

Privacy-Aware Human Mobility Prediction via Adversarial Networks

Yuting Zhan, Hamed Haddadi
Imperial College London
London, UK

Alex Kylo, Afra Mashhadi
University of Washington
Bothell, WA, USA

Abstract—As various 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. While these geolocated data could provide a rich understanding of human mobility patterns and address various societal research questions, privacy concerns for users’ sensitive information have limited their utilization. In this paper, we design and implement a novel 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. We quantify the utility-privacy trade-off of mobility datasets in terms of trajectory reconstruction risk, user re-identification risk, and mobility predictability. Our proposed architecture reports a Pareto Frontier analysis that enables the user to assess this trade-off as a function of Lagrangian loss weight parameters. The extensive comparison results on four representative mobility datasets demonstrate the superiority of our proposed architecture and the efficiency of the proposed privacy-preserving features extractor. Our results show that by exploring Pareto optimal setting, we can simultaneously increase both privacy (45%) and utility (32%).

Index Terms—mobility datasets, LSTM neural networks, mobility prediction, data privacy

I. INTRODUCTION

Geolocation and mobility data collected by location-based services (LBS), can reveal human mobility patterns and address various societal research questions [1]. For example, Call Data Records (CDR) have been successfully used to provide real-time traffic anomaly and event detection [2], and a variety of mobility datasets have been used in shaping policies for urban communities and epidemic management in the public health domain [3]. Human mobility prediction based on users’ trajectories, a popular and emerging topic, supports a series of important applications ranging from individual-level recommendation systems to large-scale smart transportation.

While there is no doubt of 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 [4]. A mobility privacy study conducted by De Montjoye et al [5] illustrates that four spatio-temporal points are enough to identify 95% of the individuals in a certain granularity. As human mobility traces are highly unique, a mechanism capable of decreasing the user re-identification risk can offer enhanced privacy protection in mobility data sharing.

In the past decade, the research community has extensively studied privacy of geolocated data via various location privacy protection mechanisms (LPPM) [6]. Some traditional privacy-preserving approaches such as k-anonymity and geo-masking have shown to be insufficient to prevent users from being re-identified [5], [7], [8]. More recently, some related works also try to apply machine-learning or deep-learning based approaches to explore the effective LPPM. Rao et al. proposed an LSTM-TrajGAN model to generate privacy-preserving synthetic mobility datasets for data sharing and publication [9]. 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 [10]. Though these state-of-the-art models provide a reasonable balance between utility and privacy, the effectiveness of the privacy mechanism and utility metrics have not been fully investigated in human mobility literature.

To this end, we posit an architecture for quantifying the utility-privacy trade-off of mobility datasets in terms of data reconstruction leakage (*i.e.*, *Privacy I*), user re-identification risk (*i.e.*, *Privacy II*), and mobility predictability (*i.e.*, *Utility*). In order to do so, we explore a novel mechanism to investigate these trade-offs and train a privacy-preserving feature extractor Enc_L based on representation learning and adversarial learning. Inspired by PAN [11] (privacy adversarial network), 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 LSTM auto-encoder (AE) 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 *maximize* the percentage of users who are re-identifiable through their trajectories by discriminative learning. We use Lagrange multipliers to vary the weights that are given to each of these objectives before combining them into a total loss, L_{sum} . The output of this model is a Pareto-Frontier analysis that would guide the user in investigating the trade-off between utility and privacy.

We report the analysis of our architecture (*i.e.*, LSTM-PAE) by a thorough evaluation on four real-world representative

