



HAL
open science

Enforcing Adaptive Location Privacy with Federated Learning

Yanis Meziani, Besma Khalfoun, Sara Bouchenak, Sonia Ben Mokhtar, Vlad Nitu

► To cite this version:

Yanis Meziani, Besma Khalfoun, Sara Bouchenak, Sonia Ben Mokhtar, Vlad Nitu. Enforcing Adaptive Location Privacy with Federated Learning. COMPAS 2020 : parallémisme, Architecture, Système/ Temps réel, Jun 2020, Lyon, France. <hal-03339363>

HAL Id: hal-03339363

<https://hal.science/hal-03339363v1>

Submitted on 9 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Enforcing Adaptive Location Privacy with Federated Learning

Yanis Meziani^{1,2}, Besma Khalfoun¹, Sara Bouchenak¹, Sonia Ben Mokhtar¹, and Vlad Nitu¹

¹INSA Lyon – LIRIS, France {firstname.lastname}@insa-lyon.fr

²Ecole Nationale Supérieure d'Informatique Algiers, Algeria

Abstract

Preserving the privacy of mobility data has been the center of active research in the last decade as this data may reveal sensitive information about individuals (e.g., home, work places, political, religious, sexual preferences). In this context, a large variety of Location Privacy Protection Mechanisms (LPPMs) have been proposed. To employ LPPMs more effectively for the benefits of the users' privacy, adaptive solutions that dynamically combine LPPMs have also been investigated. These solutions apply various LPPMs on a given trace and choose the one that better meets privacy and utility requirements. To meet this objective, adaptive solutions often rely on a trusted proxy that gathers users' traces and apply LPPMs locally. In this paper we release this assumption by designing an approach, i.e., EDEN, where mobility data never leaves the user's device before being protected by an appropriate LPPM. Experimental evaluation performed on real mobility traces shows that EDEN with its local and adaptive strategy outperforms individual LPPMs both in terms of privacy and utility metrics.

Keywords : Mobility Data, Location Privacy, Re-identification Attacks, Federated learning.

1. Introduction

The last decade has witnessed the apparition of a wide variety of location based services (LBSs), which are extensively used by increasing numbers of users. While SBSs offer useful services to the customers (e.g., finding points of interests, geo-gaming), the latter massively collect (and possibly share/sell) user location data. For instance, recent studies¹ have shown that among 30000 sampled iOS applications available in the App store, 51.1% collect location data when the application is in use. The downside is that collecting users' location data opens many privacy threats if the latter falls between the hands of curious/malicious adversaries. For instance, mobility data may very well reveal a user's home and workplace, health status or even religious or sexual preferences if the latter regularly visits health centers, worship places or libertine places (respectively).

In order to mitigate the above threats, the research community has been actively proposing Location Privacy Protection Mechanisms (LPPMs). Examples of proposed LPPMs include Geo-indistinguishability [5], which adds spacial noise to a user's GPS coordinates, Promesse [16], which removes places where the user stops and Trilateration [10], which generates dummy locations to obfuscate the user's real location. In this context, one of the issues that has been

1. 2019 study on iOS App Permissions : <https://www.wandera.com/ios-app-permissions/>

