EDEN: Enforcing Location Privacy through Re-identification Risk Assessment: A Federated Learning Approach

Besma Khalfoun, Sonia Ben Mokhtar, Sara Bouchenak, Vlad Nitu.

Abstract

Crowd sensing applications have demonstrated their usefulness in many real-life scenarios (e.g., air quality monitoring, traffic and noise monitoring). Preserving the privacy of crowd sensing app users is becoming increasingly important as the collected geo-located data may reveal sensitive information about these users (e.g., home, work places, political, religious, sexual preferences). In this context, a large variety of Location Privacy Protection Mechanisms (LPPMs) have been proposed. However, each LPPM comes with a given set of configuration parameters. The value of these parameters impacts not only the privacy level but also the utility of the resulting data. Choosing the right LPPM and the right configuration for reaching a satisfactory privacy vs. utility tradeoff is generally a difficult problem mobile app developers have to face. Solving this problem is commonly done by relying on a *trusted proxy server* to which *raw* geo-located traces are sent and privacy vs. utility assessment is performed enabling the selection of the best LPPM for each trace. In this paper we present EDEN , the first solution that selects automatically the best LPPM and its corresponding configuration without sending raw geo-located traces outside the user's device. We reach this objective by relying on a federated learning approach. The evaluation of EDEN on five real-world mobility datasets shows that EDEN outperforms state-of-the-art LPPMs reaching a better privacy vs. utility tradeoff.

Keywords: Location Privacy, Crowd Sensing Applications, Protection Mechanism, Re-identification Attack, Mobility Data, Data utility, Federated Learning

1 Introduction

The wide propagation of connected devices (75 billion in 2025 as forecast by various studies[1]) equipped with increasingly rich sensory capabilities and connected to fast networks (e.g., 5G) have contributed to the apparition of a wide variety of crowd sensing applications [13]. These applications offer useful services to their users (e.g., traffic congestion monitoring [5], noise level monitoring in urban areas [37], air quality monitoring [19] or assessing the radioactivity level near nuclear sites [12]). The principle behind a crowd sensing application is that a set of (paid or volunteer) users carry a device equipped with a GPS and an environmental sensor (e.g., an NO2 sensor), which could be their own smartphone or a dedicated device. Along their journey, the application collects timestamped, geo-located traces with the corresponding environmental measurements (e.g., pollution measurements). Then, it periodically sends this data to a central server called the Mobile Crowd sensing Server (MCS), which aggregates the collected data and provides updated maps to its clients (e.g., live pollution monitoring maps). The downside of these applications is that the collected data may constitute a serious threat to the participants' privacy if this data falls between the hands of curious/malicious adversaries. Indeed, various studies have shown that location data may leak sensitive information about their originating user [33, 61]. 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 [27, 21]. Furthermore, the disclosure of a user's identity is also jeopardized by re-identification attacks, i.e., attacks where an anonymous mobility trace is re-associated to its originating user based on previously collected data [39, 53, 23]. For example, the journalists from The New York Times were able to re-identify and track the whereabouts of ex-president Trump from a dataset of more than 50 billion location pings from more than 12 million user mobile devices [59].

To overcome the above threats, the research community has been actively proposing Location Privacy Protection Mechanisms (LPPMs). Examples of proposed LPPMs include Geo-indistinguishability [3], which enforces differential privacy [20] by adding spacial noise to a user's GPS coordinates, Promesse [54], which removes places where the user stops for a significant period of time and Trilateration [29], which generates dummy locations to obfuscate the user's real location. In this context, a problem that mobile app developers aiming at enforcing privacy-by-design have to solve is : **"how to objectively compare the privacy vs. utility tradeoff offered by different LPPMs and choose the right one ?"** For instance, how to decide whether an LPPM enforcing k-anonymity [57] (with a given value of k) is better than another one enforcing ϵ-differential privacy [20] (with a given value of ϵ)? To answer this question, the regulator (e.g., the EU General Data Protection Regulation in article 35) requires to carry out privacy risk assessment, which in our context translates into assessing

[1] https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/

which solution yields the smallest re-identification risk. In practice, solutions that have been explored in the literature to select among a set of LPPMs generally rely on re-identification attacks [31, 56]. Specifically, these solutions apply various LPPMs on a given trace and choose the LPPM (and its corresponding configuration) that better resists a given set of re-identification attacks [23, 53, 39]. The role of these attacks is to link anonymous traces to past user data. However, to reach this objective, the proposed solutions assume a *trusted proxy server* as existing re-identification attacks are centralized: they build user profiles using past unprotected mobility data and use them to estimate to whom a given protected trace belongs to.

In this paper, we overcome this assumption and propose EDEN , the first effective privacy-preserving solution for mobility data that performs re-identification risk assessment without requiring to send raw data to a remote server and that provides high data utility. Specifically, EDEN operates in two phases: (1) a phase on which a re-identification risk assessment model called FURIA is periodically trained on the users' devices and (2) a second phase where the latest computed FURIA model is used along with utility metrics to choose the best LPPM (among a set of off-the-shelf LPPMs and corresponding configurations) each time a user wants to send a geo-located trace to the MCS. To avoid centralizing raw data in a trusted proxy server, we design FURIA using a Federated Learning (FL) paradigm [8]: a relatively new machine learning technique which proposes to distributively train the models right where the data is created, i.e., on the user mobile devices.

We extensively evaluate EDEN using three real-world mobility datasets. We compare the performance of EDEN , both in terms of privacy and utility to the one of three off-the-shelf LPPMs using three configurations for each LPPM to cover the spectrum from strong privacy guarantees (despite the resulting impact on data utility) to strong utility objectives (with weaker privacy guarantees). For measuring the privacy offered by EDEN , we implement three state-of-the-art re-identification attacks (i.e., POI-attack [53], PIT-attack [23] and AP-attack [39]) and combine them in a single, stronger attack that relies on majority voting. This attack is run on the MCS (considered as an adversary) and is different from the re-identification risk assessment model trained by EDEN . For measuring utility, we use two types of metrics: a quantitative metric and a qualitative metric. The quantitative metric, i.e., area coverage, evaluates how far the area covered by a protected mobility data overlaps with the one of the original data. The qualitative metric captures the degradation in pollution measurements taken from a fourth real-world air pollution dataset [6]. An additional qualitative metric is considered, i.e., range queries. It counts the number of users going through regions. This metric is useful for analyzing traffic congestion in a city, and it is taken from a fifth cab drivers mobility dataset in the city of San Francisco. In addition to comparing EDEN to state-of-the-art LPPMs, we consider two extreme solutions: a Privacy Oracle, which knows the attack run by the MCS and chooses the best LPPM accordingly and a Utility Oracle (referred to as NOBF in the paper), which sends pseudonymized raw data to the MCS. The results show that EDEN provides a better tradeoff between privacy and data utility compared to individual LPPMs. To ease the reproducibility of our results, our code, the implemented LPPMs and attacks code as well as the scripts used to run the experiments are available as open source [2].