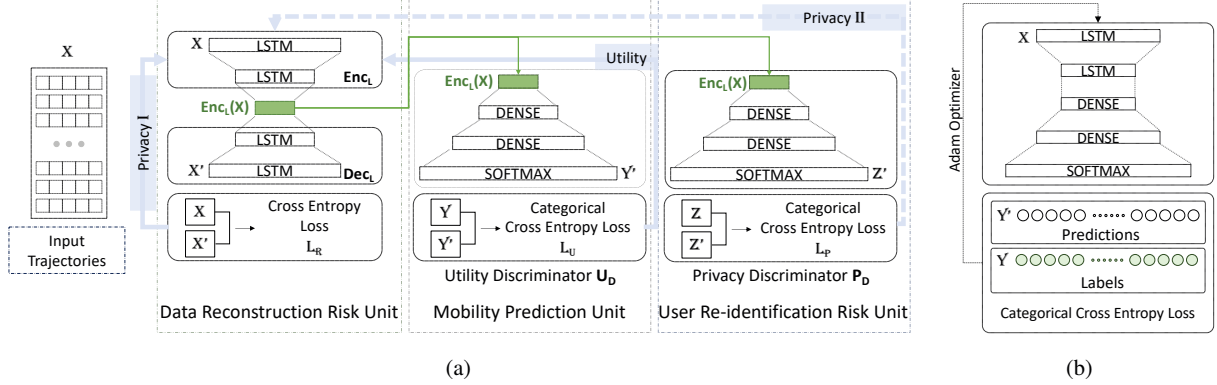


Fig. 1: (a) Schematic overview of the proposed privacy-preserving adversarial architecture with representation learning, consisting of data reconstruction risk unit, mobility prediction unit, and user re-identification risk unit. (b) The baseline LSTM network for standalone classifiers.

mobility datasets. A benchmark comparison is carried out with the state-of-the-art algorithm based on Generative Adversarial Network (GAN), namely LSTM-TrajGAN [9]. The results show that our architecture achieves better utility-privacy trade-offs than other models. That is, in the given spatial-temporal granularity, it is possible to achieve a better privacy level for a dataset with the same utility value and vice versa.

II. DESIGN OF THE ARCHITECTURE

A. Problem Definitions

Before describing our proposed LSTM-PAE model in detail, we first give a brief problem definition of the trade-off between mobility data utility and privacy in terms of mobility prediction accuracy, user re-identification efficiency and data reconstruction differences.

Data Utility: Mobility datasets are of great value for understanding human behavior patterns, smart transportation, urban planning, public health issue, pandemic management, and 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 [12]. Mobility prediction not only can be analyzed to understand personalized mobility patterns, but can also inform the allocation of public resources and community services. 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. Hence, the definition of *Utility* is concluded as followed:

Utility (U): the mobility predictability (*i.e.*, next location prediction accuracy)

Privacy Protection: With more and more intelligent devices and sensors are utilized to collect information about human activities, the trajectories also expose increasing intimate details about users' lives, from their social life to their preferences. A mobility privacy study conducted by De Montjoye et al. [5] illustrates that four spatio-temporal points are enough to identify 95% of the individuals in a certain granularity. The

capability of de-identification is important to balance the risks and benefits of mobility data usage, for all data owners, third parties, and researchers. We leverage the data reconstruction risk and user re-identification risk as our privacy metrics to evaluate our proposed privacy-aware architecture. Hence, the definition of *Privacy* is summarized as followed:

Privacy I (PI): the differences between the reconstructed data X' and the original data X , that is, information loss in the reconstruction process.

Privacy II (PII): the user re-identification inaccuracy, that is, the user de-identification effectiveness.

B. Architecture Overview

Our proposed privacy-preserving adversarial feature encoder, the LSTM-PAE, is based on representation learning and adversarial learning and aims to ease data sharing privacy concerns. As shown in Figure 1a, we design a multi-task adversarial network to learn an LSTM-based encoder Enc_L , which can generate the optimized feature representations $f = Enc_L(X)$ via lowering privacy disclosure risk of user identification information and improving the task accuracy (*i.e.*, mobility predictability) concurrently. Two potential malicious privacy leakages from the data reconstruction risk unit and the user re-identification risk unit, are attempted to retrieve sensitive information from the feature representations f .