exhibited by the research community is that there is no *one size fits all* LPPM that effectively protects users' location data in all contexts. An additional dimension that makes it more complex to protect users' mobility data is that each LPPM can be configured differently hence offering a different privacy vs utility trade-off. To deal with this issue, solutions that try to find the best LPPM (and its corresponding proper configuration) to protect a given user mobility data at a given point in time have been devised (e.g., Mood [11]). These solutions usually rely on known attacks (e.g., [8, 12]) and utility metrics (e.g., the spacial distortion induced by the LPPM compared to raw traces) to assess the effectiveness of LPPMs. Specifically, given a mobility trace on which a set of LPPMs are applied, adaptive solutions would choose the one that better resists state-of-the-art attacks while preserving data utility above a given threshold. However, these solutions often rely on the assumption that a trusted proxy server collects users' mobility data, runs the attacks and consequently apply the appropriate LPPM. In this paper we aim at releasing this assumption by designing an approach where mobility data never leaves the user's device before being protected by an appropriate LPPM. Specifically, we present EDEN a solution that relies on federated learning to train a decentralized attack model called FURIA. This attack is then used locally along with a utility metric in order to choose the most appropriate LPPM for a given mobility trace. We evaluated EDEN with a set of experiments using real mobility data. We used EDEN to choose the best LPPM and its corresponding configuration among nine possibilities (3 LPPMs and 3 different configurations). The results show that EDEN outperforms individual LPPMs both in terms of privacy (better resilience against FURIA our decentralized attack) and in terms of data utility. The remaining of this paper is structured as follows. First, we present in Section 2, Background and related work. Then, we describe EDEN design principles in section 3. Further, in section 4, we proceed to an experimental evaluation of our attack. Finally, we conclude in section 5.

2. Background and Related Work

There are two parties involved in the context of our work : the Location-Based Service (LBS) represented by an LBS server and the users represented by their mobile devices. We suppose that users send queries to the LBS server with raw mobility data (single or multiple mobility records). Therefore, the LBS provider has access to all raw mobility data of all users.

We consider the LBS as an *honest but curious* entity. It provides the users with the requested information according to their locations but it may store the collected raw mobility data and exploit it maliciously. The stored mobility traces are noted as the background knowledge of the LBS, $BK = \{t_1, t_2, \dots, t_m\}$. Where t_i corresponds to the mobility trace of user u_i .

After that, users start using Location Privacy Protection Mechanisms (LPPMs) in order to protect their mobility data. As of now, only obfuscated mobility data are sent to the LBS. The latter tries to re-identify the owner of the data based on BK.

In the literature, we have many re-identification attacks suggestion. Among them, we find the POI-attack [15], which uses the Point Of Interests (POIs) to represent mobility of a user, POIs are locations where users spent an amount of time. PIT-attack [9] also uses Mobility Markov chains where states are POIs and edges represent the probability of transition between POIs. AP-attack[12] describes user mobility as a heatmap which inspires our work.

3. EDEN Design Principles

3.1. System Model

Let $U = \{U_1, U_2, \dots, U_n\}$ be the set of users. Each user U_i holds a mobile device which collects and stores her raw mobility data. These data are used by geo-located services in order

to respond to user's requests. A user request contains the user ID (e.g., IP address), the user location (i.e., GPS latitude and longitude), the time at which the request is issued, and the actual requested method (e.g., best restaurants nearby). We consider the record $r_j = (\text{lat}_j, \text{lng}_j, t_j)$ as a mobility data record where lat_j and lng_j respectively correspond to the latitude and the longitude of GPS coordinates, and t_j is a timestamp. A sequence of mobility data records forms a mobility trace $T_j = \{r_1 = (\text{lat}_1, \text{lng}_1, t_1), r_2 = (\text{lat}_2, \text{lng}_2, t_2), \dots\}$.

3.2. Overview of EDEN

EDEN (*Enforcing aDaptive location privacy with federated lEarNing*) is a user side mobility data protection system. Thus, user sensitive data are kept locally and are not shared with other entities, which increases privacy guarantees. EDEN's main objective is to protect users mobility data in online/semi-online services. EDEN includes and uses a distributed user re-identification attack called FURIA (Federated User Re-Identification Attack). This allows EDEN to better assess the resilience of protected data to privacy attacks.

The overall execution process of EDEN and FURIA is as follows. During the day, all or a subset of users run mobile applications and get access to the LBS services at any time of the day. The mobility data received by the LBS provider is continuously protected by EDEN thanks to FURIA which improves its users mobility knowledge every day. For a given day j , let $T_i = \{T_{i_1}, T_{i_2}, \dots, T_{i_k}\}$ be the set of mobility traces of user U_i collected over the whole day. At the end of the day (e.g., at midnight), each user gets back the global model of FURIA, transforms her set of mobility traces T_i into a vector of features, trains the model locally and sends it back to FURIA maintainer. Once The FURIA maintainer gets the users updates, it aggregates the gradients and produces a updated global model which continuously learns new discriminative mobility patterns. The latter helps to uniquely identify users. User mobility data features are extracted from raw mobility traces in order to represent spatio-temporal behavior of users in a compact way. The following mobility features are considered :

Spatial feature. The spatial information (i.e. lat and lng) is projected on a heat map. A heat map is a set of cells of equal size. For each cell, the proportion of mobility records in a given trace T_j that belong to that cell is computed. This corresponds to the cell visit rate.

