mobility Prediction

Yuting Zhan
Imperial College London, UK
yuting.zhan18@imperial.ac.uk

Hamed Haddadi
Imperial College London, UK
h.haddadi@imperial.ac.uk

Afra Mashhadi
University of Washington, USA
mashhadi@uw.edu

ABSTRACT

As mobile devices and location-based services are increasingly developed in different smart city scenarios and applications, many unexpected privacy leakages have arisen due to geolocated data collection and sharing. User re-identification and other sensitive inferences are major privacy threats when geolocated data are shared with cloud-assisted applications. Significantly, four spatio-temporal points are enough to uniquely identify 95% of the individuals, which exacerbates personal information leakages. To tackle malicious purposes such as user re-identification, we propose an LSTM-based adversarial mechanism with representation learning to attain a privacy-preserving feature representation of the original geolocated data (i.e., mobility data) for a sharing purpose. These representations aim to maximally reduce the chance of user re-identification and full data reconstruction with a minimal utility budget (i.e., loss). We train the mechanism by quantifying privacy-utility trade-off of mobility datasets in terms of trajectory reconstruction risk, user re-identification risk, and mobility predictability. We report an exploratory analysis that enables the user to assess this trade-off with a specific loss function and its weight parameters. The extensive comparison results on four representative mobility datasets demonstrate the superiority of our proposed architecture in mobility privacy protection and the efficiency of the proposed privacy-preserving features extractor. We show that the privacy of mobility traces attains decent protection at the cost of marginal mobility utility. Our results also show that by exploring the Pareto optimal setting, we can simultaneously increase both privacy (45%) and utility (32%).

KEYWORDS

LSTM neural networks, mobility prediction, data privacy, adversarial learning

1 INTRODUCTION

Geolocation and mobility data collected by location-based services (LBS) [32], can reveal human mobility patterns and address various societal research questions [34]. For example, Call Data Records (CDR) have been successfully used to provide real-time traffic anomaly as well as event detection [54, 56], and a variety of mobility datasets have been used in shaping policies for urban communities [19] and epidemic management in the public health domain [44, 45].

Proceedings on Privacy Enhancing Technologies 2023(1), 556–570

https://doi.org/10.56553/popets-2023-0032

From an individual-level perspective, users can benefit from personalized recommendations when they are encouraged to share their location data with third parties or other service providers (SPs, e.g., social platforms) [16]. Human mobility prediction based on users' traces, a popular and emerging topic, supports a series of important applications. For instance, one of the prerequisites for a successful LBS-recommendation system is the ability to predict users' activities or locations ahead of time, tracking their intentions and forecasting where they will go [22].

While there is no doubt about the usefulness of predictive applications for mobility data, privacy concerns regarding the collection and sharing of individuals' mobility traces have prevented the data from being utilized to their full potential [6, 35, 51]. A mobility privacy study conducted by De Montjoye et al. [13] illustrates that four spatio-temporal points are enough to identify 95% of the individuals, which exacerbates the user re-identification risk and could be the origin of many unexpected privacy leakages. Additionally, with increasingly intelligent devices and sensors being utilized to collect information about users' locations, a malicious third party can derive increasing intimate details about users' lives, from their social life to their preferences. Hence, a mechanism capable of decreasing the chance of user re-identification against malicious attackers or untrusted SPs can offer enhanced privacy protection in mobility data applications, as human mobility traces are highly unique.

In the past decade, the research community has extensively studied the privacy of geolocated data via various location privacy protection mechanisms (LPPM) [20, 21]. Some traditional privacy-preserving approaches (e.g., k-anonymity and geo-masking) have shown to be insufficient to prevent users from being re-identified [13, 23, 39, 53]. Differential privacy (DP), another popular notion, is shown to be a limited metric for location trace privacy since temporal correlations are not taken into account [51]. [16] also states that DP and k-anonymity are meant to ensure the privacy of a single data point in time. In general, many DP for LBS (DP-L) mechanisms [2, 11, 28] attempt to protect the *user's location* instead of the *user's identity*, which is different from our problem's scope. More recently, some related works have successfully applied machine-learning- or deep-learning-based approaches to explore effective LPPMs. Rao et al. proposed a model based on Generative Adversarial Network (GAN) [24] to generate privacy-preserving synthetic mobility datasets for data sharing and publication [48]. Feng et al. investigated human mobility data with privacy constraints via federated learning, achieving promising prediction performance while preserving the personal data on the local devices [18]. Though these works provide promising architectures to protect location privacy, the mobility data's privacy protection and utility degradation have not been thoroughly investigated, especially in reducing the chance of user re-identification. Our work extends these machine-learning-based mechanisms and explores

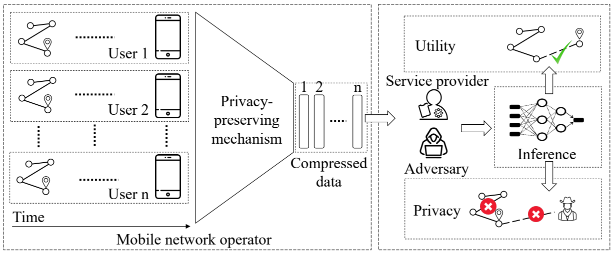

Figure 1: Privacy protection in user's location data collection and sharing. Users share their daily traces to a trusted mobile network operator; these traces are aggregated with a privacy-preserving mechanism and shared as a compressed data format; the compressed data should allow utility inference and avoid privacy inference.

the privacy-utility trade-off on mobility data in terms of declining the effectiveness of privacy inference attacks while maintaining its predictability. Moreover, research on human mobility shows that the predictability of users' location trajectories or mobility, and the particular constraints of users' movements, are sufficient to reconstruct and/or identify anonymous or perturbed locations [52]. This specific confrontation makes the trade-off between mobility predictability and users' identity more challenging.

Consider a scenario, shown in Figure 1, where users share their daily traces to a trusted mobile network operator, which then aggregates these traces in a privacy-preserving approach and sends them to third parties or other SPs with/without users' consent. These users may want to minimize the risk of being re-identified and trajectory reconstructed by those who will access these released data. However, they would like to keep receiving potential effective services from SPs. Therefore, a privacy-preserving mechanism, which can release required information for the services (i.e., utility) while features or patterns that facilitate fully data reconstruction or user re-identification are obscured (i.e., privacy), is beneficial. The compressed data encoded by this privacy-preserving mechanism is freely accessed to SPs for the inference tasks, and SPs are free to use any prediction algorithms of their choice.

To this end, we propose a **p**rivacy-aware **a**dversarial network to train a feature extractor Enc_L for **mo**bility privacy, namely Mo-PAE. It is based on representation learning and aims to ease data sharing privacy concerns from privacy inference attacks. Inspired by PAN (privacy adversarial network) [37], we employ adversarial learning to better balance the potential trade-off between privacy and utility. In contrast to PAN, which focuses on the privacy of images, our approach is designed for complex time-series data that exhibits spatial-temporal characteristics. At the core of our architecture lies an auto-encoder (AE) and long short-term memory (LSTM) layers with three branches, corresponding to the three training optimization objectives of the feature extractor Enc_L: i) to *maximize* the loss associated with the reconstructed output by generative learning, ii) to *minimize* the prediction loss using the learned representation from the Enc_L by discriminative learning, and iii) to *minimize* the percentage of users who are re-identifiable through their trajectories by discriminative learning. We explore and quantify the privacy-utility trade-off provided by Mo-PAE in terms of data reconstruction risk, user re-identification risk, and mobility predictability. The results show that our proposed mechanism achieves a better privacy level with the same utility loss and vice versa.

The contributions of our work are the following:

- We propose a **p**rivacy-aware **a**dversarial network to train an effective feature extractor Enc_L for **mo**bility privacy, namely Mo-PAE;
- We report the analysis of Mo-PAE by a comprehensive evaluation of four real-world representative mobility datasets;
- We provide an extensive analysis of different inference tasks and quantify the privacy and utility bound of the target mobility dataset, along with a trade-off analysis between these contrasting objectives;
- We compare our model with, i) a famous DP notion that developed on the idea from Geo-indistinguishability [2] (namely GI-DP); ii) a state-of-the-art GAN-based mechanism that attempts to generate synthetic privacy-preserving mobility data (namely TrajGAN [48]); iii) as well as the optimal LSTM-based inference models, and obtain favourable results.

The rest of this paper is structured as follows: we review the related work in Section 2; the proposed Mo-PAE is described in detail in Section 3; we describe the experimental settings in Section 4; we demonstrate an evaluation of our mechanism over four mobility datasets with baseline comparisons in Section 5; Section 6 reports an in-depth discussion of our setting; finally, we conclude the paper with future work directions in Section 7.

2 RELATED WORK

2.1 Notions of Location Privacy

Diverse privacy notions, *direct* or *indirect*, for the LBSs have been proposed and evaluated in the literature. In [2], various *direct notions* of location privacy and the techniques to achieve them are examined and concluded, including but not limited to expected distance error, k-anonymity, differential privacy (DP), and other location-privacy metrics. First, the expectation of distance error reflects the accuracy when an adversary guesses the user's real location in a location-obfuscation mechanism by using the available side information. In [52], an optimal LPPM strategy and its corresponding optimal inference attack are obtained by formalizing the mutual optimization of user-adversary objectives (location privacy vs correctness of localization). Second, k-anonymity is the most widely used privacy notion for the LBSs [55]. These systems aim to protect the *user's identity*, requiring that the attacker cannot infer the correct user among a set of k different users. Third, DP [14] is an emerging notion initially formulated in the context of statistical databases and aims to protect an individual's data while publishing aggregate information about the dataset. More precisely, a randomization mechanism M gives ϵ-differential privacy for all neighbouring datasets D and D', and the difference between D and D' is within a bound of e^{ϵ}. One of the popular mechanisms to achieve DP perturb the original query result using random noise that is calibrated with the privacy budget ϵ and defines a global sensitivity for all neighbouring D and D' [15]. The work in [17] reviews research works done in differential privacy targeted toward location data from a data flow perspective, including collection, aggregation, and mining. [2] proposed a Geo-indistinguishability notion based on differential privacy and a planar Laplace mechanism. Significantly, different from the systems in k-anonymity category aim at protecting *user's identity*, DP mechanisms are interested in protecting the *user's locations* [2, 11, 26, 28]. Apart from three mainstream approaches, the location cloaking mechanism tries to define the uncertainty region and measure privacy by the size of the cloak and by the coverage of sensitive regions; the inaccuracy of the sensing technology tries to achieve a certain level of privacy by increasing uncertainty; and transformation-based approach tries to make user's location invisible to the service provider.

On the other hand, *indirect* notions of location privacy arise with the emerging machine learning-based mechanism, which assesses the privacy guarantee by measuring the effectiveness of target privacy inference attacks [25, 30, 39, 40]. For instance, in [39], the authors access the privacy guarantee by measuring the difference between an adversary's belief about sensitive inferences before and after observing the released data. In [30], the proposed mechanism provides privacy guarantees by capturing how well an adversary does in terms of inferring the private variables. In general, for any LBS, their main privacy concerns can be summarized in two categories. One is the attack on the *user's identity*, which can be re-identified maliciously. For instance, even if the adversary is assumed to be unaware of the user identity of a trace, they can infer *user's identity* or additional sensitive information due to the location information leakage based on publicly accessible background information. The other attack is the one on *user's location* while the adversary has legible access to *user's identity*. In this manner, *user's locations* are sensitive, which could exert a significant impact on other sensitive personal details, such as religious affiliation, sexual orientation, economic condition, health status, and so on.