The remaining of this paper is structured as follows. First, we present the necessary background and related work and illustrate our research problem in Section 2. Then, we describe the system model and an overview of our solution in Section 3. Further, in

[2] https://github.com/MobDataPriv/EDEN

Section 4, we present a detailed description of EDEN and FURIA. An experimental evaluation of our solution is then presented in Section 5 and finally, we conclude the paper in Section 6.

2 Background and Related Work

In crowd sensing applications, users contribute geo-located data, which contains the user ID (e.g., the device MAC address), the user location (i.e., GPS latitude and longitude), the time at which the data has been collected, and the actual environmental measurement (e.g., explosion alert, NO_2 measurements). Despite the pseudonymization of the user identity (i.e., replacing the user ID with a pseudonym) and techniques to hide the IP address of the originating device (e.g., by using anonymous communication protocols such as TOR [55]), sharing geo-located data may still leak information about users as discussed in the following section.

2.1 User Re-Identification Attacks

User mobility data is a fingerprint that can be used to perform user re-identification attacks [18]. A user re-identification attack (or a de-anonymization attack) aims at associating an anonymized (and/or obfuscated) mobility trace to its originating user, based on a previously collected background knowledge from which the attack builds a set of user profiles. Each time an anonymous mobility trace is received by the attacker, the latter tries to re-associate it to the closest previously built user profile. Various user re-identification attacks have been proposed in the literature. What distinguishes these attacks is the way they represent user profiles. For instance, POI-attack [53], uses a set of Points Of Interest (POIs) to represent user profiles. POIs are locations where users spend a significant amount of time such as at home. PIT-attack [23] uses Mobility Markov Chains to synthesize mobility traces. In this attack, nodes represent POIs while edges represent the transition probability between POIs. Finally, AP-attack [39] represents user profiles as a heatmap. A heatmap is an aggregate representation of mobility data inspired by [46]. To construct the heatmap, first, the world map is divided into cells (i.e., geographical regions) of relatively equal size. Then, for each cell, the probability of visiting the cell for a mobility trace T is computed. More precisely, it corresponds to the number of records in T present in the cell divided by the total number of records in T. At the end, we obtain for each user, a probability distribution of visited cells.

User re-identification attacks are generally used as a privacy metric to compare LPPMs. However, existing attacks are generally centralized as they need to collect past users' data to build user profiles. In this paper, we propose a federated way to conduct re-identification risk assessment for comparing LPPMs and thus choosing automatically the right one, i.e., the one which passes the re-identification test.

2.2 Related Work on Location Privacy Protection Mechanisms

Mitigating threats affecting location privacy has attracted the interest of many researchers. As a result, various Location Privacy Protection Mechanisms (LPPMs) have been proposed in the literature [1, 2, 29, 3, 54, 52]. An LPPM can be defined as

a function which takes as input one or multiple mobility records of a given user and produces as output an obfuscated version of this data.

LPPMs can be classified in two categories according to the data they need to protect a given mobility trace: (1) those that need knowledge about the mobility of other users and (2) those that do not need any external knowledge. The first category of LPPMs leverage the mobility of other users either to obfuscate a mobility trace a way that makes it closer to the profile of another user than the originating user [38], this is known as a profile conversion technique, or to achieve formal guarantees such as k-anonymity [57, 1, 2]. The concept of k-anonymity is first introduced by Sweeney [57]. It states that a user is hidden among at least $k-1$ other users with similar properties. In the context of location privacy, this translates to cloak a given user exact location in a geographical zone (called cloaking region) where there are at least $k-1$ other users. Numerous LPPMs that enforce k-anonymity are proposed in the state-of-the-art, among them, CliqueCloak [25], Casper [44] which compute cloaking regions on each geo-located point without considering a mobility trace as a whole. Alternatively, NeverWalkAlone [1] and its extension W4M (Wait for Me) [2] extend the concept of k-anonymity to k-δ-anonymity. They ensure that a user mobility trajectory is always hidden among $k-1$ other trajectories inside a cylindrical volume of radius δ where users move. Furthermore, solutions enforcing k-anonymity have been criticized as sensitive information can be leaked if the k users are co-located in the same semantically sensitive location (e.g., a hospital). For this reason, l-diversity in general and location diversity in particular are introduced [62, 36]. The latter states that in a given cloaking region, there should be at least l different semantic locations. One of the major limitations of these LPPMs is that they rely on a centralized trusted proxy server, since they need to access the overall raw data. Another limitation is that they lack an actual privacy risk assessment before sharing the data with the MCS.

In contrast, the second category of LPPMs are implemented on the user-side as they do not need an external knowledge to protect a given trace. For instance, Huang et al. [29] propose Trilateration (*TRL*), a new way to achieve k-anonymity by generating dummy location data on the user's device. For each real location l of the user, it randomly generates three locations l_1, l_2 and l_3 in a range of r from the real one. The fake locations are sent to the MCS instead of the real position. Dwork introduces the concept of differential privacy as a formal privacy guarantee [20] for database systems. The idea is that an aggregate result over a database is not affected by the presence or absence of a single element in the table. Andres et al. [3] adapt this concept in the context of location privacy. They propose Geo-indistinguishably (*Geoi*), a mechanism that perturbs the spatial information of data by adding Laplacian spatial noise to each GPS coordinate. The amount of noise is calibrated by a privacy budget ϵ (the lower the ϵ, the higher the privacy level). This perturbation is done on the user-side without the implication of any external proxy server. In the same direction, PROMPTS [10], exploits the coresets theory, another specific privacy guarantee that helps users to evaluate their privacy exposure locally, before sharing their geo-located data. Promesse (*PROM*) [54] uses a speed smoothing technique to erase user POIs from the mobility trace. This ensures that spatiotemporal points in the obfuscated mobility trace are equidistant thanks to a parameter α which tunes this distance. FOUGERE includes several LPPMs that can be chosen by the users [41], but it does not provide any means for selecting the most appropriate LPPM nor its proper configuration. This may result in under protected data, or in low data utility. Finally, these solutions are applied blindly without any privacy risk assessment. Thus, other existing systems aim at combining existing LPPMs to propose a personalized protection approach with

privacy risk assessment. For example, MooD is a user-centric fine-grained protection system based on the combination of multiple off-the-shelf LPPMs to protect a mobility dataset in front of re-identification attacks [31]. ALP is a framework that enables an automatic configuration of the LPPM parameters using optimization algorithms such as simulated annealing [51]. PULP is another system, which automatically configures LPPMs according to users' objectives in terms of privacy and utility [14]. For the same purpose, authors in [56] leverage machine learning techniques (i.e., GANs) to obfuscate mobility data. They use a generator network to produce noise for mobility data perturbation, and a discriminator classifier to evaluate the re-identification of the perturbed data. However, these systems follow a centralized approach with a trusted third party since they need to access the overall raw data. In contrast, EDEN does not require a trusted third party, and finds for each mobility data the best LPPM and its configuration among various LPPMs, thanks to FURIA risk and utility assessment.