Temporal feature. The temporal information is considered to differentiate similar mobility patterns between day-night shifts, e.g., a user living near the working place of another user. If temporal information is not taken into account, they may have similar heat maps but at different times of the day. A simple, yet effective, temporal information that is the average time of the day of all the records in a given T_j . This is convenient in the case of online/semi-online LBS, where mobility traces are minutes length (e.g., online session-based services) or hours length (e.g., crowd-sensing systems), without exceeding a day.

Additional features. Other types of information are extracted to enrich the user mobility profile. For instance, the number of mobility data records available in the trace T_j is considered. This allows to represent location-based service usage intensity. We also extract the centroid of the mobility trace T_j (i.e., centroid's latitude and centroid's longitude), to capture the average position of the user in the map.

3.3. EDEN Architecture

The architecture of EDEN is depicted in the left part of Figure 1. EDEN takes as input a user mobility trace T and a set of LPPMs with various configuration (i.e. low, medium and high impact on mobility data) and returns as output an obfuscated mobility trace T'_i which will be publicly released for crowd-sensing campaign or sent to LBS server for a particular purpose (e.g., search for nearby market). It has four main components, the first component *Apply LPPMs* aims at applying the considered LPPMs on the raw mobility data T stored on the user

mobile device (step ①). The second component *Format Data* transforms the different obfuscated versions of the raw mobility trace, i.e., $\{T'_1, T'_2, \dots, T'_n\}$ into feature vectors as enumerated in section 3.1, (step ②). the third component *Global Model FURIA* is the crucial component in the system. Once a day, the user retrieves the latest version of the Global Model FURIA. The latter evaluates the re-identification of each vector of features of the obfuscated data, (step ③). If the attack succeeds in predicting the right identity label associated to the mobility trace, the latter will be deleted and not shared outside the user device. Otherwise, the protected mobility trace is potentially elected to be sent to the LBS server. Finally, the last component *Best coverage*, only the protected mobility trace against the most up to date joint model of FURIA with the maximal area coverage is retained. Note that the area coverage is computed between the original and the obfuscated mobility trace, T and T'_i respectively, to measure the covered cells by T'_i in total cells of T [13], (step ④). Once the protected mobility trace with best utility is produced and sent to the LBS server, the latter treats the user's request and returns back the response to the user's device, (step ⑤).