In our work, we are interested in protecting *user's identity* as the privacy scope, which is similar to the location privacy notion defined by the k-anonymity, and taking the real/distorted *user's location* as input for the personal recommendation model to provide contextual services for their future travels. In general, DP paradigms have the most formal privacy guarantee when compared with the others, however, they are not immune to inference attacks [10, 27]. We will also compare our proposed model with one popular DP paradigm on location privacy to illustrate the ineffectiveness of DP to our research question. More details on our privacy definitions are in Section 3.

2.2 Privacy Preserving Techniques for Spatial-Temporal Data

Current privacy-preserving techniques for spatial-temporal data focus on two research streams. One is the DP approach to grouping and mixing the trajectories from different users so that the identification of individual trajectory data is converted into a k-anonymity problem [1, 2, 57]. For example, a recent Privacy-Preserving Trajectory Framework (PPTPF) [60] applies the k-indistinguishability to anonymize trips for each user by condensing them into $k-1$ trajectories and determining $k-1$ anonymized clusters of trips.

The other stream focuses on synthetic data generation [9, 31, 46, 49]. Synthetic data generation methods have been extensively studied in recent years as a way of tackling privacy concerns of location-based datasets. The majority of existing mobility synthesis schemes are mainly categorized into two approaches: one is a more traditional, simulation-based approach, while the other is a more recent, neural network-based generative modeling approach that utilizes recurrent autoencoders and generative adversarial networks to produce realistic trajectories [50]. Simulation-based approaches generate mobility traces by modeling overall user behavior as a stochastic process, such as a Markov chain model of transition probabilities between locations, and then drawing random walks, potentially with additional stochastic noise added, as demonstrated in Xiao et al [58]. These approaches require considerable feature engineering effort and struggle to capture longer-range temporal and spatial dependencies in the data [38] and are thus limited in their ability to preserve the utility of the original datasets. In contrast, the generative neural network approach synthesizes user mobility traces by learning via gradient descent back-propagation, and then the optimal weights are utilized for decoding a high-dimensional latent vector representation into sequences that closely resemble the original data. Such traces can maintain important statistical properties of the original data while taking advantage of noise introduced in the reconstruction process, to improve data subject anonymity. Huang et al [31] demonstrates the use of a variational autoencoder network to reconstruct trajectory sequences, while Ouyang et al [46] utilizes a convolutional GAN, but neither work directly makes a quantitative assessment of the extent of privacy protection that their algorithms provide [31, 46]. The TrajGAN by Rao et al [48] is a state of the art example of the generative trajectory modeling approach, which quantifies its privacy protection by demonstrating a significant decline in the performance of a second user ID classifier model on the synthetic outputs compared to the original input trajectories. For these reasons, we used TrajGAN as a baseline for comparison.

Our proposed model takes the neural network-based generative modeling approach, but differs from existing methods, where we utilize a combined, multi-task adversarial neural network to simultaneously reconstruct trajectories, predict next locations, and re-identify users, from the same learned latent vector representation. We seek an optimal trade-off between the three tasks' individual losses by optimizing a sum loss function with per-task weights, improving the controllability of the relative utility and privacy of the outputs.

3 DESIGN OF MO-PAE

3.1 Definition of Important Terms

3.1.1 Mobility Trace. The raw geolocated data or other mobility data commonly contain three elements: user identifiers $u \in U$, timestamps $t \in T$, and location identifiers $l \in L$. Hence, each location record r could be denoted as $r_{(u,i)} = [u, t_i, l_i]$, while each location sequence S is a set of ordered location records $S_{(u,n)} = \{r_{(u,1)}, r_{(u,2)}, r_{(u,3)}, \cdots, r_{(u,n)}\}$, namely *mobility trace* of *user u*. Therefore, given the past mobility trace $S_{(u,n)}$, the mobility prediction task is to infer the most likely location l_{n+1} at the next timestamp t_{n+1} for the user u. The data fed into the proposed architecture are a list of traces with a specific sequence length (i.e., *SL*). For instance, if the sequence length is 10, that indicates each trace contains 10 history location records r, $S_{10} = \{r_1, r_2, r_3, \cdots, r_{10}\}$, and $SL = 10$. In this paper, we assume that different users' mobility traces are collected and aggregated (denoted as data X) by trusted telecom operators or social platforms and shared with third-party SPs.

3.1.2 User Re-identification. The user re-identification risk arises because of the high uniqueness of human traces [13] and could be the origin of many unexpected privacy leakages. We assume each trace S is originally labeled with a corresponding user identifier u, and the user re-identification is to infer the user u to whom the target trace $S_n = \{r_1, r_2, r_3, \cdots, r_n\}$ belongs. We thereby leverage the user identifiers u as the ground-truth values for the user identity classes. This identity information is what we want to protect in the proposed adversarial network.

3.2 Problem Definition

3.2.1 Definition of Utility and Privacy. On the one hand, mobility datasets are of great value for understanding human behaviour patterns, smart transportation, urban planning, public health issue, pandemic management, etc. Many of these applications rely on the next location forecasting of individuals, which in the broader context, can provide an accurate portrayal of citizens' mobility over time and inform the allocation of public resources and community services. Therefore, we focus on the capability of *mobility prediction* (*next location forecasting*) in this paper, and leverage the accuracy of the prediction as an important metric for quantifying the *data utility*. On the other hand, with increasing intelligent devices and sensors being utilized to collect information about human activities, the traces also increasingly expose intimate details about users' lives, from their social life to their preferences. In this manner, the capability of user re-identification is important to balance the risks and benefits of mobility data usage, for all data owners, third parties, and researchers. We then leverage the efficient reduction of data reconstruction risk and user re-identification risk as the *privacy protection* metrics. Moreover, research in [52] shows that the predictability of users' mobility, and the particular constraints of users' movements, are sufficient to reconstruct and/or identify anonymous or perturbed locations [52]. This confrontation makes the trade-off between keeping mobility predictability and reducing the chance of user re-identification more interesting. For instance, an adversary can re-identify anonymous users' traces given the users' mobility profile [12]; infer the users' next activities from the frequency of location visits [22]; even obtain the personal home or working address from the trajectories [47].

In this work, we design a model to protect location privacy regarding users' identity and data integrity while simultaneously minimizing the service quality (i.e., accuracy of next location forecasting) degradation stemming from the obfuscation of true data. Specifically, users' mobility traces are collected and fed into this proposed model, encoded as privacy-preserving representations that allow third parties and other SPs freely access. At the same time, two built-in adversaries, which try to achieve maximum accuracy in user re-identification and trace reconstruction during the adversarial training, are simulating the strong privacy adversaries that can attain disclosed sensitive information and examine the quality of feature representations instantly. Overall, the encoded privacy-preserving representations should retain as little user identifiable information as possible, as well as the data reconstruction information, to decrease the user re-identification accuracy and increase the location obfuscation.

Hence, we summarize the *Utility*, *Privacy I* and *Privacy II* of the encoded feature representations as follows:

Utility (U): the encoded representations should retain information about mobility predictability (i.e., forecasting accuracy, higher accuracy indicates higher utility).

Privacy I (*PI*): the encoded representations should contain little information advantage to the data reconstruction (i.e., more information loss in the reconstruction process); represented as the distortion increment (i.e., *Euclidean* [3] and *Manhattan* distance [7]) between the reconstructed data X' and the original data X.

Privacy II (*PII*): the encoded representations should contain little information advantage to the user re-identification task (i.e., the user de-identification effectiveness); measured by the degradation of the user re-identification accuracy.

3.2.2 Privacy vs. Utility Trade-off. An effective LPPM must consider three fundamental elements: i). the privacy requirements of the users (namely *privacy gain*); ii). the adversary's knowledge and capabilities; iii). and maximal tolerated service quality degradation stemming from the obfuscation of true locations (namely, *utility loss*) [52]. There is an inherent trade-off between location privacy protection and utility degradation [33]. That is, achieving a better level of privacy protection may require sacrificing the service quality provided by the data. Such trade-off is omnipresent in various privacy protection mechanisms, especially in location obfuscation mechanisms. Higher privacy protection is achieved when the probability of an adversary inferring the true location of the user decreases, however, the result of a query based on the obfuscated location is significantly different from the actual interest of the user. The privacy-utility trade-off, hence, needs to be examined and analyzed to guarantee the efficiency of the privacy protection mechanism.

In this work, *utility loss* denotes the accuracy degradation after the proposed privacy protection mechanism is applied, and the *privacy gain* (in terms of *PI* and *PII*) quantifies the protected privacy information. To be specific, a more obfuscated dataset will tend to perform better at preserving privacy, but worse at preserving utility, and vice versa. Hence, monitoring these two performance metrics in tandem allows users to select the optimal privacy-utility trade-off for their use cases, given their hyperparameter selections. Our Mo-PAE model is designed to train a features Encoder $Enc_L(X)$ that could convey more information on the utility but less on privacy and investigate a better trade-off between them. More details will be discussed in the following sub-section.

3.3 Mo-PAE Overview

Our proposed **p**rivacy-preserving **a**dversarial feature **e**ncoder on mobility data, denoted as the *Mo-PAE*, is based on representation learning and adversarial learning and aims to ease data sharing privacy concerns. Figure 2 presents the basic workflow of the proposed Mo-PAE. It composes of three crucial units: data reconstruction risk unit (DRU), mobility prediction unit (MPU), and user re-identification risk unit (URU).

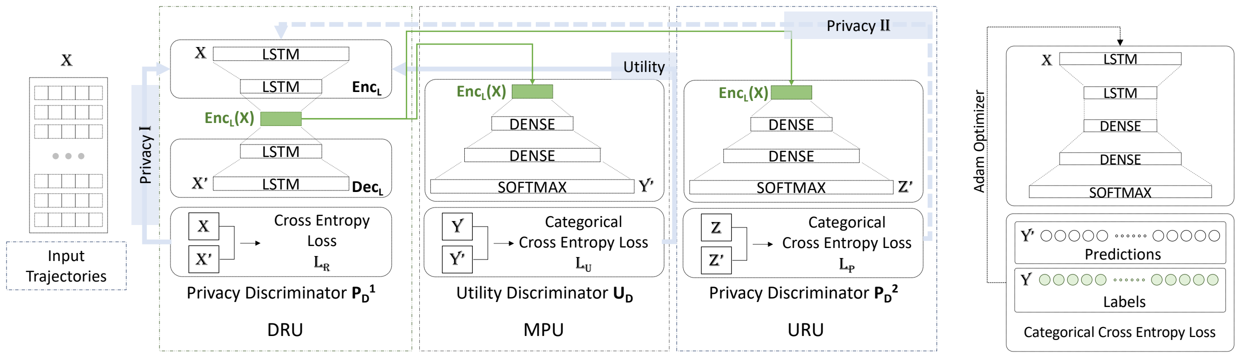

(a) Mo-PAE **(b) Optimal-IMs**

Figure 2: (a) Schematic overview of the proposed privacy-preserving adversarial architecture (Mo-PAE), consisting of data reconstruction risk unit (DRU), mobility prediction unit (MPU), and user re-identification risk unit (URU); (b) The baseline LSTM network for optimal classifiers (Optimal-IMs).