Given mobility raw data \mathcal{X} for *Privacy I* (*e.g.*, data reconstruction risk unit), the ground-truth label z_i for *Privacy II* (*e.g.*, user re-identification risk unit), and the ground-truth label y_i for utility (*e.g.*, mobility prediction), we train the LSTM encoder Enc_L of this multi-task network by adversarial learning to learn the representation $\mathcal{F} = Enc_L(\mathcal{X})$. We design a specific loss function, namely *sum loss* \mathcal{L}_{sum} , for this optimization process. Specifically, when reconstructing the data \mathcal{X}' , an LSTM decoder Dec_L attempts to recreate the data based on the features \mathcal{F} , that is $Dec_L : \mathcal{F} \rightarrow \mathcal{X}'$. This data reconstruction unit is trained by maximizing the reconstruction loss \mathcal{L}_R while minimizing the \mathcal{L}_{sum} . The mobility prediction

Dataset-City	Bounding Box				Record Counts		Number			
	Latitude		Longitude		Train	Test	User ID	POI	Weekday	Hour
MDC-Lausanne	46.50	46.61	6.58	6.73	77393	19429	143	149	7	24
Privamov-Lyon	45.70	45.81	4.77	4.90	62077	16859	58	129	7	24
GeoLife-Beijing	39.74	40.07	116.23	116.56	95038	24578	145	960	7	24
FourSquare-NYC	40.55	40.99	-74.28	-73.68	43493	11017	466	1712	7	24

TABLE I: Overview of four mobility datasets.

unit, that is the utility discriminator U_D , is trained to output a probability distribution of the next location of interest, and this distribution has Y potential classes. Discriminative training here is to maximize the prediction accuracy by minimizing the utility loss \mathcal{L}_U , denoted as $\min \mathcal{L}_U$. The user re-identification risk unit, that is the privacy discriminator P_D , is trained to re-identify whom the target trajectory belongs. Then in this privacy discriminator, the user re-identification loss \mathcal{L}_P is maximized, denoted as $\max \mathcal{L}_P$.

In general, the encoder Enc_L should satisfy high predictability ($\min \mathcal{L}_U$) and low user re-identification accuracy ($\max \mathcal{L}_P$) of the mobility data when maximizing the reconstruction loss ($\max \mathcal{L}_R$) in reverse engineering, where the training objective can be written as:

$$\begin{aligned}
& \min (\mathcal{L}_{sum}) \\
& = -\lambda_1 (\max \mathcal{L}_R) + \lambda_2 (\min \mathcal{L}_U) - \lambda_3 (\max \mathcal{L}_P) \\
& = -\lambda_1 \|Dec_L(\mathcal{F}) - \mathcal{X}\|^2 + \lambda_2 \left(-\sum_{i=1}^Y y_i \log(U_D(\mathcal{F}))\right) \quad (1) \\
& \quad - \lambda_3 \left(-\sum_{i=1}^Z z_i \log(P_D(\mathcal{F}))\right)
\end{aligned}$$

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 Lagrange multipliers [13].

The overall training is to achieve privacy-utility trade-off by adversarial learning on \mathcal{L}_R , \mathcal{L}_U , and \mathcal{L}_P concurrently. The gradient of the loss (*i.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. The Lagrange multipliers are utilized to find the maxima or minima of a constrained problem. When two of them are set to zero, the model is transformed to a specific evaluation tool for a specific task. When three multipliers are utilized together, they control the relative importance of each unit and guide the overall model to find the maxima or minima given the specific trade-off choices.

III. EXPERIMENTAL SETTING

A. Datasets

Experiments are conducted on four representative mobility datasets: Mobile Data Challenge Dataset (MDC) [14], PRIVA'MOV [15], FourSquare NYC [16], and GeoLife [17]. Once imported into our architecture, each dataset was filtered and preprocessed individually to derive their respective test and training sets illustrated in Table I.

MDC: The MDC dataset, recorded from 2009 to 2011, 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 [14].

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 the purpose of this project, we only used the GPS traces from the dataset [15].

GeoLife: The GeoLife dataset was 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, mostly at a 5-second sampling rate [17].

FourSquare NYC: The Foursquare NYC dataset contains check-ins in NYC and Tokyo collected during the approximately 10 months from 12 April 2012 to 16 February 2013, containing 227,428 check-ins from 1,083 subjects in New York City [16].