2.3 Problem Illustration

Consider an app developer, say Bob, who has to integrate privacy-by-design in a crowd sensing application. Bob needs to choose an LPPM with the appropriate configuration to protect users' mobility data before sharing it with a MCS. Bob does not want to implement yet another LPPM as there already exists a variety of LPPMs proposed by the research community. However, the actual level of privacy offered by each LPPM and the impact of the latter on data utility can dramatically vary according to how these LPPMs are configured.

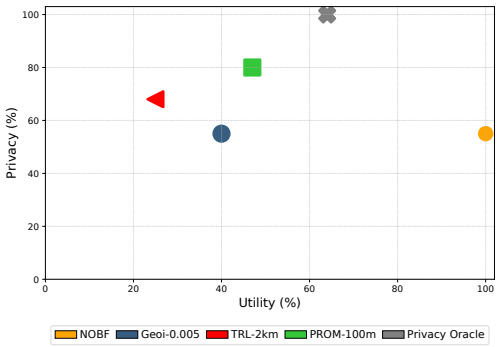

Fig. 1: Impact of LPPMs on Privacy vs. Utility on PrivaMov Dataset illustration.

To better illustrate this problem, we perform (on behalf of Bob) an experiment on the PrivaMov Dataset [45], a real world mobility dataset. In this experiment, we compare the privacy vs. utility tradeoff offered by three state-of-the-art LPPMs : Geo-Indistinguishability[3], Trilateration[29] and Promesse[54] noted *Geoi, TRL* and *PROM*, respectively. In order to find a satisfactory compromise between privacy and utility, we configure each of these LPPMs for enforcing an average privacy level, i.e., $\epsilon = 0.005$, $r = 2km$ and $\alpha = 100m$ for *Geoi, TRL* and *PROM*, respectively. We provide more details about further configurations in Section 5.3.4. In addition to these three LPPMs, we also evaluate two alternatives: NOBF, which represents pseudonymized raw data without any additional obfuscation and Privacy Oracle, which is an oracle

that selects the best LPPM for each individual trace (the LPPM which prevents re-identification and maximizes the data utility).

We measure privacy as the ratio of user mobility data which is not re-identified by the MCS-side attacker over the overall mobility data. The attacker applies a majority voting over three state-of-the-art attacks, namely, *(1)* POI-Attack [53], *(2)* PIT-Attack [23] and *(3)* AP-Attack [39]. We refer to this attack as *Mv-Attack*. On the other hand, we evaluate utility with the Area Coverage (AC) metric, which computes the overlap between the obfuscated and the original mobility trace using F-score [51]. In Figure 1, an aggregate value of AC is depicted in the x-axis. It is computed as in Equation 1; where $F = \begin{pmatrix} 0 & 0.25 & 0.5 & 0.75 & 1 \end{pmatrix}$ refers to the vector of utility factors and $U = \begin{pmatrix} u_0 & u_{0.25} & u_{0.5} & u_{0.75} & u_1 \end{pmatrix}$ refers to the data proportion with $AC = 0$ or $AC \in$]0 0.25] ,]0.25 0.5] ,]0.5 0.75] and]0.75 1] of the raw mobility data, respectively.

$$Utility = F^T.U \tag{1}$$

Note that Equation 2 is verified:

$$\sum_{i \in F} u_i = 1 \tag{2}$$

If no LPPM protects against *Mv-Attack*, the Privacy Oracle chooses to drop the data instead of sending it to the MCS. Privacy Oracle constitutes the best choice that can be done (from a privacy perspective) if Mv-attack is known by the defender and if all the data is centralized in a trusted proxy server. On the other side of the spectrum, NOBF is the best choice that can be done to preserve the data utility.

Results are depicted in Figure 1. In these results, we can observe that the Privacy Oracle protects 100% of the data and provides 65% of AC, which is better both in terms of privacy and data utility compared to individual LPPMs. In practice, the Privacy Oracle finds an LPPM resisting the attack for 96% of the traces and drops 4% of the remaining traces.

Therefore, we conclude that a Privacy Oracle has the potential to outperform all other state-of-the-art LPPMs in terms of privacy vs. utility tradeoff. However, the latter assesses the privacy using the *Mv-Attack*, which needs to centralize the raw data to a proxy server. The goal of this paper is to design a solution that is as close as possible to the Privacy Oracle without centralizing the raw data in a proxy server.

3 System Model

In this section, we first introduce federated learning (Section 3.1). Then, we provide an overview of EDEN (Section 3.2) and present the threat model (Section 3.3).

3.1 Overview of Federated Learning

Federated learning (FL) is a machine learning paradigm where many clients (i.e., workers) collaboratively train a model under the coordination of a central server, also known as the federator [40]. Each client's raw data is stored locally at the client side, and it is not transmitted to any other party. In a traditional FL protocol, the federator server initializes the global model and sends it to the clients. A client trains his local model based on his local data, and then the model gradients (i.e., weights) from the individual models are sent to the central server, that aggregates the gradients and updates the global model. The latter is sent back to the clients' devices and the training process can then be repeated until a desired level of accuracy is attained. It also exists

more advanced FL protocols where a central server is not required and the model is decentralized over multiple parties where a Secure Multiparty Computation (SMC) is made to update the global model [30]. Nowadays, federated learning has several applications such as Google Gboard [28, 16, 63].

3.2 EDEN Overview

EDEN is a user-side mobility data protection system for crowd sensing applications, which operates as depicted in Figure 2. Let us consider Alice (depicted in the center of the figure), a participant of a crowd sensing campaign. Along her journey, Alice collects geo-located environmental measurements. At a given point in time, the crowd sensing application decides to send the collected data (depicted at the bottom of the figure) to the MCS. Before sending this data to the MCS, EDEN automatically sanitizes the data without the implication of Alice. It applies a given LPPM among a set of available choices. For each LPPM, various configurations are considered by EDEN going from configurations that enforce strong privacy to configurations that rather try to preserve data utility.