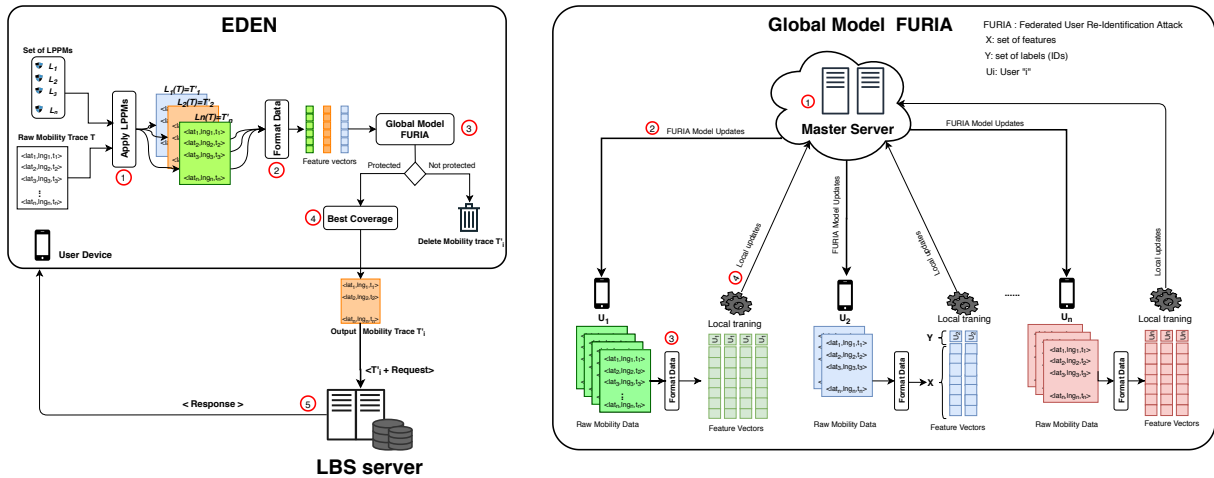


FIGURE 1 – EDEN and FURIA architecture

3.4. FURIA Architecture

FURIA's global model is an essential part of EDEN. It applies federated machine learning over decentralized user mobility data, in order to learn user mobility models from different users, in a privacy preserving way. It involves the following parties, c.f., Figure 1 - right side : (i) mobile user devices where raw mobility data are kept at user's own device, (ii) a master server responsible of the Federated modeling and processing of the user re-identification attack, i.e., an attack which aims to associate an anonymous mobility trace to its originating user identity. First, The *Master server* initializes randomly a straightforward classification algorithm which is *Logistic regression* and sends it to all participants (step ① and ②). This model is denoted as $FURIA_0$. Each participant U_j transforms its raw mobility data of the day to feature vectors (step ③) and trains the model $FURIA_0$ locally on the generated vectors. Notice that the training phase is done in the user device and not in the *Master server*. Once all participants have finished the first learning round, they send their local updates (i.e. gradients) of their current local models to the *Master server*, (step ④). The *Master server* aggregates users' gradients and produces a new model, denoted $FURIA_1$, ready for use at the following day. This process is iteratively done once at the end of the day (e.g., midnight) for continuous learning of users' mobility behavior.

3.5. Threat Model

In this work, the following assumptions are made. First, as EDEN uses FURIA's global model to assess its effectiveness, the latter requires a trusted proxy server, the so called *Master Server*. Its main task is to aggregate the local mobility updates of each participant and produce a joint updated model. The latter can be implemented in Trusted Execution Environment. Mobile user devices which participate collaboratively in building FURIA's global model are not trusted. This means that a malicious participant can use poisoning attacks to introduce backdoor functionality into the global model and infer sensitive information. Many current researchers are working on this aspect which is orthogonal to our research objectives [7]. In order to update FURIA's global model, users need to send their transformed mobility data at night to the Master Server. User devices are supposed to be charging, idle and connected to WIFI. Finally, the LBS is honest but curious. It responds to users' requests on one hand but on the other hand it may exploit the mobility data maliciously to infer personal information about the users.

4. Experimental Evaluation

4.1. Experimental Setup

EDEN is developed in Python by using Pytorch library [2] for developing the model and Pysyft[1] for emulating the federated protocols and participants. For data preprocessing, we use Scikit-learn [4] to scale data, and S2Geometry[3] to form the features vector, by gathering mobility data on cells with edge length ranges from 850m to 1185m .

In our experiments, the training is done on multiple round, each round represents a day. We assume that users train the model each midnight. On this way, on day x we train the model on data collected the day before (i.e. day $x-1$) and we test the model on data collected that day. We use the simple gradient descent algorithm as an optimizer for the logistic regression, with a value of 0.001 for the learning rate parameter.

State-of-the-art LPPMs.In our scenario, we consider several existing LPPMs (location privacy protection mechanisms), and compare them to our federated learning-based protection mechanism. In order to apply LPPMs in our dataset, we used an open source library [14], with configurations described in Table 1.

Mobility Dataset. In our experiments, we use the Privamo dataset which contains real mobility data of 48 users in Lyon [6]. We extracted the most active month (i.e. 30 days). For each user, we formed traces by gathering all mobility data in chunks of 30 minutes, which corresponds to a LBS that collects users' mobility data every 30 minutes.