B. Training

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.*, $SL = 10$). After the pre-processing of the datasets, 70% of the records of each user are segmented as the training dataset, 10% as the validation dataset and the remaining 20% as the testing dataset. We utilize the mini-batch learning method with 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 1) in terms of L_R , L_U , and L_P concurrently. Meanwhile, the *sum loss* L_{sum} is also optimized by using the Adam optimizer with learning rate of 0.0003. All the experiments were performed with the Tesla V100 GPU; a round of training would take 30 seconds on average and each experiment trains for 1000 rounds.

C. Metrics

We set *Euclidean* and *Manhattan* distance as our evaluation metrics for the data reconstruction unit to evaluate the quality of the reconstructed data X' generated from extracted features f . *Euclidean* distance gives the shortest or minimum distance between two points, while *Manhattan* distance applies only

if the points are arranged in the form of 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' . We leverage the top-n accuracy as our evaluation metric for both mobility prediction and user re-identification risk units.

IV. ARCHITECTURE EVALUATION

In this section, we present the comparison results of the proposed architecture LSTM-PAE and two baseline models under the same training setting.

Baseline Models

I. Standalone Model: It comprises three independent sub-models, namely data reconstruction sub-model, mobility prediction sub-model, and user re-identification sub-model. Each of the sub-models have a similar layer design as the corresponding unit in the LSTM-PAE, however, the results of the three sub-models are completely independent and have no effect on others. Differently from the LSTM-PAE, which leverages adversarial learning to finally attain an extracted feature representation f that satisfies the utility requirements and privacy budgets simultaneously, the standalone models only are trained for optimal inference accuracy at the individual tasks.

II. LSTM-TrajGAN [9]: 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 common geomasking 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 is in a line with the mobility prediction based utility assessment in our work. Hence, we train a 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 LSTM-PAE is training an encoder Enc_L that forces the extracted representations f to convey maximal utility while minimizing private information about user identities, via adversarial learning.

A. Performance Comparison

We first compare our proposed models with the standalone model and the LSTM-TrajGAN model on four representative mobility datasets, as details shown in Table II. The overall performance is evaluated in terms of the *utility level* provided by the mobility prediction unit and the *privacy threat* provided by two risk units. The *Model I* is our proposed architecture but without applying the Lagrange multipliers (*i.e.*, where each losses are weighed equally), and the *Model II* is the one with Lagrange multipliers (*i.e.*, $\lambda_1 = 0.1$, $\lambda_2 = 0.8$, $\lambda_3 = 0.1$ for the results in Table II). The results in Table II are based on the input data with trajectory sequence length 10 (that is

$SL = 10$). Because the standalone models are trained without the consideration for the utility-privacy trade-offs, the results on the standalone models can be leveraged to explain the best inference accuracy (*i.e.*, mobility prediction accuracy and user re-identification accuracy) that each composition unit could achieve. Differently from the standalone model, the TrajGAN and LSTM-PAE are both taking the utility-privacy trade-offs into consideration and we compare their trade-offs with the standalone version. Hence, in Table II, these results are shown in utility decline and privacy gain, both of which are in a percentage format. The similarity indexes are leveraged to intuitively represent the difference between the original data X and reconstructed data X' , where the larger value indicates numerical differences between them.

Table II demonstrates that our proposed models, especially the one with Lagrangian multipliers, outperforms the LSTM-TrajGAN model across various datasets. For instance, when models are trained with the MDC dataset, our *Model II* achieves the best privacy-utility trade-offs among different models, as the utility decline is only 13.43% but with 65.51% privacy gain, while 46.32% utility decline and 20.32% privacy gain with the TrajGAN. The similarity indexes also indicate the reconstructed data X' via *Model II* has bigger differences than the one via the TrajGAN. Second, although the utility decline of the TrajGAN on the Priva'Mov dataset is 4.21% higher than our *Model II*, both two privacy metrics of the TrajGAN are worse than the *Model II*. Our model has better overall trade-offs in utility requirements and privacy budgets. The performance on Geolife and FourSquare are similar but inverse, where the utility of our model is better than TrajGAN and with slightly weaker privacy preservation. We leverage the composite metrics to score the overall performance of different models, and demonstrates that our model achieves better utility-privacy trade-offs. The comparisons between *Model I* and *Model II* also illustrate the importance of the Lagrange multipliers in not only providing flexibility to our proposed architecture that enable its application in different scenarios, but also enhancing the utility-privacy trade-offs in this special case.