Specifically, EDEN applies each LPPM to the raw trace and evaluates both : (1) the re-identification risk of the trace using this LPPM and (2) the corresponding data utility. For evaluating the re-identification risk, EDEN uses FURIA, a federated learning model, which is trained as depicted in the left part of Figure 2. When Alice's device fulfills a set of predefined requirements to participate in the FURIA training process (e.g., her device is idle, charging and connected to WiFi), it downloads from a server called the FURIA Master Server, the latest FURIA global model, trains the model using its locally collected data and sends the updated model back to the server. Once the server receives a predefined number of user responses, it aggregates them into a new version of FURIA global model. This way, FURIA continuously learns and dynamically improves its global knowledge with new discriminating mobility patterns of users. The training process of FURIA is performed by night independently from the process of protecting mobility traces using EDEN which is depicted on the right part of Figure 2. In this part of the figure, EDEN prepares batches of protected mobility data and sends a batch periodically (if any) to the MCS. This batch of geo-located data has been protected by using an LPPM for which the re-identification risk assessment performed using FURIA passed; i.e., FURIA could not re-identify Alice as the originating user of this trace. If two LPPMs (or two variants of the same LPPM) pass the risk assessment, the one that has the best data utility is chosen. If no LPPM passes the FURIA risk assessment, then EDEN makes a decision according to a given configuration policy. For instance, it would drop the data if it is configured with a conservative policy. We describe other policies than the conservative one in Section 4.2.1. In EDEN , The configuration policy choice can be set by the participant according to her preferences.

In this paper, we focus on three state-of-the-art LPPMs each having three configurations and we use three utility metrics as further described in Section 5. Though, EDEN is not tight to a given set of LPPMs or utility metrics. More LPPMs, with their corresponding configurations and more utility metrics whether quantitative or application-dependent can be easily integrated in EDEN .

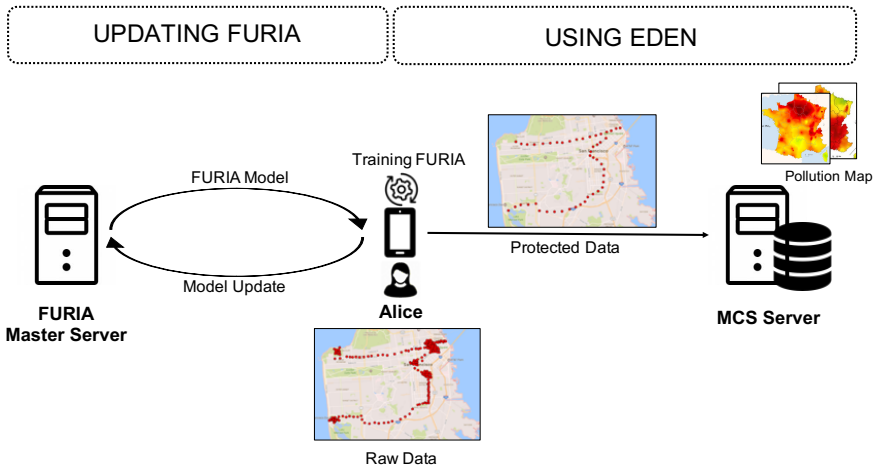

Fig. 2: System Architecture

3.3 Threat Model

As depicted in Figure 2, EDEN uses FURIA to assess the re-identification risk of a protected trace. FURIA is designed following a traditional federated learning protocol as presented in Section 3.1 using a master server. The communication channels between the clients and this server are encrypted. This server never accesses users' raw data, but rather aggregates model updates coming from various users.

Furthermore, we assume that users' devices are trusted, and that the data protected by EDEN is sent anonymously to the MCS (e.g., using an anonymous communication protocol such as TOR [55]). We consider the MCS as an adversary and we assume it to be *honest-but-curious* (i.e., semi-honest [50]). Specifically, the MCS collects and processes geo-located environmental measurements to produce aggregate data to its clients. It performs this task honestly, i.e., without deleting or altering the received data. However, it is curious because he may exploit the received data to learn valuable private information which may interest him or any third party (to whom he might sell the data for advertising purposes). For instance, the adversary may conduct inference attacks on the received data and consequently reveal the user identity or other sensitive attributes (e.g., POIs, social links, etc.), even if the participant is using the application anonymously [32]. In particular, we assume that the MCS tries to link the received data to previous user profiles he has built from leaked background knowledge. To reach this objective, we consider that the MCS implements the latest available user re-identification attacks he finds in the literature. In this paper, we assume that the MCS combines the three attacks described in Section 2, namely, AP-Attack, POI-Attack and PIT-Attack. The attacker runs these attacks by relying on user profiles built from past mobility data of the participating users. Then upon receiving a trace protected by EDEN , it runs the three attacks and performs a majority voting between the predicted values of these attacks and returns the identity label that received more votes. Finally, complementary techniques can be applied with EDEN in order to counter attacks that attempt to leak information from model updates in federated learning [26, 48], or to counter malicious users that poison data to introduce

a backdoor into the global model [22, 4].

4 EDEN Detailed Description

In this section, we dive into a detailed description of how the re-identification risk assessment is done with FURIA and how mobility traces are protected with EDEN.

4.1 Re-identification Risk Assessment with FURIA

FURIAS's global model is a crucial part of our solution. It applies federated learning to build a re-identification risk assessment model. The latter learns discriminating mobility patterns that uniquely identify users and help our system to assess LPPMs in a privacy-preserving way. As depicted in Figure 4, FURIA involves two parties: (1) *mobile user devices* where raw mobility data is stored and where model updates are computed, (2) the *FURIA Master Server* where model updates provided by various users are aggregated. FURIA operates as follows. First, the *FURIA Master Server* initializes a classification algorithm with random values. In this paper we use *Logistic Regression*, a simple yet effective classification model that satisfies very well our objectives, after an empirical experiment. Precisely, we compare three methods of classification, namely, Logistic Regression (LR) [42], Random Forest (RF) [11] and a multi-layer perceptron Neural Network (NN) [47]. Figure 3(a) shows the re-identification rate over three unprotected datasets: Geolife, MDC and PrivaMov, described in Section 5.2. LR is slightly better than RF with +3% of re-identification rate in Geolife and both LR and RF are better than NN with up to +11%. Thus, the retained model for the rest of our work is LR.