LPPM	LPPM's configuration
Geo-Indistinguishability (Geo-I) [5]	0.01, 0.005, 0.001
Promesse (PROM) [16]	50 m, 100 m , 200 m
Trilateration (TRL) [10]	1 km, 2 km, 3 km

TABLE 1 – State-of-the-art LPPMs and their configurations

4.2. Comparison of EDEN with State-of-the-Art LPPMs

In the following, we run our algorithm attack on protected data and compute the attack success rate each day. Then, We compare EDEN with Geo-I, TRL, and PROM. Due to space constraints, we represent in Figure 2 only results regarding to the best configuration in term of privacy for each LPPM. We plot the curve associated to the attack over none obfuscated data (i.e. NOBF) as a reference for the protection efficiency. For each case, we chose one color : orange, blue, red,

green and pink for NOBF, Geo-I, TRL, PROM and EDEN respectively. The attack success rate in our protection does not exceed 4%, as a comparison the accuracy for Geo-I range from 31% to 75%, for TRL is from 10% to 63% and for PROM it does not exceed 10%. EDEN protect 100% of the traces during 7 days and at 99% during 17 days. The accuracy decreases in all cases in some days because of the appearance of new users.

Figure 3 presents the utility of data protected with state-of-the-Art LPPMs and compares it to EDEN. Here, utility metric is the area coverage of the protected data compared to the raw (i.e., non-protected) data. We consider three levels of utility : high utility if the area coverage is higher than 60%, utility if it is between 30% to 60%, and low in other cases. We observe that apart from the case of EDEN and PROM (configured with a 50 m distance), more than the half of protected data have a low utility. Furthermore, EDEN has the highest ratio of protected data with high utility.

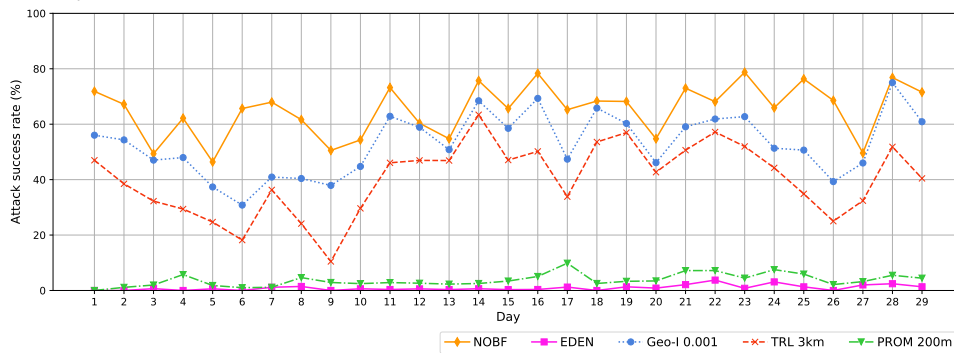


FIGURE 2 – Attack success rate with EDEN vs. state-of-the-art LPPMs

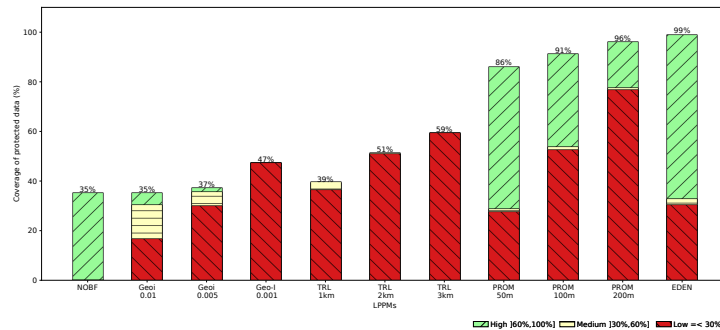


FIGURE 3 – Data utility with EDEN vs. state-of-the-art LPPMs

5. Conclusion

In this paper we presented EDEN a solution that protects mobility data by choosing the best among a set of off-the-shelf LPPMs, without relying on a trusted proxy. Specifically, EDEN relies on federated learning to train a decentralized attack model called FURIA. This attack is then used locally along with a utility metric in order to choose the most appropriate LPPM for a given mobility trace. We evaluated EDEN by performing a set of experiments on real mobility traces over a period of one month. We used EDEN to choose the best LPPM and its corresponding configuration among nine possibilities. The results shown that EDEN outperforms individual LPPMs both in terms of privacy (better resilience against our FURIA decentralized attack) and in terms of data utility.