B. Utility-Privacy Analysis

In this section, we first present the utility-privacy trade-off analysis between TrajGAN and our proposed LSTM-PAE in terms of the mobility prediction accuracy (*i.e.*, U) and user de-identification efficiency (*i.e.*, PII). We then discuss the privacy risk (*i.e.*, PII) of our proposed framework among four representative mobility datasets.

1) Trade-off Comparison: Figure 2 presents the trade-off comparisons of the four datasets in terms of the U and PII under different Lagrangian settings, where the *hollow squares* and *hollow diamonds* show the tradeoffs provided by the proposed LSTM-PAE in $SL = 5$ and $SL = 10$, respectively. The *solid points* present the utility-privacy trade-off provided by the TrajGAN under the same spatial granularity and same trajectory sequence length. As can be seen from these results, in all four cases the synthetic dataset generated

Datasets	Models	Privacy I		Utility (% for decline)			Privacy II (% for gain)		
		Euc(log)	Man(log)	top-1	top-3	top-5	top-1	top-3	top-5
MDC	Standalone	0.001	0.002	0.9347	0.9837	0.9922	0.9247	0.9819	0.9911
	TrajGAN	3.526	5.456	-46.32%	-24.16%	-15.98%	+20.32%	+8.13%	+4.02%
	Our Model	I 2.294	4.542	-54.56%	-34.74%	-25.10%	+69.80%	+50.44%	+39.95%
		II 3.732	6.023	-13.43%	-6.26%	-3.95%	+65.51%	+45.11%	+34.86%
Priva'Mov	Standalone	1.281	2.554	0.9482	0.9878	0.9954	0.5643	0.8215	0.8765
	TrajGAN	3.704	5.779	-6.60%	-1.89%	-0.93%	+14.17%	+14.35%	+8.88%
	Our Model	I 1.740	3.433	-3.36%	-1.59%	-0.81%	+27.02%	+14.19%	+9.19%
		II 4.164	5.803	-10.81%	-6.83%	-4.91%	+35.29%	+14.97%	+10.05%
Geolife	Standalone	1.903	3.804	0.4705	0.6842	0.7636	0.6572	0.8690	0.9294
	TrajGAN	4.581	6.680	-62.31%	-50.45%	-43.72%	+66.73%	+47.89%	+37.22%
	Our Model	I 1.776	3.469	-31.45%	-25.02%	-21.90%	+54.88%	+39.59%	+30.81%
		II 4.616	6.928	-21.13%	-18.78%	-17.11%	+55.49%	+40.40%	+32.34%
FourSquare	Standalone	2.357	4.464	0.6468	0.8210	0.8823	0.8780	0.9735	0.9892
	TrajGAN	4.795	6.710	-26.30%	-22.30%	-18.75%	+51.86%	+32.49%	+23.49%
	Our Model	I 2.418	4.533	-51.05%	-41.45%	-35.20%	+53.47%	+35.26%	+25.86%
		II 4.541	6.638	-2.54%	-3.14%	-2.84%	+51.08%	+34.39%	+26.16%

TABLE II: Performance comparison between our proposed models with standalone model and the TrajGAN model. The *Model I* is our proposed architecture without Lagrange multipliers, and the *Model II* is the one with multipliers ($\lambda_1 = 0.1$, $\lambda_2 = 0.8$, and $\lambda_3 = 0.1$). The results shown in this table are all with trajectory 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 declines; and the *Privacy II (%)* represents the privacy gains calculated via the user re-identification *inaccuracy* rate.