The *FURIA Master server* sends this model to all participants (step ① and ② in Figure 4). This model is denoted as AF_0 (Attack Federated). Each participant U_j transforms its raw mobility data of the day to feature vectors (step ③) and trains the model AF_0 locally on the generated feature vectors. Once all participants have finished the first learning round, they send their local updates (i.e., gradients) of their current local models to the *FURIA Master Server*, (step ④). Upon receiving model updates, the *FURIA Master Server* aggregates users' gradients and produces an updated model, denoted AF_1, ready to use at the following day. This process is iteratively done and generally takes place at night time in order to avoid any interference with other applications running on the user's device. Indeed, model updates are computed when the user's device is idle, plugged in and connected to WIFI, which is generally the case at night time. FURIA processing is inspired by active/online learning where daily mobility data is unrolling between the train set of the current learning round and the testing set of the following one.

In FURIA, three types of features are considered: (1) spatial features (2) temporal features and (3) aggregated features. ***Spatial features***. To synthetically capture spatial information, records (i.e. lat and lng) are projected on a heatmap. A heatmap is a set of cells of equal size. For each cell, the proportion of mobility records in a given trace T that belongs to that cell is computed. This corresponds to the cell visit rate. ***Temporal features.*** 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, the above two users would have similar heatmaps but at different times of the day. In FURIA we use a simple, yet effective temporal information, which is the average time of the day

of all the records in a given trace T. This is convenient in the case of crowd sensing applications, where mobility traces are generally in the order of minutes or hours length without exceeding a day. ***Aggregated 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 is considered. This allows to represent service usage intensity. We also extract the centroid of the mobility trace T (i.e., centroid's latitude and centroid's longitude), to capture a central position of the user's mobility in a map.

The spatial features are usually considered in related work [39, 35, 46]. In addition to this type of features, we explore several other types of features (e.g., temporal and aggregated features) and evaluate their benefits for a user re-identification attack, as presented in Figure 3(b). The results show that temporal and aggregated features improve the re-identification rate of the attack. Specifically, the use of combined spatial, temporal and aggregated features increases the re-identification success rate with +2%, +7% and +8% in comparison to the use of only spatial features in Privamov, Geolife and MDC datasets, respectively.

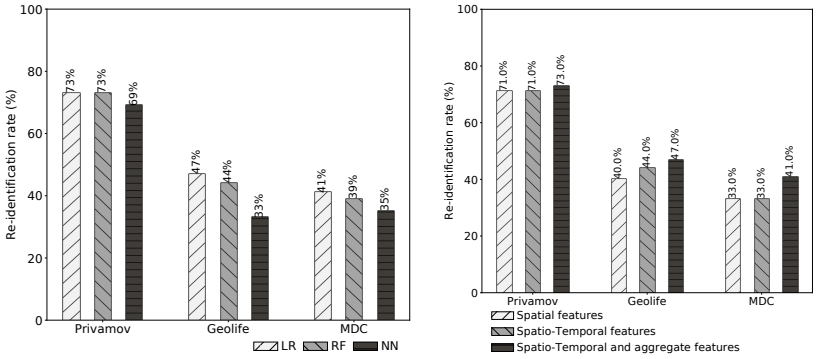

(a) Model variation (b) Feature variation

Fig. 3: Empirical experiments.

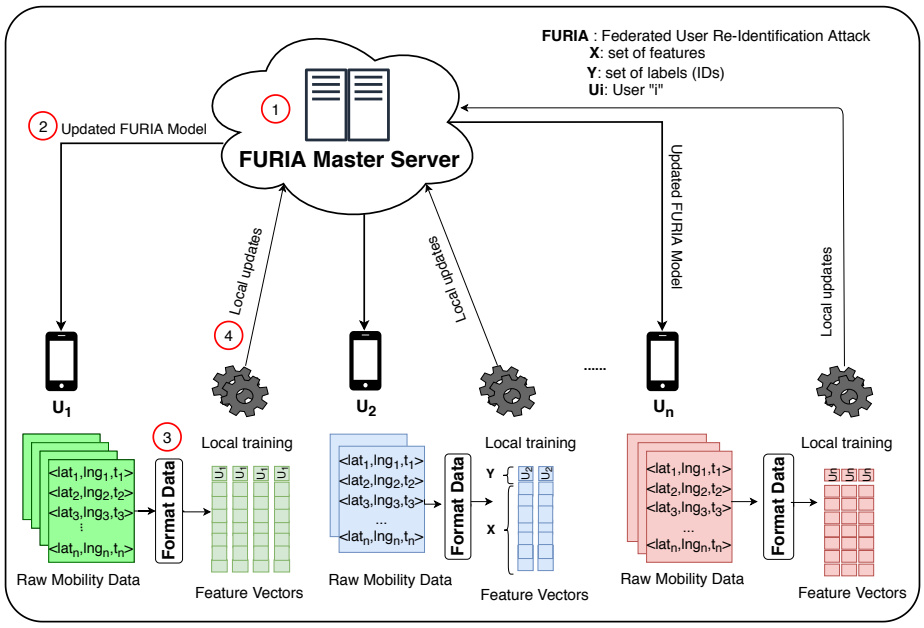

Fig. 4: FURIA architecture

4.2 Protecting Mobility Traces with EDEN

The detailed architecture of EDEN is depicted in Figure 5 and its behavior is described in Algorithm 1 in the Appendix A. EDEN takes as input a user mobility trace T and a set of LPPMs $\mathbb{L}$ with various configurations (i.e., low, medium and high impact on privacy vs. utility tradeoff) and returns as output a protected mobility trace T_i', which will be sent to the MCS. EDEN has four main components, the first component *"Apply LPPMs"* iteratively applies all the LPPMs implemented in EDEN with their configuration variants on the raw mobility trace T stored on the user smart phone (step ① in Figure 5). As a result, it produces a set of obfuscated versions of the same raw mobility trace i.e., $C = \{T_1', T_2', \ldots, T_n'\}$. The second component *"Format Data"* transforms the different obfuscated traces available in the set C into feature vectors as described in Section 4.1 (step ②). The third component *"Global Model FURIA"* uses the latest version of FURIA's Global Model AF_i and uses it as a privacy metric. Specifically, it evaluates the user re-identification risk of each feature vector of the obfuscated data, (step ③). If the model fails in predicting the right identity label associated to the transformed mobility trace, the latter is potentially elected to be sent to the MCS. Otherwise, if the model succeeds in predicting the right identity label associated to the mobility trace for all the considered LPPMs, three different policies can be adopted by EDEN (step ④) as described in Section 4.2.1.

Finally, the last component in EDEN is *"Best coverage"*. It selects the protected mobility trace candidate that maximizes data utility. EDEN can consider various utility metrics. In the current version of our system, we use the AC metric [51]. It is computed between the original and the obfuscated mobility trace, T and T_i' respectively, to measure how much the alteration caused by an LPPM affects the

regions visited by a user (step ⑤). We provide more details about AC metric in Section 5.3.1. Finally, the obfuscated mobility trace that better resists the re-identification test performed by FURIA and that has the best utility is sent to the MCS. The latter processes the received data and produces useful information to users (step ⑥).