3.3.1 Composition Units of Mo-PAE.

I. Mobility Prediction Unit (MPU):

The MPU unit is composed of three parts, the input part with the multi-modal embedding of trace information, the sequential part with LSTM layers [29], and an output part with the softmax activation function. As per the definition mentioned earlier, the traces in this work are shown as location sequences S. First, the location identifiers l and timestamps t are converted into one-hot vectors. We then employ LSTM layers to model the mobility patterns and sequential transition relations in these mobility traces. As a prominent variant of the recurrent neural network, LSTM networks exhibit brilliant performance in modeling the entire data sequences, especially for learning long-term dependencies via gradient descent [61]. Following the sequential module, the softmax layer outputs the probability distribution of the prediction results. This probability distribution is converted to the top-n accuracy metrics to illustrate the unit performance.

II. Data Reconstruction Risk Unit (DRU):

The DRU is the encoder Enc_L unit in reverse, also denoted as Dec_L, which is regarded as the first *privacy discriminator* P_D^1 in the proposed architecture. It is designed to evaluate the distance $d(\cdot,\cdot)$ (i.e., *Privacy I*) between the reconstructed data X' and the original input data X. A malicious party is free to explore any machine learning model and reconstruct the data if they have the shared extracted features f. We use a layer-to-layer reverse architecture of our encoder Enc_L to build the *data reconstruction unit* to act as a robust built-in adversary. To compare with baseline models and keep the comparison in a line, we measure the distance $d(\cdot,\cdot)$ between the X and X' by leveraging the *Euclidean* and *Manhattan* distance as our metrics. Both of them are widely used in location privacy literature [2, 16].

III. User Re-identification Risk Unit (URU):

The URU is regarded as the second *privacy discriminator* P_D^2 in the proposed architecture. The unit is composed of three parts, the input part with the one-hot embedding of user identity, the sequential part with LSTM layers, and an output part with softmax function. First, the user identity list is converted into one-hot vectors. Similar to the MPU, the URU also applies LSTM layers to better extract the spatial and temporal characteristics of the context. A softmax function with a cross-categorical entropy loss function is applied to output a categorical probability distribution of the user re-identification task. We then use the top-n accuracy of this classifier as the metric of user re-identification privacy risk (i.e., *Privacy II*). The more accurately a classifier can re-identify the user when given a trajectory, the higher the risk of disclosing private data. Same as P_D^1, P_D^2 is designed as the built-in adversary to infer the ability of generated features in protecting users' sensitive information.

The overall architecture eventually learns to fool both built-in adversaries, P_D^1 and P_D^2, while maintaining mobility predictability. In this manner, both adversaries are assumed to be free to access the exclusive feature representations and the entire encoder network, which allows them to have the optimal decoder setting. We will discuss the effectiveness of two privacy inference attacks in Section 5.3.

3.3.2 Overall Design.

When three units train concurrently, the MPU is regarded as the *utility discriminator* U_D, while DRU and MPU act as two built-in adversaries and are regarded as the two *privacy discriminators*, P_D^1 and P_D^2, respectively. The built-in adversary has been used as an effective adversarial regularization to prevent inference attacks, *e.g.* in the classification setting [43] or in various GAN models [42] for privacy-preserving purpose. In this work, P_D^1 and P_D^2 are simulating malicious SPs, who attempt to obtain sensitive information (i.e., maximize the accuracy of privacy inference tasks), while the encoder Enc_L is trained to produce feature representations f to the advantage of U_D but to the disadvantage of P_D^1 and P_D^2, by jointly optimizing the hybrid losses of the DRU, MPU, and URU simultaneously, during adversarial training. Therefore, in the Mo-PAE, the encoder Enc_L, and three discriminators U_D, P_D^1, P_D^2 play a multi-player game to minimaximize the value function $V(Enc_L, U_D, P_D^1, P_D^2)$:

$$\min_{(Enc_L, U_D)} \max_{(P_D^1, P_D^2)} V(Enc_L, U_D, P_D^1, P_D^2) = \mathbb{E}_{x\sim X}[log U_D(Enc_L(x))] + \mathbb{E}_{x\sim X}[log(1 - P_D^1(Enc_L(x)))] + \mathbb{E}_{x\sim X}[log(1 - P_D^2(Enc_L(x)))] \tag{1}$$

As described in the Eq.1, we design a multi-task adversarial network to learn an LSTM-based encoder $Enc_L(X;\theta)$ with parameter set $\theta \in \Theta$, which can generate the optimized feature representations $f = Enc_L(X;\theta)$ via lowering the privacy disclosure risk of user identification information and improving the task accuracy (i.e., mobility predictability) concurrently. Two potential malicious privacy leakages from URU and DRU, are attempted to retrieve sensitive information from the feature representations f. As built-in adversaries, they have full access to the feature representations f and the entire encoder network with parameter set $f = Enc_L(X;\theta)$. In this manner, they have the optimal decoder setting. Hence, the notion of privacy (*privacy gain*), is measured by the decline of the effectiveness of target inference attacks (i.e., user re-identification attack and data reconstruction attack).

3.3.3 Details of Mo-PAE.

We define the raw mobility data we want to protect as $\mathcal{X}$, trained features as $\mathcal{F}$, and reconstructed data as $\mathcal{X}'$. Given mobility raw data $\mathcal{X}$ for P_D^1 (DRU), the ground-truth label z_i for P_D^2 (URU), and the ground-truth label y_i for utility U_D (MPU), we train the encoder Enc_L to learn the representation $\mathcal{F} = Enc_L(\mathcal{X};\theta_E)$. We design a specific loss function, namely *sum loss* $\mathcal{L}_{sum}$, for this optimization process.

Specifically, when reconstructing the data $\mathcal{X}'$, a decoder Dec_L attempts to recreate the data based on the features $\mathcal{F}$, that is $Dec_L(\mathcal{F};\theta'_D) : \mathcal{F} \rightarrow \mathcal{X}'$. This DRU, the first privacy discriminator P_D^1, is trained as a built-in adversary and tries to achieve sensitive information as much as possible. Hence, the DRU is primarily trained by minimizing the reconstruction loss $\mathcal{L}_\mathcal{R}$:

$$\min \mathcal{L}_\mathcal{R} \Rightarrow \mathcal{L}_\mathcal{R} = d(\mathcal{X}, \mathcal{X}') = \underset{\mathcal{F};\theta'_R}{\operatorname{argmin}} \|Dec_L(\mathcal{F}, \theta'_R) - \mathcal{X}\|^2 \tag{2}$$

The URU, the second privacy discriminator $P_D^2(\mathcal{F};\theta')$, is trained to re-identify whom the target trajectory belongs to. It outputs a probability distribution of predicted user identifiers among Z potential classes. Then in this privacy discriminator, the user re-identification loss $\mathcal{L}_\mathcal{P}$ is primarily trained to minimize, denoted as $\min \mathcal{L}_\mathcal{P}$:

$$\min \mathcal{L}_\mathcal{P} \Rightarrow \mathcal{L}_\mathcal{P} = \underset{\mathcal{F};\theta'_P}{\operatorname{argmin}} \sum_{i=1}^{Z} z_i \log(P_D^1(\mathcal{F};\theta'_P)) \tag{3}$$

The MPU, the utility discriminator $U_D(\mathcal{F};\theta')$, is trained to output a probability distribution of the next location of interest, and this distribution has Y potential classes. Discriminative training of U_D maximizes the prediction accuracy by minimizing the utility loss $\mathcal{L}_\mathcal{U}$ concurrently with minimizing the $\mathcal{L}_{sum}$, denoted as $\min \mathcal{L}_\mathcal{U}$.

$$\min \mathcal{L}_\mathcal{U} \Rightarrow \mathcal{L}_\mathcal{U} = \underset{\mathcal{F};\theta'_U}{\operatorname{argmin}} \sum_{i=1}^{Y} y_i \log(U_D(\mathcal{F};\theta'_U)) \tag{4}$$

Algorithm 1: Training of the Mo-PAE (*Model II*)

Input : Mobility data **X**, real mobility prediction labels **Y**, real user identification labels **Z**, weights: λ_1, λ_2, λ_3

Output : Adversarial Encoder $Enc_L(X;\theta_E,\theta_R,\theta_U,\theta_P)$

Initialize model parameters $\theta_E,\theta_R,\theta_U,\theta_P$;

for *n epochs* **do**

  for $k = 1, \cdots, K_t$ **do**

    1. Sample a mini-batch of mobility trajectories x, prediction labels y, identification labels z

    2. Update θ_E with Adam optimizer on mini-batch loss $L_{sum}(\theta_E,\theta_R,\theta_U,\theta_P,\lambda_1,\lambda_2,\lambda_3)$

    3. Update θ_R with Adam optimizer on mini-batch loss $L_R(f;\theta_R)_{(x,\hat{x})}$: $\min L_R$

    4. Update θ_U with Adam optimizer on mini-batch loss $L_U(f;\theta_U)_{(y,\hat{y})}$: $\min L_U$

    5. Update θ_P with Adam optimizer on mini-batch loss $L_P(f;\theta_P)_{(z,\hat{z})}$: $\min L_P$

  end

  Update with the gradient descent on $L_{sum}(\theta_E,\theta_R,\theta_U,\theta_P,\lambda_1,\lambda_2,\lambda_3)$: $\min L_{sum}$

end

The overall training is to achieve a privacy-utility trade-off by adversarial learning on $\mathcal{L}_\mathcal{R}$, $\mathcal{L}_\mathcal{U}$, and $\mathcal{L}_\mathcal{P}$, concurrently. The encoder $Enc_L(\mathcal{X};\theta_E)$ should satisfy high predictability (*min* $\mathcal{L}_\mathcal{U}$) and low user re-identification accuracy (*max* $\mathcal{L}_\mathcal{P}$) of the mobility data when maximizing the reconstruction loss (*max* $\mathcal{L}_\mathcal{R}$) in reverse engineering, where the training objective transformed from Eq.1 can be written as:

$$\min \mathcal{L}_{sum} = \min_{\mathcal{L}_U} \max_{\mathcal{L}_R,\mathcal{L}_P} \left(\sum_{x=i}^{\mathcal{X}} \left(\mathcal{L}_U(f_i), \mathcal{L}_P(f_i), \mathcal{L}_R(f_i)\right)\right) \tag{5}$$

We use Eq. 5 to guide the first version of Mo-PAE, denoted as *Model I*. In order to fully investigate the range of trade-offs, we leveraged the Lagrange multipliers [4] as hyperparameters to control the privacy-utility trade-offs in the Mo-PAE, and this weighted-controlled model is denoted as *Model II*. Accordingly, the optimization function of the training objective is:

$$\begin{aligned}
\min \mathcal{L}_{sum} &= \min_{\mathcal{L}_U} \max_{\mathcal{L}_R,\mathcal{L}_P} \left(\sum_{x=i}^{\mathcal{X}} \left(\lambda_1\mathcal{L}_R(f_i), \lambda_2\mathcal{L}_U(f_i), \lambda_3\mathcal{L}_P(f_i)\right)\right) \\
&= \underbrace{-\lambda_1\left(\max \mathcal{L}_\mathcal{R}(f_i)\right)}_{\text{Privacy I}} + \underbrace{\lambda_2\left(\min \mathcal{L}_\mathcal{U}(f_i)\right)}_{\text{Utility}} \underbrace{-\lambda_3\left(\max \mathcal{L}_\mathcal{P}(f_i)\right)}_{\text{Privacy II}} \\
&= -\lambda_1\|Dec_L(\mathcal{F}) - \mathcal{X}\|^2 + \lambda_2\left(\sum_{i=1}^{Y} y_i \log(U_D(\mathcal{F}))\right) \\
&\quad - \lambda_3\left(\sum_{i=1}^{Z} z_i \log(P_D(\mathcal{F}))\right)
\end{aligned} \tag{6}$$

where y_i is the ground-truth label for *Utility*, z_i is the ground-truth value for *Privacy II*; λ_1, λ_2 and λ_3 are non-negative, real-valued weights, as the hyperparameters that control the privacy-utility trade-off in the Mo-PAE.