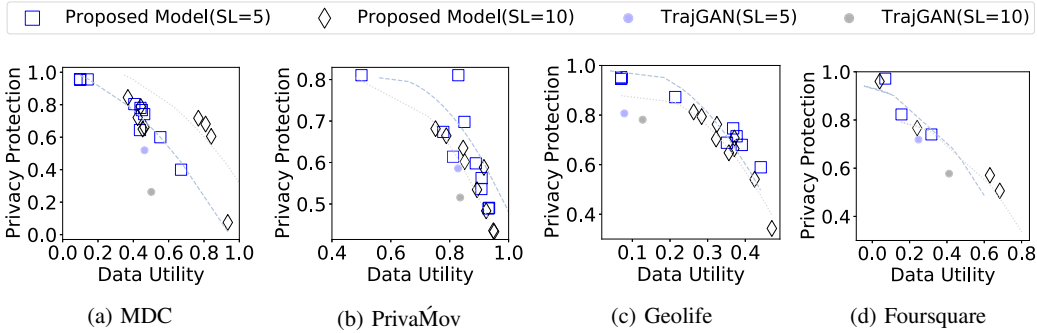


Fig. 2: Pareto Frontier trade-off of Utility and Privacy on four datasets. The hollow squares and diamonds present results of the proposed models. The solid points present results of the TrajGAN. Blue color means $SL = 5$. Black color means $SL = 10$.

by the TrajGAN is not *pareto-optimal*. That is, in that given spatial-temporal granularity, the proposed architecture is able to achieve a better privacy level for a dataset with the same utility value. Compared with the TrajGAN, our proposed architecture improves utility and privacy at the same time on four datasets. Especially for the performance on the MDC dataset, the privacy improves 45.21% than the TrajGAN while the utility also increases 32.89%. These results illustrate that our proposed model achieves promising performance in training a privacy-sensitive encoder Enc_L for different datasets.

2) *Privacy Leakage Risk Analysis*: After evaluating the superior performance of our proposed framework, we discuss the privacy leakage risks among four representative mobility datasets in terms of user re-identification inaccuracy (PII , *privacy gain* in Figure 3). We use five different combinations of Lagrangian multipliers to evaluate the comprehensive performance of the proposed model, namely setting *I*, *II*, *III*, *IV*, and *V* in Figure 3.

Figure 3 presents *Utility* and *Privacy II* (i.e., user re-identification risk) trade-offs of the proposed system on the four datasets. The *Zero* line (i.e., $y = 0\%$) in each sub-figure is leveraged to indicate the original utility rate (U) and privacy rate (PII) of the raw data. The blue line with square marker is the privacy gain rate with top-1 accuracy and the blue line without marker is the top-5 accuracy. The orange lines with and without triangle marker present the utility decline rate with top-1 and top-5 accuracies, respectively. Hence, the orange area represents the utility decline while the light-green area represents the privacy gain when compared with original results. The dark-green area represents the trade-offs between utility and privacy budgets. The x-axis shows five different settings of the model, and the y-axis shows the trade-offs (i.e., *trade-offs* = *privacy gain* + *utility decline*).

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

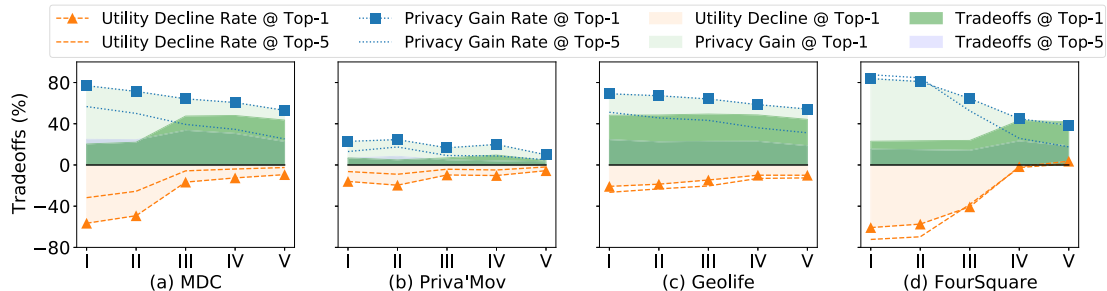


Fig. 3: User re-identification privacy gain (PII) versus utility decline (U) on four datasets. The orange area represents the utility decline while the light-green area represents privacy gain. The dark-green area represents the trade-offs between utility achievement and privacy budgets. The x-axis shows five different settings of the model, and the y-axis shows the trade-offs.

decline but more than 50% privacy gains. The performance on MDC and FourSquare also show the promising utility-privacy trade-offs, especially for setting V on the FourSquare dataset, 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 [5]. Our results emphasize that the concern of user re-identification risk could be alleviated effectively with our proposed model.