4.2.1 EDEN Policies

Three policies are considered by EDEN if a mobility trace is re-identified by FURIA. The first policy is **EDEN-pessimistic** (EDEN-pes): this policy is the most conservative policy as it simply deletes a mobility trace that FURIA is able to re-identify. The rationale behind this policy is that if FURIA is able to re-identify a mobility trace, an external attacker could very well reach the same result. The downside of this solution is that it causes data loss from the application perspective.

The second policy considered in EDEN is **EDEN-optimistic** (EDEN-opt): an opposite solution to the previous one where the mobility traces that FURIA is able to re-identify are also sent to the MCS without applying any LPPM or by applying a default LPPM. We use this policy as a baseline to assess the impact of sending traces despite the red flag raised by FURIA regarding the re-identification risk of some mobility traces.

The third and last intermediary solution is **EDEN-balanced** where sending or dropping mobility traces is based on a local metric evaluating how a given mobility trace under consideration is similar to past mobility traces of the same user. Towards this purpose, we use the Topsoe divergence metric [15]. The latter is computed between two probability distributions : (1) the heatmap of the current raw mobility trace T and the heatmap corresponding to the past mobility data of the same user, which is stored on the user's device, (see Equation 3). It is a derived symmetric version of the Kullback Leibler divergence [15] which measures the information deviation between the user past mobility and the current one. If the deviation is high (greater than a threshold empirically set to 0.8), the mobility trace is sent to the MCS, otherwise it will be deleted. The set value of the threshold fits well the distribution of Topsoe deviation values.

$$d_{Topsoe}(P,Q) = \sum_i \left[P_i \ln\left(\frac{2P_i}{P_i+Q_i}\right) + Q_i \ln\left(\frac{2Q_i}{P_i+Q_i}\right) \right] \tag{3}$$

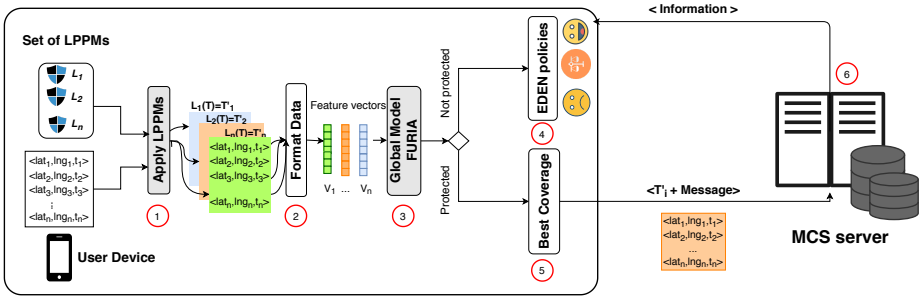

Fig. 5: EDEN architecture

5 Experimental Evaluation

We start this section by describing our experimental setup (Section 5.1), the mobility datasets we use to evaluate EDEN (Section 5.2), and our evaluation scenario (Section 5.3). Then, our evaluation answers the following questions: (1) how does EDEN compare to state-of-the-art LPPMs against an attack performed by the MCS? (Section 5.4); (2) what is the impact of EDEN on data utility compared to state-of-the-art LPPMs? (Section 5.5); (3) what is the privacy vs. utility tradeoff reached by EDEN (Section 5.6)? After answering these questions, our evaluation dives into a fine-grained analysis of EDEN and its configuration policies (Section 5.7) and finishes with the run-time overhead of EDEN and training FURIA (Section 5.8).

5.1 Experimental Setup

All the experiments related to the MCS attacker are carried out on a server running an Ubuntu 14.04 OS with 50GB of RAM and 16 cores of 1.2GHz each. Both EDEN and FURIA are developed in Python using the Pytorch library [3]. We used S2Geometry library [4] for the decomposition of the map into cells of approximately equal size. The cell edge length ranges from 212m to 296m . To accelerate the training process of our federated learning model, we use a machine with NVIDIA TESLA V100 GPU. Participating users are simulated by considering the data coming from real-world mobility datasets.

5.2 Mobility Datasets

In our experiments, we use three real-world publicly available mobility datasets with a summary given in Table 1. These datasets are : *(1)* Geolife [64] that contains the mobility of users in the city of Beijing; (2) MDC [34] that contains the mobility of users in the city of Lausanne and Privamov [45] that contains the mobility of users in the city of Lyon. In our experiments, we extract only the most active month (i.e., 30 days) of each dataset for a fair comparison. In the context of location privacy, these datasets are used by many state-of-the-art LPPMs in order to assess the effectiveness of their approach [54, 31, 38, 53]. That is why, we decide to evaluate our approach on these datasets to be in line with the community.

Tab. 1: Description of datasets

Name	Geolife	MDC	Privamov
# users	42	144	48
location	Beijing	Geneva	Lyon
# records	1,468,989	904,282	774,401
area (km^2)	16,808	41.37	47.87

[3] https://www.pytorch.org
[4] https://www.s2geometry.io

5.3 Evaluation Scenario

We simulate mobile users that correspond to the users of the three datasets described in Section 5.2. We assume that the data corresponding to the first 15 days of each dataset have been leaked to the MCS. Using this data, the MCS builds user profiles. To be in the same conditions as the adversary, the same data is used to train the first FURIA model (i.e., AF_0). The remaining 15 days of each dataset is then used as a test set. Specifically, FURIA is inspired by active/online machine learning. Its training/testing mobility data is unrolled with a time window of 24 hours. For example, the 16th day mobility data is used to test AF_0 and to train AF_1, the 17th day mobility is used to test AF_1 and train AF_2 and so on. It means that our training set of mobility data is incremented day by day and many phases of test occur on each newly trained FURIA model AF_i.

Moreover, as sharing data with the MCS in real time is energy-consuming [24, 60], we assume that the user's crowd sensing application prepares batches of 30 minutes length to be as close as real-time data transmission use cases, protects this data using EDEN and periodically sends it to the MCS.

Upon receiving a geo-located trace, the MCS uses this trace to update its target map. Simultaneously, the MCS tries to re-associate the received trace to one of the user profiles it has previously built. To this end, it uses *Mv-Attack*. We compare EDEN to cases where the considered LPPMs are applied blindly on user mobility data or assisted with a Privacy Oracle.