Dataset-City	Bounding Box				Record Counts		Number	
	Latitude		Longitude		Train	Test	User ID	POI
MDC-Lausanne	46.50	46.61	6.58	6.73	77393	19429	143	149
Priva'Mov-Lyon	45.70	45.81	4.77	4.90	62077	16859	58	129
GeoLife-Beijing	39.74	40.07	116.23	116.56	95038	24578	145	960
FourSquare-NYC	40.55	40.99	-74.28	-73.68	43493	11017	466	1712

Table 1: Overview of four mobility datasets after pre-processing. The bounding box represents the range of the considered locations/traces.

As shown in the Algorithm 1, the gradient of the loss (i.e., θ_E, θ_R, θ_U, θ_P) back-propagates through the LSTM network to guide the training of the encoder Enc_L. The encoder is updated with the *sum loss* function $\mathcal{L}_{sum}$ until convergence. It is tricky to investigate all possible weight combinations practically, and we look for the optimal combinations through training [39] with the Eq.6 by brute-force evaluation. Then we approximate the required data utility reserved and reformulate the optimization problem in Eq.6 as a maxima privacy optimization problem.

$$\min_{Enc_L(P_D^1, P_D^2)} \max V_{\lambda \to U_D}(Enc_L, P_D^1, P_D^2) = \mathbb{E}_{x \sim X}[log(1 - P_D^1(Enc_L(x)))] + \mathbb{E}_{x \sim X}[log(1 - P_D^2(Enc_L(x)))] \tag{7}$$

Additionally, another key contribution is the flexibility of the *sum loss* function $\mathcal{L}_{sum}$, which could be regulated to satisfy different requirements on privacy protection level and service quality. That is, different combinations of weights control the relative importance of each unit and guide the overall model to find the maxima or minima given the specific trade-off choices.

4 EXPERIMENTAL SETTING

4.1 Datasets

Experiments are conducted on four representative mobility datasets: Mobile Data Challenge Dataset (MDC) [36], Priva'Mov [5], GeoLife [62], and FourSquare [59].

MDC: it is recorded from 2009 to 2011 and contains a large amount of continuous mobility data for 184 volunteers with smartphones running a data collection software in the Lausanne/Geneva area. Each record of the *gps-wlan* dataset represents a phone call or an observation of a WLAN access point collected during the campaign [36].

Priva'Mov: the PRIVA'MOV crowd-sensing campaign took place in the city of Lyon/France, from October 2014 to January 2016. Data collection was contributed by roughly 100 participants, including university students, staff, and family members. The crowd-sensing application collected GPS, WiFi, GSM, battery, and accelerometer sensor data. For this paper, we only used the GPS traces from the dataset [5].

GeoLife: it is collected by Microsoft Research Asia from 182 users in the four and a half year period from April 2007 to October 2011 and contains 17,621 trajectories [62]. This dataset recorded a broad range of users' outdoor movements, including life routines like going home and going to work and some entertainment and sports activities, such as shopping, sightseeing, dining, hiking, and cycling. It is widely used in many research fields, such as mobility pattern mining, user activity recognition, location-based social networks, location privacy, and location recommendation.

FourSquare NYC: it contains check-ins in NYC and Tokyo collected during the approximately ten months from 12 April 2012 to 16 February 2013, containing 227,428 check-ins from 1,083 subjects in New York City [59].

Once imported into our architecture, each dataset was filtered and pre-processed individually to derive their respective train and test sets illustrated in Table 1. We filter locations to a bounding box defining a city or region of interest and then transform continuous GPS coordinates by tessellating the space and encoding location as a discrete grid position to attain the location identifiers (i.e., POI). In these spatial transformations, we convert the GPS coordinates to the discretizing locations via the Geohash algorithm [41] with rectangular cells. For instance, each bounding box defines the grid size of the interested region, and the grid granularity is 0.01 degrees, where each grid represents a 0.01 longitude x 0.01 latitude area.

4.2 Baseline Models

4.2.1 I. Optimal Inference Models (Optimal-IMs). Optimal-IMs comprise three independent inference models, namely data reconstruction model, mobility prediction model, and user re-identification model. Each model has a similar layer design as the corresponding unit in the Mo-PAE, however, these three models are completely independent and have no effect on each other. Unlike the Mo-PAE, which leverages adversarial learning to finally attain an extracted feature representation f that satisfies the utility requirements and privacy budgets simultaneously, the Optimal-IMs are only trained for optimal inference accuracy at the individual tasks to characterize the original data.

4.2.2 II. LSTM-TrajGAN (TrajGAN) [48]. It is an end-to-end deep learning model to generate synthetic data which preserves essential spatial, temporal, and thematic characteristics of the real trajectory data. Compared with other standard geo-masking methods, TrajGAN can better prevent users from being re-identified. The TrajGAN work claims to preserve essential spatial and temporal characteristics of the original data, verified through statistical analysis of the generated synthetic data distributions, which aligns with the mobility prediction-based utility metric in our work. Hence, we train an optimal mobility prediction model for each dataset and evaluate the mobility predictability of synthetic data generated by the TrajGAN. In contrast to the TrajGAN that aims to generate synthetic data, our proposed Mo-PAE is training an encoder Enc_L

that forces the extracted representations f to convey maximal utility while minimizing private information about user identity via adversarial learning.

4.2.3 III. GI-DP [8]. The principle of geo-indistinguishability (i.e., GI) [2], is a formal notion of privacy that protects the user's exact location with a level of privacy that depends on radius r, which corresponds to a generalized version of differential privacy (DP). GI-DP is a mechanism for achieving geo-indistinguishability when the user releases his location repeatedly throughout the day. It fulfils desired protection level by perturbing the actual location with random noises and achieving an optimal trade-off between privacy and utility (i.e., service quality). We re-implement the geo-indistinguishability of optimal utility with graph spanner [8], namely GI-DP in this paper, to attain the released version data that satisfied the DP guarantees. We then train a series of Optimal-IMs to evaluate the effectiveness of target inference attacks on the released version data in a line to compare with our proposed mechanism.

4.3 Training

4.3.1 Training of Mo-PAE. The main goal of the proposed adversarial network is to learn an efficient feature representation based on the utility and privacy budgets, using all users' mobility histories. In most experiments in this work, the trajectory sequences consist of 10 historical locations with timestamps (i.e., $SL = 10$), and the impact of the varying sequence lengths is discussed in Section 6.2. After data pre-processing, 80% of each user's records are segmented as the training set and the remaining 20% as the testing set. We utilize the mini-batch learning method with the size of 128 to train the model until the expected convergence. We take a gradient step to optimize the *sum loss* L_{sum} (i.e., Equation 6) in terms of L_R, L_U, and L_P concurrently. Meanwhile, the *sum loss* L_{sum} is optimized by using the Adam optimizer. All the experiments are performed with the Tesla V100 GPU; a round of training would take 30 seconds on average, and each experiment trains for 1000 rounds.

4.3.2 Training of the TrajGAN. To provide a state-of-the-art machine learning-based model for comparison, we re-implement the TrajGAN model described in [48] using the same hyperparameters, setting latent vector dimension to 100, using 100 LSTM units per layer, a batch size of 256, utilizing the Adam optimizer with learning rate 0.001 and momentum 0.5, and training for 200 epochs (where one epoch is a pass through the entire training set). We train TrajGAN independently on the training split of each benchmark mobility dataset, and then use it to generate synthetic trajectories from the test set. Then we train the proposed Mo-PAE on the same training data and use it to generate a feature extraction from the same test data. Finally, we evaluate the performance of the user re-identification unit and mobility prediction unit on the real and synthetic test sets generated by TrajGAN, and compare the changes in accuracy to assess the relative utility and privacy of the TrajGAN and Mo-PAE.

4.3.3 Training of the DP-GI. We re-implement the DP-GI model described in [8] using the default settings. That is, we set *epsilon* $\epsilon = 0.5$, *dilation* $\delta = 1.1$, the *distance matrix* d_x is defined by Euclidean distance. From [8], let X be a set of locations with metric d_x, and let $G(X, E)$ be a $\delta - spanner$ of X, if a mechanism K for X is $\frac{\epsilon}{\delta} d_G$-private, then K is ϵd_x-private. The dilation of G is calculated as:

$$\delta = \max_{x \neq x' \in X} \frac{d_G(x, x')}{d_x(x, x')} \tag{8}$$

$$d_G(x, x') \geq d_x(x, x') \qquad \forall x, x' \in X \tag{9}$$

We re-implement the GI-DP to attain the released version data that satisfied the DP guarantees. We then train a series of Optimal-IMs to evaluate the effectiveness of target inference attacks on the released version data in a line to compare with our proposed mechanism.

4.4 Metrics

We set *Euclidean* [3] and *Manhattan* distance [7] as our evaluation metrics for the DRU to evaluate the quality of the reconstructed data X' generated from extracted features f. Both distance metrics are widely used in location privacy literature [2, 16]. For instance, the work introducing Geo-Indistinguishability [2] utilizes a privacy level that depends on the *Euclidean* distance. *Euclidean* distance gives the shortest or minimum distance between two points. In contrast, *Manhattan* distance applies only if the points are arranged in a grid, and both definitions are feasible for the problem we are working on. Note that these two distances have limited capability in showing the quality of the reconstructed data X', however, they intuitively capture the differences between the original data X and the reconstructed data X'.

For both MPU and URU, we leverage the top-n accuracy as our evaluation metric. The accuracy of the MPU is one of the most important factors in evaluating the utility of the extracted feature representation *f*, where predictability of the f increases as much as it can during the adversarial training. On the other hand, the competing training objective is to decrease the accuracy of the user re-identification unit to enhance the privacy of *f*. The top-n metric computes the number of times the correct label appears among the predicted top n labels. The top-n metric takes n predictions with higher probability into consideration, and it classifies the prediction as correct if one of them is an accurate label. The top-1 to top-5 accuracies are leveraged in our paper to discuss the performance of the proposed model.