V. CONCLUSION

In this paper, we presented a privacy-preserving architecture based on the adversarial networks. Our model takes into account 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 hyper-parameters using four real-world mobility datasets. The Lagrange multipliers λ_1 , λ_2 , and λ_3 bring more flexibility to our framework that enable 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. 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 model 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, for efficiency in searching for the optimal Lagrange multipliers.

ACKNOWLEDGMENT

The authors are grateful to Imperial College London and the China Scholarship Council (CSC) for financial support. Hamed Haddadi was partially supported by the UK-Dementia Research Institute.

REFERENCES

[1] K. W. Kolodziej and J. Hjelm, *Local positioning systems: LBS applications and services*. CRC press, 2017.

[2] S. Wang, J. Cao, and P. Yu, "Deep learning for spatio-temporal data mining: A survey," *IEEE Transactions on Knowledge and Data Engineering*, 2020.

[3] N. Oliver, B. Lepri, H. Sterly, R. Lambiotte, S. Deletaille, M. De Nadai, E. Letouzé, A. A. Salah, R. Benjamins, C. Cattuto *et al.*, "Mobile phone data for informing public health actions across the covid-19 pandemic life cycle," 2020.

[4] J. Krumm, "A survey of computational location privacy," *Personal and Ubiquitous Computing*, vol. 13, no. 6, pp. 391–399, 2009.

[5] Y.-A. De Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," *Scientific reports*, vol. 3, p. 1376, 2013.

[6] B. Gedik and L. Liu, "Protecting location privacy with personalized k-anonymity: Architecture and algorithms," *IEEE Transactions on Mobile Computing*, vol. 7, no. 1, pp. 1–18, 2007.

[7] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human mobility," *Science*, vol. 327, no. 5968, pp. 1018–1021, 2010.

[8] Y. Zhan, A. Killo, A. Mashhadi, and H. Haddadi, "Privacy-aware human mobility prediction via adversarial networks," *arXiv preprint arXiv:2201.07519*, 2022.

[9] J. Rao, S. Gao, Y. Kang, and Q. Huang, "Lstm-trajgan: A deep learning approach to trajectory privacy protection," *arXiv preprint arXiv:2006.10521*, 2020.

[10] J. Feng, C. Rong, F. Sun, D. Guo, and Y. Li, "Pmf: A privacy-preserving human mobility prediction framework via federated learning," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 4, no. 1, pp. 1–21, 2020.

[11] S. Liu, J. Du, A. Shrivastava, and L. Zhong, "Privacy adversarial network: representation learning for mobile data privacy," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 3, no. 4, pp. 1–18, 2019.

[12] Y. Zhan and H. Haddadi, "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*, 2019, pp. 413–417.

[13] B. Beavis and I. Dobbs, *Optimisation and stability theory for economic analysis*. Cambridge university press, 1990.

[14] J. K. Laurila, D. Gatica-Perez, I. Aad, O. Bornet, T.-M.-T. Do, O. Dousse, J. Eberle, M. Miettinen *et al.*, "The mobile data challenge: Big data for mobile computing research," Tech. Rep., 2012.

[15] S. B. Mokhtar, A. Boutet, L. Bouzouina, P. Bonnel, O. Brette, L. Brunie, M. Cunche, S. D'Alu, V. Primault, P. Raveneau *et al.*, "PRIVA'MOV: Analysing human mobility through multi-sensor datasets," in *NetMob 2017*, 2017.

[16] D. Yang, D. Zhang, V. W. Zheng, and Z. Yu, "Modeling user activity preference by leveraging user spatial temporal characteristics in LBSNs," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 45, no. 1, pp. 129–142, 2014.

[17] Y. Zheng, H. Fu, X. Xie, W.-Y. Ma, and Q. Li, *Geolife GPS trajectory dataset - User Guide*, geolife gps trajectories 1.1 ed., July 2011, geolife GPS trajectories 1.1. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/geolife-gps-trajectory-dataset-user-guide/>