5.3.1 Utility Metrics

To evaluate the impact of EDEN and its competitors on the quality of the generated data, there are two categories of utility metrics, proposed in the literature [52]. *(1) Data-centric or quantitative* utility metrics which measure the distortion between the original and the obfuscated mobility data. Examples of such metrics include spatial distortion [54], spatiotemporal distortion [38] where a spatial error is computed under a temporal constraint and finally the AC metric [51]. In this paper, as previously mentioned in Section 4.2, we use the AC metric. To recall, it computes the overlap between the obfuscated and the original mobility trace using the F-score. This metric is able to capture the degradation in data utility caused both by LPPMs that remove data points (e.g., PROM), and the degradation caused by LPPMs that add data points or move them spatially (e.g., TRL and Geoi). Thus, this metric can be used in various applications, such as in transportation mobile applications where a data analyst can use the AC metric to adapt the availability of public transportation in areas according to visiting user density. The AC utility metric can also be used in the context of a pandemic, where a non-infected user can be aware of an area that she visits at the same time as other infected users.

(2) Application-centric or qualitative utility metrics, which compare the result of a given application before and after applying an LPPM for a specific application. In this paper, we consider two real-world use cases. In the first use case, we visualize the air pollution degradation map before and after the application of EDEN and state-of-the-art LPPMs on a crowd sensing air pollution dataset [6]. This map allows to detect areas where the level of gaseous pollutants is high (NO_2, CO) (i.e., hotspots). This dataset is described in Section 5.5.2. In the second use case, we use range queries metric, a classical operation which compares the number of unique users who go through areas during a time window before and after obfuscation [54]. An illustrative

example is provided in Figure 6, where two range queries **Q1** and **Q2** of different radius are performed (the temporal dimension is not represented). Before the obfuscation process, **Q1** and **Q2** return 3 and 1 users, respectively, whereas after obfuscation, **Q1** and **Q2** return 1 and 3 users. Thus, the utility is measured as the range query distortion defined in [2]. In our example, the distortion of **Q1** is $\frac{|3-1|}{3} = \frac{2}{3}$ and the distortion of **Q2** is $\frac{|1-3|}{1} = 2$. Then, the average query distortion is computed, i.e., $\frac{\frac{2}{3}+2}{2} \approx 1.333$.

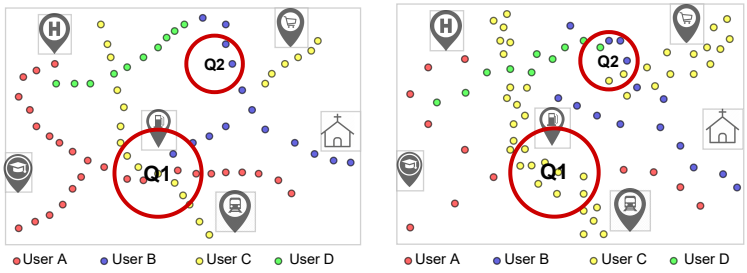

(a) Before obfuscation (b) After obfuscation

Fig. 6: Illustrative example of range queries metric.

5.3.2 FURIA Configuration

To build FURIA's global models, we opt for the *multi-class logistic regression* algorithm. The training of our global model is done over multiple rounds (R_i). Each round represents a 1-day training, except the first round (R_0) where a training set of 15 days is used. The latter is considered as historical data stored on the user device and previously leaked to the adversary (i.e., *Mv-Attack* has access to 50% of the mobility dataset to construct user profiles). Thus, we decide to start the training process with the same knowledge of the adversary to be in the same conditions. We assume that users train the model at the end of the day (every night) with the data collected during that day in order to prepare the model of the next day. Thus, the model is actively trained by incoming mobility data and improves its global view day by day. In each round, we run 100 epochs/user with a variable batch size. We tune the batch size according to the number of collected traces per user participating in the given round. In addition, we use the Stochastic Gradient Descent (SGD) algorithm with a learning rate of 0.001. These values are fixed after several experiments.

5.3.3 User Re-identification Attack Configuration

Mv-Attack is made by the combination of three state-of-the-art attacks, namely: AP-Attack [39], POI-Attack [53] and PIT-Attack [23]. By combining three state-of-the art attacks, we obtain an attack that is stronger than considering the attacks separately as the adversary gets more confidence about the result of the re-identification process. Each attack has a set of parameters, described below. POI-Attack and PIT-Attack have two parameters for the extraction of POIs from mobility traces, [65]. These parameters are the diameter of the clustering area and the minimum time spent inside a

POI. These parameters are respectively set to 500 meters and 5 minutes to accommodate small traces. AP-Attack has a configuration parameter that corresponds to the square cell size. It was set to 800 meters (default value).

5.3.4 Competitors

To evaluate EDEN, we select three LPPMs from the state-of-the-art: *(1) Geo-Indistinguishably* (Geoi) [3], *(2) Trilateration* (TRL) [29] and *(3) Promesse* (PROM) [54]. We select these LPPMs because they can be run on the user side (i.e., without external knowledge) and they provide diverse guarantees: differential privacy, dummy-based obfuscation and POI erasure, respectively. These LPPMs are described in Section 2.2. Each LPPM has its own configuration parameters. These parameters have an impact on the balance between privacy and data utility. In our experiments, for *Geoi*, we set the privacy parameter ϵ to 0.01, 0.005 and 0.001. A lower value of ϵ leads to a higher level of noise added to mobility records and consequently ensures a higher privacy guarantee). For *TRL*, there is a circular region with a radius of r, that surrounds the real location of the user. The chosen values of this parameter are 1km, 2km and 3km. A higher value of r generates a bigger region for location dummies and consequently ensures a higher protection level. And finally, for *PROM*, which has a parameter α that specifies the distance between two locations points, we set α to 50, 100 and 200 meters. A higher value of α leads to a larger distance between locations in a mobility trace and thus ensures a higher protection level. However, the latter can cause serious data loss especially if a mobility trace is recorded in a short distance over a short period of time. In addition to the above three LPPMs and corresponding configurations, we also evaluate two baselines: a Utility-centric baseline, referred to as *NOBF*, which corresponds to sending the data to the MCS without obfuscation (i.e., sending raw, pseudonymized data); and a Privacy-centric baseline, referred to as Privacy Oracle. This baseline is only used in the evaluation of the privacy vs. utility tradeoff. It represents a solution where the selection of the best LPPM is driven by the attack performed by the adversary. As such, perfect privacy can be reached but the chosen LPPMs can still degrade data utility.