5 ARCHITECTURE EVALUATION

In this section, we present the comparison results between the proposed Mo-PAE and three baseline models under the same training setting. Our evaluation is mainly on the trade-off between U and PII, as the main contribution of the Mo-PAE is to protect users' identities, while we also consider the scope of users' locations as an auxiliary measurement.

5.1 Performance Comparison

We first compare our proposed models with the Optimal-IMs, TrajGAN, and DP-GI on four representative mobility datasets, as shown in Table 2. The overall performance is evaluated in terms of the *utility level* provided by the MPU and the *privacy protection* provided by DRU and URU. The *Model I* is our proposed architecture but without applying the Lagrange multipliers (i.e., where each unit is weighted equally). The *Model II* is the one with Lagrange multipliers (i.e., λ_1, λ_2, λ_3) and different weights are given to units

Datasets	Models		Privacy I		Utility (% for loss)			Privacy II (% for gain)			Utility-PII trade-offs (%)
			Euc	Man	top-1	top-3	top-5	top-1	top-3	top-5	
MDC	Optimal-IMs		0.0000	0.0000	0.9347	0.9837	0.9922	0.9247	0.9819	0.9911	-
	TrajGAN		0.0434	3.6923	-46.32%	-24.16%	-15.98%	+20.32%	+8.13%	+4.02%	-26.00%
	GI-DP		**0.2341**	**56.8764**	-97.34%	-93.25%	-89.44%	**+97.47%**	**+93.71%**	**+90.12%**	0.13%
	Our Model	I	0.0025	0.4501	-54.56%	-34.74%	-25.10%	+69.80%	+50.44%	+39.95%	15.24%
		II	0.0697	13.6168	**-13.43%**	**-6.26%**	**-3.95%**	+65.51%	+45.11%	+34.86%	**52.08%**
Priva'Mov	Optimal-IMs		0.0003	0.0058	0.9482	0.9878	0.9954	0.5643	0.8215	0.8765	-
	TrajGAN		0.0815	9.6843	-6.60%	-1.89%	-0.93%	+14.17%	+14.35%	+8.88%	7.57%
	GI-DP		0.1899	**38.6712**	-91.20%	-83.53%	-72.37%	**+85.49%**	**+63.80%**	**+53.31%**	-5.71%
	Our Model	I	0.0009	0.0437	**-3.36%**	**-1.59%**	**-0.81%**	+27.02%	+14.19%	+9.19%	23.66%
		II	**0.2347**	10.2239	-10.81%	-6.83%	-4.91%	+35.29%	+14.97%	+10.05%	**24.48%**
Geolife	Optimal-IMs		0.0008	0.0670	0.4705	0.6842	0.7636	0.6572	0.8690	0.9294	-
	TrajGAN		0.4010	50.3620	-62.31%	-50.45%	-43.72%	+66.73%	+47.89%	+37.22%	4.42%
	GI-DP		**1.2332**	**312.9972**	-97.74%	-96.56%	-95.36%	**+91.57%**	**+84.13%**	**+78.65%**	-6.17%
	Our Model	I	0.0006	0.0310	-31.45%	-25.02%	-21.90%	+54.88%	+39.59%	+30.81%	23.43%
		II	0.4351	89.2209	**-21.13%**	**-18.78%**	**-17.11%**	+55.49%	+40.40%	+32.34%	**34.36%**
FourSquare	Optimal-IMs		0.0052	0.6691	0.6468	0.8210	0.8823	0.8780	0.9735	0.9892	-
	TrajGAN		**1.4341**	**117.9181**	-26.30%	-22.30%	-18.75%	+51.86%	+32.49%	+23.49%	25.56%
	GI-DP		0.5826	86.096	-69.35%	-59.23%	-53.36%	**+77.29%**	**+66.58%**	**+59.82%**	7.94%
	Our Model	I	0.0060	0.7845	-51.05%	-41.45%	-35.20%	+53.47%	+35.26%	+25.86%	2.42%
		II	0.7985	99.9212	**-2.54%**	**-3.14%**	**-2.84%**	+51.08%	+34.39%	+26.16%	**48.54%**

Table 2: Performance comparison between Mo-PAE with other baseline models. The *Model I* is our proposed architecture without weights, and the *Model II* is the one with multipliers (λ_1 = 0.1, λ_2 = 0.8, and λ_3 = 0.1). The results shown in this table are all with trace sequence length 10 (i.e., *SL* = 10). The *Privacy I* intuitively shows the difference between the raw data and reconstructed data; the *Utility* (%) represents the utility loss; and the *Privacy II*(%) represents the privacy gain calculated via the decline of the user re-identification accuracy.

(i.e., $\lambda_1 = 0.1,\ \lambda_2 = 0.8,\ \lambda_3 = 0.1$ for the results in Table 2). In this table, the sequence length of the input traces is 10, that is $SL = 10$. We will discuss why we choose SL=10 and the impact of the SL in Section 6.

As we mention in Section 4.2, Optimal-IMs are trained without considering the privacy-utility trade-offs; hence, they can be leveraged to explain the optimal inference accuracy achieved. That is, before any privacy-preserving mechanism applies, the accuracy of the target or private tasks with raw data. For instance, the accuracy of the *Privacy II* (0.9247 (MDC), 0.5643 (Priva'Mov), 0.6572 (GeoLife), 0.8780 (FourSquare)) demonstrates that an adversary can accurately infer user identity from raw data before any privacy protection.

Different from Optimal-IMs, the other models consider privacy-utility trade-offs, and we measure the privacy protection and data utility by the effectiveness of the inference units. First, for the *Privacy I*, the distance indexes (i.e., *"Euc"* and *"Man"*) are leveraged to intuitively represent the difference between the original data X and reconstructed data X', where a larger value indicates numerical differences between X and X'. For the distance index, we are interested in the distance between each trace, hence, we consider the quantity of trace for datasets X_D and get these distance indexes by averaging the corresponding record numbers N_D, that is (take *"Euc"* for example):

$$Euc(X_D, X'_D) = \frac{\sqrt{\sum_{i=1}^{N_D}(x_i - x'_i)^2}}{N_D}, \qquad N_D = N'_D \tag{10}$$

Different to the *Privacy I*, the *Utility loss* and *Privacy II gain* are in a percentage format (%), compared with the accuracy of Optimal-IMs. To compare the trade-off between them more intuitively, we list the "Utility-PII trade-off" column, where "trade-offs = Utility (% for loss) + Privacy II (% for gain)". Table 2 demonstrates that our proposed models, especially *Model II*, outperform the TrajGAN and GI-DP across various datasets. For instance, with the MDC dataset, our *Model II* achieves the best trade-offs when compared with other models, as the utility loss is only 13.43% but with 65.51% privacy gain, while 46.32% utility loss and 20.32% privacy gain with the TrajGAN, and 97.34% utility loss and 97.47% privacy gain with the GI-DP. The extreme performance on the GI-DP illustrates that while the DP paradigm is a robust privacy-preserving technique in protecting *user's location*, it is not appropriate in protecting the *user's identities.*

More intuitively, in the column of "trade-off", *Model II* achieves all the best trade-offs among four datasets (52.08% (MDC), 24.48% (Priva'Mov), 34.36% (GeoLife) and 48.54 (FourSquare)). *Model I* has worse performance than *Model II*, in general, but is still superior to TrajGAN and DP-GI, where the latter two might even get *negative* trade-offs (i.e., TrajGAN got -26.00% with MDC and GI-DP got -5.71% with Priva'Mov). Moreover, for the Priva'Mov dataset, although the utility loss of the TrajGAN is 4.21% smaller than our *Model II*, both two privacy metrics of the TrajGAN are worse than the *Model II*. Again, our model has better overall trade-offs, as 23.66% for *Model I* and 24.48% for *Model II*. The performance on Geolife and FourSquare are similar but inverse, where the utility of our

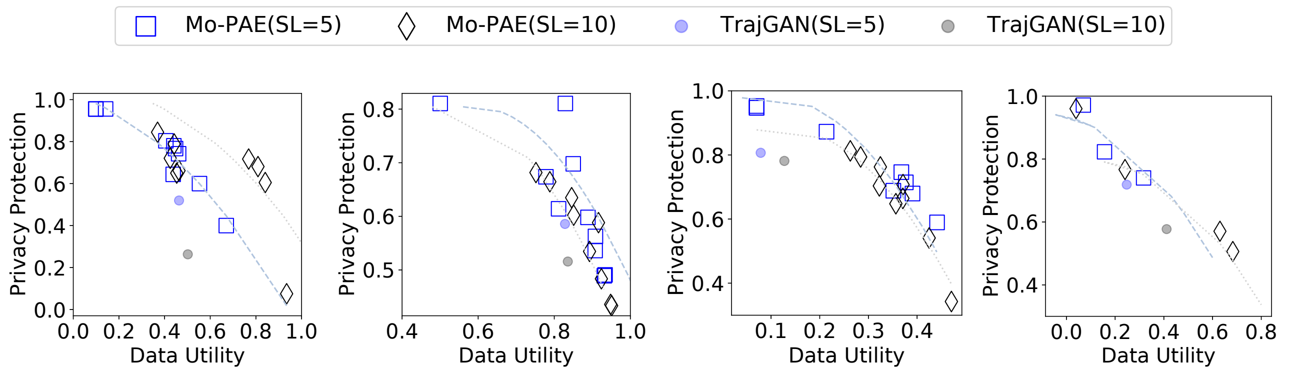

(a) MDC (b) PrivaḾov (c) Geolife (d) Foursquare

Figure 3: Pareto Frontier trade-off analysis on four datasets. The hollow squares and diamonds present the results of the proposed models Mo-PAE. solid points present the results of the TrajGAN. Blue color means $SL = 5$. Black color means $SL = 10$.

model is better than TrajGAN and with slightly weaker privacy preservation.

In summary, GI-DP always has the highest privacy gain among the four datasets, however, the utility loss is also very high, resulting in inadequate and unexpected privacy-utility trade-offs. This trend also shows that the DP mechanism is not an appropriate metric for the location privacy of *user's identity*, which is also in line with the conclusions from other related work [16, 51]. The comparisons between *Model I* and *Model II* also illustrate the importance of the Lagrange multipliers, which provides flexibility to our proposed architecture, enabling its application in different scenarios and enhancing the privacy-utility trade-offs in this case.

5.2 Trade-off Comparison

In this section, we present the privacy-utility trade-off analysis between the proposed Mo-PAE and TrajGAN in terms of mobility prediction accuracy (i.e., *U*) and user de-identification efficiency (i.e., *PII*). Figure 3 presents the trade-off comparisons of the four datasets, where the *hollow squares* and *hollow diamonds* show the trade-offs provided by the proposed Mo-PAE in $SL = 5$ and $SL = 10$, respectively. The *solid points* present the results of the TrajGAN under the same experimental setting. As can be seen from these results, in all four cases, the synthetic dataset generated by the TrajGAN is not *Pareto-optimal*. That is, the proposed Mo-PAE is able to achieve a better privacy level for a dataset with the same utility value. Compared with the TrajGAN, Mo-PAE improves utility and privacy simultaneously on four datasets. Especially for the performance of MDC, the privacy improves 45.21% more than the TrajGAN, while the utility also increases by 32.89%. These results illustrate that our proposed model achieves promising performance in training a privacy-sensitive encoder Enc_L for different datasets.