References

1. PySyft is a python library for secure and private deep learning. – <https://www.github.com/OpenMined/PySyft>. Accessed : 2020-04-01.
2. PyTorch an open source machine learning library based on the torch library. – <https://www.pytorch.org>. Accessed : 2020-04-01.
3. S2 is a library for spherical geometry. – <http://www.s2geometry.io/>. Accessed : 2020-04-01.
4. Scikit-learn is a free software machine learning library for the python programming language. – <https://www.scikit-learn.org/stable/>. Accessed : 2020-04-01.
5. Andrés (M. E.), Bordenabe (N. E.), Chatzikokolakis (K.) et Palamidessi (C.). – Geoindistinguishability : Differential privacy for location-based systems. – In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pp. 901–914, 2013.
6. Boutet (A.), Mokhtar (S. B.) et Primault (V.). – Uniqueness assessment of human mobility on multi-sensor datasets. 2016.
7. Fung (C.), Yoon (C. J. M.) et Beschastnikh (I.). – Mitigating sybils in federated learning poisoning. *CoRR*, vol. abs/1808.04866, 2018.
8. Gambs (S.), Killijian (M.-O.) et del Prado Cortez (M. N.). – Show me how you move and i will tell you who you are. – In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Security and Privacy in GIS and LBS*, pp. 34–41, 2010.
9. Gambs (S.), Killijian (M.-O.) et del Prado Cortez (M. N.). – De-anonymization attack on geolocated data. *Journal of Computer and System Sciences*, vol. 80, n8, 2014, pp. 1597–1614.
10. Huang (Y.), Cai (Z.) et Bourgeois (A. G.). – Search locations safely and accurately : A location privacy protection algorithm with accurate service. *Journal of Network and Computer Applications*, vol. 103, 2018, pp. 146–156.
11. Khalfoun (B.), Maouche (M.), Mokhtar (S. B.) et Bouchenak (S.). – Mood : Mobility data privacy as orphan disease : Experimentation and deployment paper. – In *Proceedings of the 20th International Middleware Conference, Middleware 2019, Davis, CA, USA, December 9-13, 2019*, pp. 136–148. ACM, 2019.
12. Maouche (M.), Mokhtar (S. B.) et Bouchenak (S.). – Ap-attack : a novel user re-identification attack on mobility datasets. – In *Proceedings of the 14th EAI International Conference on Mobile and Ubiquitous Systems : Computing, Networking and Services*, pp. 48–57, 2017.
13. Primault (V.), Boutet (A.), Mokhtar (S. B.) et Brunie (L.). – Adaptive location privacy with ALP. – In *35th IEEE Symposium on Reliable Distributed Systems, SRDS 2016, Budapest, Hungary, September 26-29, 2016*, pp. 269–278. IEEE Computer Society, 2016.
14. Primault (V.), Maouche (M.), Boutet (A.), Mokhtar (S. B.), Bouchenak (S.) et Brunie (L.). – Accio : How to make location privacy experimentation open and easy. – In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pp. 896–906. IEEE, 2018.
15. Primault (V.), Mokhtar (S. B.), Lauradoux (C.) et Brunie (L.). – Differentially private location privacy in practice. *arXiv preprint arXiv :1410.7744*, 2014.
16. Primault (V.), Mokhtar (S. B.), Lauradoux (C.) et Brunie (L.). – Time distortion anonymization for the publication of mobility data with high utility. – In *2015 IEEE Trustcom/BigDataSE/ISPA volume 1*, pp. 539–546. IEEE, 2015.