After evaluating the superior performance of our proposed model, we discuss the privacy guarantee that Mo-PAE provided in terms of data reconstruction (*PI*, *"Euc"* in Table 3) and user re-identification (*PII*, *privacy gain* in Figure 4). As we mentioned in Section 2, the privacy guarantee of Mo-PAE differs from that of DP paradigms and is given in the declined effectiveness of inference attacks.

5.3 Privacy Guarantee Analysis: Effectiveness of Privacy Inference Attacks

In this section, we discuss the impact of Mo-PAE on the effectiveness of two privacy inference attacks (i.e., *PI* and *PII*), respectively.

5.3.1 Effectiveness of Data Reconstruction Attacks - PI. Table 3 shows the impact of the proposed mechanisms on the data reconstruction accuracy(*PI*). The *"Euc"* in the table follows the definition in Eq.10. Overall, *Model II* performs better than *Model I* in limiting the accuracy of data reconstruction regardless of the value of weights. Take the result of GeoLife dataset as an example, *Model II-i* achieves bigger distance than *Model I* (i.e., 0.4343 > 0.0057), while it still gets better utility (i.e., -9.94% > -17.9%). Nevertheless, both *Model I* and *Model II* have effectively defended the data reconstruction attack (MDC: 0.0697 > 0.0017 > 0.0000; Priva'Mov: 0.0453 > 0.0009 > 0.0003; GeoLife: 0.4343 > 0.0057 > 0.0008; FourSquare: 0.7933 > 0.0069 > 0.0052), while only at the marginal cost of mobility utility (MDC: -12.55%; Priva'Mov: -2.71%; GeoLife: -9.94%; FourSquare: -1.64%). The data of the Optimal-IMs are in Table 2. We list four representative settings here to make a comprehensive comparison of PII and U. From setting i to iv, one can expect more original data features loss to result in a more significant utility loss. This trend is indeed the case with different weights' combinations. However, as the results show, especially for setting i, the privacy of the traces attains decent protection at the marginal cost of mobility utility.

5.3.2 Effectiveness of User Re-identification Attacks - PII. Figure 4 presents the impact of the *Model II* on the user re-identification accuracy(*PII*). In this figure, we list five different settings, I to V ($\lambda_1 = 0.1$, over the range of $\lambda_2 = \{0.5, 0.6, 0.7, 0.8, 0.9\}$), respectively. The *Zero* line (i.e., y = 0%) in each sub-figure is leveraged to indicate the original accuracy of the raw data (i.e., Optimal-IMs). The *"Privacy Gain Rate"* (blue square line) shows the effectiveness of the Mo-PAE in defending the user re-identification attacks. That is, after applying *Model II*, the decline range of effectiveness of user re-identification attacks. For instance, with the MDC dataset, in setting I, the effectiveness of user re-identification attacks declines as

Settings		λ_1	λ_2	λ_3	MDC Euc	MDC Utility	Priva'Mov Euc	Priva'Mov Utility	Geolife Euc	Geolife Utility	FourSquare Euc	FourSquare Utility
Model I	-	-	-	-	+0.0017	-30.27%	+0.0009	-2.72%	+0.0057	-17.9%	+0.0069	-33.75%
Model II	i	0.1	0.8	0.1	+0.0697	-12.55%	+0.0453	-2.71%	+0.4343	-9.94%	+0.7933	-1.64%
	ii	0.2	0.6	0.2	+0.0791	-33.29%	+0.0738	-10.72%	+0.4889	-18.21%	+1.2722	-50.50%
	iii	0.3	0.4	0.3	+0.0889	-58.10%	+0.0782	-16.56%	+0.5220	-29.95%	+1.9586	-60.71%
	iv	0.1	0.6	0.3	+0.0822	-49.27%	+0.0776	-10.28%	+0.4717	-18.64%	+1.4139	-57.40%

Table 3: Impact of Mo-PAE on the data reconstruction accuracy (*PI*) and relative utility loss (*U*) on four mobility datasets. We list *Model I* and four different settings of *Model II*'s weight combinations to discuss the potential range of the trade-offs.

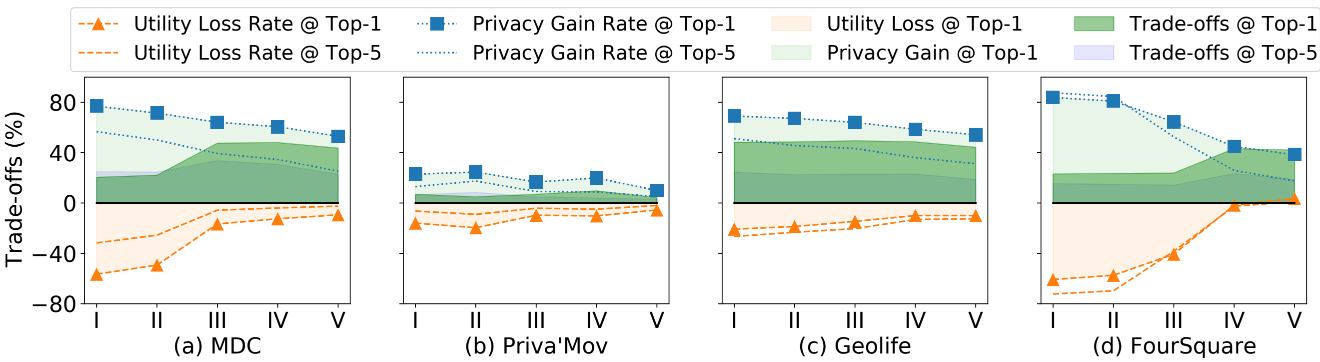

Figure 4: Impact of Mo-PAE on the user re-identification accuracy (PII) and relative utility loss (*U*) on four datasets. The orange area represents the utility loss, while the light-green area represents privacy gain. The dark-green area represents the trade-offs between utility achievement and privacy budgets. The x-axis shows five different model settings, and the y-axis shows the trade-offs.

high as 80%. At the same time, this high privacy protection is at the cost of nearly 55% of utility (orange triangle line). Things are better in setting V, where the Mo-PAE can get 60% privacy protection only at the cost of less than 10% utility. The x-axis shows five settings of the model, and the y-axis shows the trade-offs (i.e., *trade-offs* = *privacy gain* + *utility loss*). The orange area represents the utility loss while the light-green area represents the privacy gain when compared with Optimal-IMs. The dark-green area represents the trade-offs between utility and privacy budgets.

In summary, these trade-offs are all positive in different model settings on four different datasets. The performance on the GeoLife data is the best, while less than 20% utility loss but more than 50% privacy gains. The performance on MDC and FourSquare also show promising privacy-utility trade-offs, especially for setting V on the FourSquare dataset, and both the utility and privacy increase. The uniqueness of human mobility trajectories is high, and these trajectories are likely to be re-identified even with a few location data points [13]. Our results emphasize that the concern of user re-identification risk could be alleviated effectively with our proposed model.

In real applications, the trade-off of Mo-PAE is achieved continuously over time. New trajectories will be encoded with the pre-trained encoder to attain respective feature representation and utilized by SP for following task-oriented scenarios (no need to retrain). The pre-trained encoder and discriminators are assumed to be updated offline within a fixed duration for best performance purposes. Additionally, while the architecture focuses on specific application scenarios (i.e., mobility prediction), it could generally be applicable to different task-oriented scenarios.

6 DISCUSSIONS

In this section, we further discuss the impact of the temporal granularity of traces, the varying sequence length and weights on the composition units on the Mo-PAE performance.

6.1 Impact of Temporal Granularity

The timestamp is one of the essential components of the mobility trace, and different choices on the temporal granularity affect the final performance of any dataset. Figure 5 shows the impact of the varying temporal granularity on the proposed architecture. We present the top-1, top-5, and top-10 accuracies for both utility and privacy dimensions. For instance, when temporal granularity is 10-min, it indicates a location record r is taken every 10 minutes from the raw data. When using more coarser temporal granularity, the number of points of interest decreases so does the difficulty of mobility prediction. However, the uniqueness of the trajectory decreases due to ignoring many of the unique locations of each user, resulting in better privacy. To summarize from Figure 5, the impact of temporal granularity on the Priva'Mov is minimal. In terms of utility (mobility prediction), Priva'Mov is the only dataset for which accuracy decreases with increasing temporal granularity. This subtle decline emphasizes the trajectory features only have a

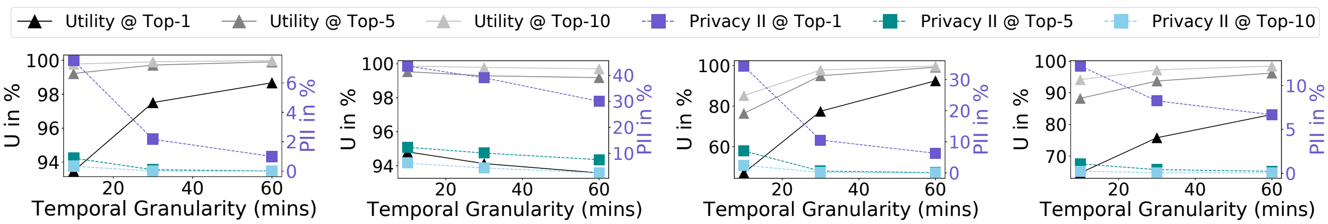

(a) MDC (b) PrivaḾov (c) Geolife (d) Foursquare

Figure 5: The effect of temporal granularity on the model performance of four mobility datasets.

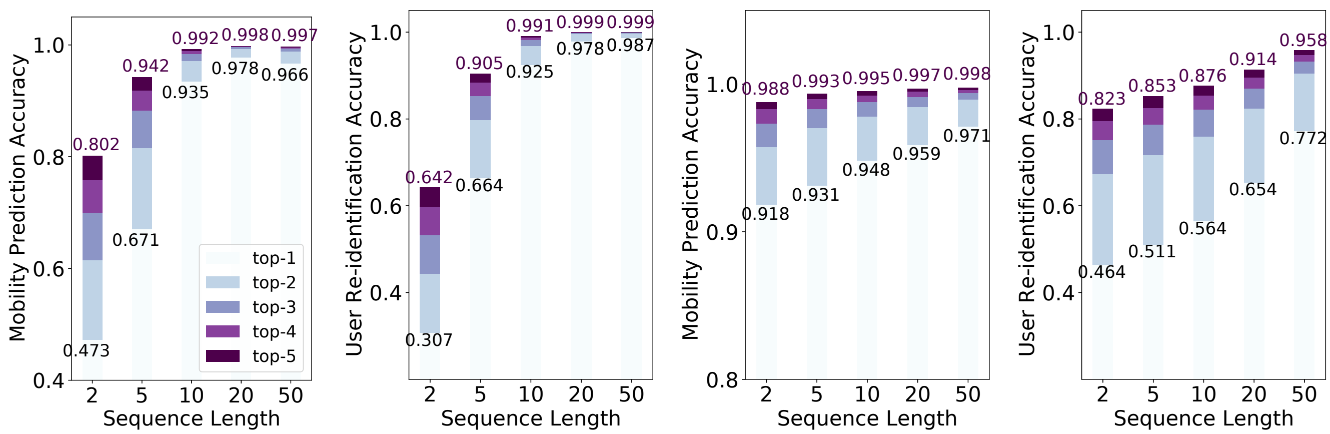

(a) U_D Performance on MDC (b) P_D^2 Performance on MDC (c) U_D Performance on Priva'Mov (d) P_D^2 Performance on Priva'Mov

Figure 6: Mobility prediction accuracy and user re-identification accuracy change with the trace sequence length (SL) in our proposed U_D and P_D^2. The color bars indicate the accuracy from top-1 to top-5, the black texts indicate the top-1 accuracy, and the purple texts indicate the top-5 accuracy. For instance, the top-1 mobility prediction accuracy on MDC with *SL* = 2 is 0.473, and the top-5 one is 0.802.

small change when varying granularity, in line with the university students' mobility.

6.2 Impact of Varying Sequence Lengths

The performance of the utility discriminator U_D (MPU) and the privacy discriminator P_D^2 (URU) exert a significant impact on the overall performance of the proposed Mo-PAE. The trace length is the most critical factor affecting these units' performance. We use two representative datasets (i.e., MDC and Priva'Mov) to present the impact of the varying sequence length on both discriminators.

As shown in Figure 6, by changing the lengths of trace sequence SL from 1 ($SL = 1$) to 50 ($SL = 50$), we observe that SL has a significant impact on different tasks' accuracy (i.e., mobility prediction accuracy for U_D and user re-identification accuracy for P_D^2) of two different datasets. Overall, the impact in the MDC dataset is much higher than in the Priva'Mov dataset. Comparing the Figure 6a and Figure 6c, there is a much sharper increase on the MDC dataset. More specifically, when the sequence length is increased from 2 to 20, the top-1 mobility prediction accuracy on MDC increases from 0.473 to 0.978 (i.e., +50.5%), while accuracy on Priva'Mov increases from 0.918 to 0.959 (i.e., only +4.1%). Similarly, more rapid growth appears in the user re-identification accuracy on MDC, which is +68.0%, while the increase for Priva'Mov is only +30.8%. We conclude that the mobility predictability and user re-identification accuracy of a dataset might have a special link. The mobility predictability of Priva'Mov is very high, almost higher than 90%, but the user re-identification accuracy is always lower than 80%, which also means the uniqueness of trajectories in this dataset is low. This low uniqueness suggests that the users in this dataset might share similar daily routes, which is reasonable, as we know these trajectories are collected from students at the same university. For the MDC dataset, when $SL = 10$, the user re-identification accuracy is relatively high, indicating that the locations are more sparse in this dataset. However, the mobility predictability here is also high, which emphasizes that this sparseness does not affect predictability. These phenomena indicate that the deep training of MPU and URU might share similar extracted features, while our proposed architecture attempts to extract features more suitable for mobility predictability but less suitable for user re-identification.

We note that the varying trace sequence length not only exerts impacts on the model performance but also has a significant influence on the computation time. For instance, the computation

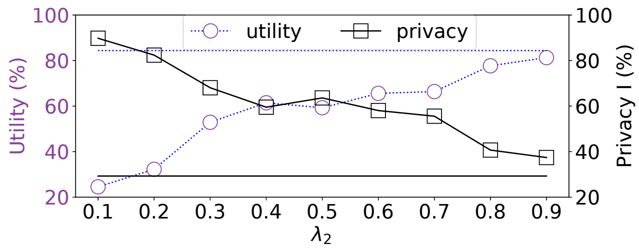

(a) $\lambda_1 = 1 - \lambda_2,\ \lambda_3 = 0$

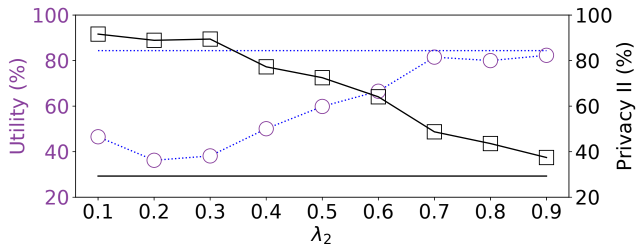

(b) $\lambda_3 = 1 - \lambda_2,\ \lambda_1 = 0$

Figure 7: Varying weights can tune the privacy-utility trade-offs. The primary y-axis (dashed line) represents utility, and the secondary y-axis (solid line) represents privacy. The x-axis represents the value of the target λ.

time at $SL = 50$ costs six times as much as that at $SL = 5$. The computation time also varies between datasets. Hence, an appropriate choice of trajectory sequence length can avoid time-consuming computation and achieve expected task inference accuracy. In our work, we place greater focus on the trace sequence lengths ranging from 5 to 10, which exhibits great performance in both the U_D and P_D while also keeping a low computation time.

6.3 Impact of Varying Weights

As we discussed in Section 3.3, the *sum loss function* L_{sum} of *Model II* is a linear combination of L_R, L_U, and L_P with different weights (i.e., Lagrange multipliers). We evaluate the influence of different weights' combinations (λ_1, λ_2, and λ_3) on the *Model II*, as the results shown in Figure 7.

We compare the overall model performance in U_D and P_D^1 by fixing the $\lambda_3 = 0$, and vary the other two multipliers by subjecting to $\lambda_1 = 1 - \lambda_2$, as shown in Figure 7a. Figure 7b illustrates the effect between U_D and P_D^2 by setting the $\lambda_1 = 0$. We could observe in both settings that the utility increases with a larger λ_2, which means when the MPU is given more weight in the Mo-PAE model, it would exert a positive impact on the data utility. We conclude that the privacy-utility trade-offs could be tuned by varying these weights; the results in Figure 7 also verify the effectiveness of our adversarial architecture. We note that the balance of three units is far more complicated than the balance of two. From the extensive experiment we conducted, initialing $\lambda_1 = 0.1$, $\lambda_2 = 0.6$, $\lambda_3 = 0.3$ can guide the model achieve the trade-off most efficiently. However, as the experiment results show, there is no dataset independent privacy interpretation for λ_1, λ_2 and λ_3, and we leave a more efficient approach using reinforcement learning to initialise these hyperparameters for different datasets in future work.

7 CONCLUSION

In this paper, we presented a privacy-preserving architecture Mo-PAE based on adversarial networks. Our model considers three different optimization objectives and searches for the optimum trade-off for utility and privacy of a given dataset. We reported an extensive analysis of our model performances and the impact of its hyperparameters using four real-world mobility datasets. The weights λ_1, λ_2, and λ_3 bring more flexibility to our framework, enabling it to satisfy different scenarios' requirements according to the relative importance of utility requirements and privacy budgets. We evaluated our framework on four datasets and benchmarked our results against an LSTM-GAN approach and a DP mechanism. The comparisons indicate the superiority of the proposed framework and the efficiency of the proposed privacy-preserving feature extractor Enc_L. Expanding this work, we will consider other utility functions for our models, such as community detection based on unsupervised clustering methods or deep embedded clustering methods. In future work, we will leverage automated search techniques, such as deep deterministic policy gradient algorithm and reinforcement learning, for efficiency in searching for optimal weight combinations.

ACKNOWLEDGMENTS

We wish to acknowledge the constructive suggestions from PETS reviewers. The authors are grateful to Imperial College London and the China Scholarship Council (CSC) for financial support. This work was partially supported by EPSRC Open Plus Fellowship EP/W005271/1.

REFERENCES

[1] Ahmet Aktay, Shailesh Bavadekar, Gwen Cossoul, John Davis, Damien Desfontaines, Alex Fabrikant, Evgeniy Gabrilovich, Krishna Gadepalli, Bryant Gipson, Miguel Guevara, Chaitanya Kamath, Mansi Kansal, Ali Lange, Chinmoy Mandayam, Andrew Oplinger, Christopher Pluntke, Thomas Roessler, Arran Schlosberg, Tomer Shekel, Swapnil Vispute, Mia Vu, Gregory Wellenius, Brian Williams, and Royce J Wilson. 2020. Google COVID-19 Community Mobility Reports: Anonymization Process Description (version 1.1). , arXiv:2004.04145 pages. https://doi.org/10.48550/ARXIV.2004.04145

[2] Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geo-Indistinguishability: Differential Privacy for Location-Based Systems. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security* (Berlin, Germany) *(CCS '13)*. Association for Computing Machinery, New York, NY, USA, 901–914. https://doi.org/10.1145/2508859.2516735

[3] Walter William Rouse Ball. 1960. *A short account of the history of mathematics.* Courier Corporation, New York, NY, USA.

[4] Brian Beavis and Ian Dobbs. 1990. *Optimisation and Stability Theory for Economic Analysis.* Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9780511559402

[5] Sonia Ben Mokhtar, Antoine Boutet, Louafi Bouzouina, Patrick Bonnel, Olivier Brette, Lionel Brunie, Mathieu Cunche, Stephane D 'Alu, Vincent Primault, Patrice Raveneau, Herve Rivano, and Razvan Stanica. 2017. PRIVA'MOV: Analysing Human Mobility Through Multi-Sensor Datasets. In *NetMob 2017*. HAL-Inria, Milan, Italy. https://hal.inria.fr/hal-01578557

[6] Alastair R Beresford and Frank Stajano. 2003. Location privacy in pervasive computing. *IEEE Pervasive computing* 2, 1 (2003), 46–55.

[7] Paul E Black. 2020. *DADS:The On-Line Dictionary of Algorithms and Data Structures.* Technical Report. National Institute of Standards and Technology. https://doi.org/10.6028/nist.ir.8318

[8] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. 2015. Geo-indistinguishability: A Principled Approach to Location Privacy. In *Distributed Computing and Internet Technology*, Raja Natarajan, Gautam Barua, and Manas Ranjan Patra (Eds.). Springer International Publishing, Cham, 49–72.

[9] Seongjin Choi, Jiwon Kim, and Hwasoo Yeo. 2021. TrajGAIL: Generating urban vehicle trajectories using generative adversarial imitation learning. *Transportation Research Part C: Emerging Technologies* 128 (2021), 103091.

[10] Chris Clifton and Tamir Tassa. 2013. On syntactic anonymity and differential privacy. In *2013 IEEE 29th International Conference on Data Engineering Workshops (ICDEW)*. IEEE, Brisbane, QLD, Australia, 88–93.

[11] Teddy Cunningham, Graham Cormode, and Hakan Ferhatosmanoglu. 2021. *Privacy-Preserving Synthetic Location Data in the Real World.* Association for Computing Machinery, New York, NY, USA, 23–33. https://doi.org/10.1145/3469830.3470893

[12] Ekler P de Mattos, Augusto CSA Domingues, and Antonio AF Loureiro. 2019. Give me two points and i'll tell you who you are. In *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, Paris, France, 1081–1087.

[13] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports* 3 (2013), 1376.

[14] Cynthia Dwork. 2008. Differential Privacy: A Survey of Results. In *Theory and Applications of Models of Computation*, Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–19.

[15] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3-4 (2014), 211–407.

[16] Ecenaz Erdemir, Pier Luigi Dragotti, and Deniz Gündüz. 2021. Privacy-Aware Time-Series Data Sharing With Deep Reinforcement Learning. *IEEE Transactions on Information Forensics and Security* 16 (2021), 389–401. https://doi.org/10.1109/TIFS.2020.3013200

[17] Fatima Zahra Errounda and Yan Liu. 2019. An analysis of differential privacy research in location data. In *2019 IEEE 5th Intl Conference on Big Data Security on Cloud (BigDataSecurity), IEEE Intl Conference on High Performance and Smart Computing,(HPSC) and IEEE Intl Conference on Intelligent Data and Security (IDS)*. IEEE, Washington, DC, USA, 53–60.

[18] Jie Feng, Can Rong, Funing Sun, Diansheng Guo, and Yong Li. 2020. PMF: A privacy-preserving human mobility prediction framework via federated learning. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 1 (2020), 1–21.

[19] Danielle L. Ferreira, Bruno A. A. Nunes, Carlos Alberto V. Campos, and Katia Obraczka. 2020. A Deep Learning Approach for Identifying User Communities Based on Geographical Preferences and Its Applications to Urban and Environmental Planning. *ACM Transactions on Spatial Algorithms and Systems* 6, 3 (May 2020), 1–24. https://doi.org/10.1145/3380970

[20] Bugra Gedik and Ling Liu. 2005. Location privacy in mobile systems: A personalized anonymization model. In *25th IEEE International Conference on Distributed Computing Systems (ICDCS'05)*. IEEE, Columbus, OH, USA, 620–629.

[21] Bugra Gedik and Ling Liu. 2007. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Transactions on Mobile Computing* 7, 1 (2007), 1–18.

[22] João Bártolo Gomes, Clifton Phua, and Shonali Krishnaswamy. 2013. Where Will You Go? Mobile Data Mining for Next Place Prediction. In *Data Warehousing and Knowledge Discovery*, Ladjel Bellatreche and Mukesh K. Mohania (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 146–158.

[23] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. 2008. Understanding individual human mobility patterns. *nature* 453, 7196 (2008), 779–782.

[24] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc., Montreal, Canada. https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf

[25] Mehmet Emre Gursoy, Ling Liu, Stacey Truex, Lei Yu, and Wenqi Wei. 2018. Utility-Aware Synthesis of Differentially Private and Attack-Resilient Location Traces. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (Toronto, Canada) *(CCS '18)*. Association for Computing Machinery, New York, NY, USA, 196–211. https://doi.org/10.1145/3243734.3243741

[26] M Emre Gursoy, Vivekanand Rajasekar, and Ling Liu. 2020. Utility-Optimized Synthesis of Differentially Private Location Traces. In *2020 Second IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*. IEEE, Atlanta, GA, USA, 30–39.

[27] Jihun Hamm. 2017. Minimax filter: Learning to preserve privacy from inference attacks. *The Journal of Machine Learning Research* 18, 1 (2017), 4704–4734.

[28] Xi He, Graham Cormode, Ashwin Machanavajjhala, Cecilia M Procopiuc, and Divesh Srivastava. 2015. DPT: differentially private trajectory synthesis using hierarchical reference systems. *Proceedings of the VLDB Endowment* 8, 11 (2015), 1154–1165.

[29] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.

[30] Chong Huang, Peter Kairouz, Xiao Chen, Lalitha Sankar, and Ram Rajagopal. 2017. Context-aware generative adversarial privacy. *Entropy* 19, 12 (2017), 656.

[31] 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 *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE, San Jose, CA, USA, 425–430.

[32] Haosheng Huang, Georg Gartner, Jukka M Krisp, Martin Raubal, and Nico Van de Weghe. 2018. Location based services: ongoing evolution and research agenda. *Journal of Location Based Services* 12, 2 (2018), 63–93.

[33] Hongbo Jiang, Jie Li, Ping Zhao, Fanzi Zeng, Zhu Xiao, and Arun Iyengar. 2021. Location privacy-preserving mechanisms in location-based services: A comprehensive survey. *ACM Computing Surveys (CSUR)* 54, 1 (2021), 1–36.

[34] Krzysztof W Kolodziej and Johan Hjelm. 2017. *Local positioning systems: LBS applications and services.* CRC press, Boca Raton.

[35] John Krumm. 2009. A survey of computational location privacy. *Personal and Ubiquitous Computing* 13, 6 (2009), 391–399.

[36] J. K. Laurila, Daniel Gatica-Perez, I. Aad, Blom J., Olivier Bornet, Trinh-Minh-Tri Do, O. Dousse, J. Eberle, and M. Miettinen. 2012. The Mobile Data Challenge: Big Data for Mobile Computing Research. In *Proceedings of the Conjunction with the 10th International Conference on Pervasive Computing*. EPFL, Newcastle, UK, 1–8. http://infoscience.epfl.ch/record/192489

[37] Sicong Liu, Junzhao Du, Anshumali Shrivastava, and Lin Zhong. 2019. Privacy adversarial network: representation learning for mobile data privacy. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 4 (2019), 1–18.

[38] Massimiliano Luca, Gianni Barlacchi, Bruno Lepri, and Luca Pappalardo. 2021. A Survey on Deep Learning for Human Mobility. arXiv:2012.02825 [cs.LG]

[39] Mohammad Malekzadeh, Richard G Clegg, Andrea Cavallaro, and Hamed Haddadi. 2020. Privacy and utility preserving sensor-data transformations. *Pervasive and Mobile Computing* 63 (2020), 101132.

[40] Soumia Menasria, Jianxin Wang, and Mingming Lu. 2018. The purpose driven privacy preservation for accelerometer-based activity recognition. *World Wide Web* 21, 6 (2018), 1773–1785.

[41] Guy M Morton. 1966. A computer oriented geodetic data base and a new technique in file sequencing. (1966).

[42] Sumit Mukherjee, Yixi Xu, Anusua Trivedi, Nabajyoti Patowary, and Juan L Ferres. 2021. privGAN: Protecting GANs from membership inference attacks at low cost to utility. *Proceedings on Privacy Enhancing Technologies* 2021, 3 (2021), 142–163.

[43] Milad Nasr, Reza Shokri, and Amir Houmansadr. 2018. Machine Learning with Membership Privacy Using Adversarial Regularization. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security* (Toronto, Canada) *(CCS '18)*. Association for Computing Machinery, New York, NY, USA, 634–646. https://doi.org/10.1145/3243734.3243855

[44] Nuria Oliver, Bruno Lepri, Harald Sterly, Renaud Lambiotte, Sébastien Deletaille, Marco De Nadai, Emmanuel Letouzé, Albert Ali Salah, Richard Benjamins, Ciro Cattuto, et al. 2020. Mobile phone data for informing public health actions across the COVID-19 pandemic life cycle.

[45] Nuria Oliver, Aleksandar Matic, and Enrique Frias-Martinez. 2015. Mobile network data for public health: opportunities and challenges. *Frontiers in public health* 3 (2015), 189.

[46] Kun Ouyang, Reza Shokri, David S. Rosenblum, and Wenzhuo Yang. 2018. A Non-Parametric Generative Model for Human Trajectories. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, Stockholm, Sweden, 3812–3817. https://doi.org/10.24963/ijcai.2018/530

[47] Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. 2018. The long road to computational location privacy: A survey. *IEEE Communications Surveys & Tutorials* 21, 3 (2018), 2772–2793.

[48] Jinmeng Rao, Song Gao, Yuhao Kang, and Qunying Huang. 2020. LSTM-TrajGAN: A Deep Learning Approach to Trajectory Privacy Protection. https://doi.org/10.48550/ARXIV.2006.10521

[49] Aria Rezaei, Chaowei Xiao, Jie Gao, Bo Li, and Sirajum Munir. 2018. Application-driven Privacy-preserving Data Publishing with Correlated Attributes. https://doi.org/10.48550/ARXIV.1812.10193

[50] Seungjae Shin, Hongseok Jeon, Chunglae Cho, Seunghyun Yoon, and Taeyeon Kim. 2020. User Mobility Synthesis based on Generative Adversarial Networks: A Survey. In *2020 22nd International Conference on Advanced Communication Technology (ICACT)*. IEEE, Phoenix Park, Korea (South), 94–103.

[51] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. 2011. Quantifying location privacy. In *2011 IEEE symposium on security and privacy*. IEEE, Oakland, CA, USA, 247–262.

[52] Reza Shokri, George Theodorakopoulos, Carmela Troncoso, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. 2012. Protecting Location Privacy: Optimal Strategy against Localization Attacks. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security* (Raleigh, North Carolina, USA)

(CCS '12). Association for Computing Machinery, New York, NY, USA, 617–627. https://doi.org/10.1145/2382196.2382261

[53] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. 2010. Limits of predictability in human mobility. *Science* 327, 5968 (2010), 1018–1021.

[54] Eran Toch, Boaz Lerner, Eyal Ben-Zion, and Irad Ben-Gal. 2019. Analyzing large-scale human mobility data: a survey of machine learning methods and applications. *Knowledge and Information Systems* 58, 3 (2019), 501–523.

[55] Jinbao Wang, Zhipeng Cai, and Jiguo Yu. 2019. Achieving personalized k-anonymity-based content privacy for autonomous vehicles in CPS. *IEEE Transactions on Industrial Informatics* 16, 6 (2019), 4242–4251.

[56] Senzhang Wang, Jiannong Cao, and Philip Yu. 2020. Deep Learning for Spatio-Temporal Data Mining: A Survey. *IEEE Transactions on Knowledge and Data Engineering* 14, 8 (2020), 1–1. https://doi.org/10.1109/TKDE.2020.3025580

[57] Yonghui Xiao and Li Xiong. 2015. Protecting Locations with Differential Privacy under Temporal Correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (Denver, Colorado, USA) *(CCS '15)*. Association for Computing Machinery, New York, NY, USA, 1298–1309. https://doi.org/10.1145/2810103.2813640

[58] Yonghui Xiao, Li Xiong, Si Zhang, and Yang Cao. 2017. LocLok: location cloaking with differential privacy via hidden markov model. *Proceedings of the VLDB Endowment* 10 (08 2017), 1901–1904. https://doi.org/10.14778/3137765.3137804

[59] Dingqi Yang, Daqing Zhang, Vincent W Zheng, and Zhiyong Yu. 2014. Modeling user activity preference by leveraging user spatial temporal characteristics in LBSNs. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45, 1 (2014), 129–142.

[60] Jianxi Yang, Manoranjan Dash, and Sin G. Teo. 2021. PPTPF: Privacy-Preserving Trajectory Publication Framework for CDR Mobile Trajectories. *ISPRS International Journal of Geo-Information* 10, 4 (April 2021), 224. https://doi.org/10.3390/ijgi10040224 Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.

[61] Yuting Zhan and Hamed Haddadi. 2019. Towards Automating Smart Homes: Contextual and Temporal Dynamics of Activity Prediction. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers* (London, United Kingdom) *(UbiComp/ISWC '19 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 413–417. https://doi.org/10.1145/3341162.3349307

[62] Yu Zheng, Hao Fu, Xing Xie, Wei-Ying Ma, and Quannan Li. 2011. *Geolife GPS trajectory dataset - User Guide* (geolife gps trajectories 1.1 ed.). Microsoft Research Asia. https://www.microsoft.com/en-us/research/publication/geolife-gps-trajectory-dataset-user-guide/ Geolife GPS trajectories 1.1.