5.4 EDEN vs. Competitors

In this section, we evaluate the effectiveness of EDEN in terms of privacy, in comparison to state-of-the-art LPPMs. For that purpose, we measure the data protection rate of EDEN 's variants and its competitors against *Mv-Attack*. To recall from Section 2.3, data protection rate is the percentage of mobility data that is not re-identified by the MCS. Results are depicted in Figure 7. From this figure, we observe that on the Privamov Dataset, (Figure 7(a)), 55% of the data sent without obfuscation (*NOBF*) is not re-identified by the MCS. This percentage is the same when *Geoi-0.01* is used and slightly increases when *Geoi-0.005* and *Geoi-0.001* are used, i.e., 56% and 64% of protected data, respectively. The most privacy-protective LPPMs from the literature are *TRL* with an increased range r or *PROM* with a large distance α between points in the trajectory. The proportion of protected traces reach up to 74% when *TRL-3km* is used and 91% when *PROM-200m* is used. In the case of EDEN , 85%, 86% and 87% of protected mobility data are recorded with the optimistic, balanced and pessimistic variants of EDEN , respectively. We notice that *PROM-200m* outperforms EDEN 's variants with +5% on average, this is due to the fact that *PROM* is based on the re-sampling of mobility traces by suppressing points. Thus, if a mobility trace

does not exceed 50m, 100m or 200m of traveling distance, the latter will be deleted. The deleted data is not sent to the MCS provider and is thus considered as protected. However, this dramatically degrades data utility as further discussed in Section 5.5.

In the Geolife Dataset, (Figure 7(b)), 61% of mobility data is naturally protected against the Mv-attack. The application of *Geoi* and *TRL* with their different configurations does not improve the protection rate compared to the baseline. Using *PROM-50m*, *PROM-100m* and *PROM-200m* increases the protection rate with +4%, +6% and +10%, respectively. This is also due to the suppression of chunks where the mobility does not exceed 50m, 100m, or 200m in a laps of 30 minutes. On this dataset, using EDEN significantly improves the data protection rate reaching 87 % and 90% of protected mobility data with *EDEN-balanced* and *EDEN-pes*, respectively. Finally, in the MDC Dataset, (Figure 7(c)), 68 % of the mobility data is naturally protected. The application of EDEN ' variants improves the protection rate with +10% on average compared to the NOBF baseline. This result has the same trend for the other LPPMs except with *promesse* (*PROM-50m, PROM-100m* and *PROM-200m*) which provides a higher protection rate (+4 % in average). This result is due to the suppression of mobility traces.

Here, we notice that PROM-50m provides a better protection with the MDC dataset in comparison with Privamov and Geolife datasets. MDC involves 144 users in a quite small area (around 41 km^2). Such a high population density naturally reduces user uniqueness which, thus, enables higher protection. In contrast, the Privamov dataset has three times less users than MDC, and the Geolife dataset has only 42 users in a large geographical area (16,808 km^2). This makes users in these two datasets more distinguishable and, thus, harder to protect.

5.5 Impact of EDEN on Data Utility

As described in Section 5.3.1, we measure the data utility using two types of utility metrics : a quantitative metric which is AC metric (Section 5.5.1) and a qualitative metric where we capture the degradation in pollution measurements taken from a real-world air pollution dataset [6] (Section 5.5.2). In addition, we measure the number of cab drivers going through regions taken from a real-world mobility dataset [49] (Section 5.5.3).

5.5.1 Utility Evaluation Using AC Metric

In Figure 8, we evaluate the data quality of sent/not sent data to the MCS (the privacy dimension is not considered in this figure). In this figure, an AC equal to 0 corresponds either to the data deleted by the LPPM (e.g., in the case of promesse) or to protected data that has no intersection with the original data. Further, we split Area Coverage in four intervals. The best LPPMs are solutions that maximize AC in the interval $[0.75, 1]$ while minimizing the data loss, i.e., $AC = 0$.

In the Privamov Dataset, we notice that EDEN's heuristics produce a balanced data quality. Specifically, EDEN-opt, EDEN-balanced and EDEN-pes loose an average of 21.4 % of mobility data and produce an average of 46 % of the generated data with an AC ¡ 0.75. And finally, an average of 32 % of the generated mobility traces has an utility value greater than 0.75. However, *PROM* with its different configurations has predominately darker bars. Specifically, with *PROM-100m* and *PROM-200m*, respectively, 50 % 74 % of the generated mobility data has $AC = 0$. Promesse chooses to not share with the MCS a large proportion of the generated data. In contrast,

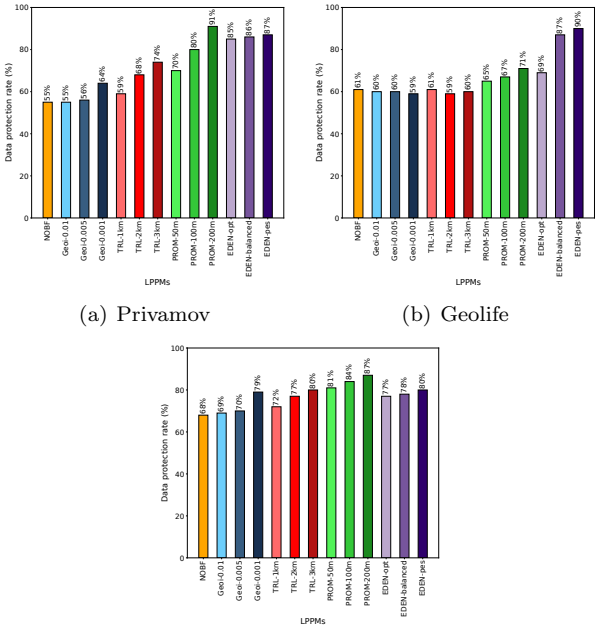

(a) Privamov (b) Geolife

(c) MDC

Fig. 7: Mv-Attack evaluation on EDEN vs. LPPM competitors.

the data chosen to be shared closely mirror the original mobility data: up to 65 % of protected mobility data by *PROM-50m* has an *AC* ¿ 0.75. *Geoi* and *TRL* degrade the quality of almost all generated data. Specifically, in *Geoi-0.01*, *Geoi-0.005*, *Geoi-0.001*, an average of 92 % of the resulting mobility data has an *AC* ¡ 0.75. Even worst with *TRL* where an average of 95 % of mobility data have an *AC* ¡ 0.5.

In the Geolife Dataset, the difference between EDEN's policies is more prominent. The data loss is reduced from 41 % in EDEN-opt to around 10 % in EDEN-pes. More than 79 %, 58 % and 48 % of the data generated by EDEN-opt, EDEN-balanced and EDEN-pes, respectively, have an *AC* ¿ 0.75. However, *PROM* with 50m, 100m and 200m cause a data loss of 10 %, 16 % and 25 %, respectively. The remaining data (i.e., 85 %, 76 % and 62 %) have an *AC* ¿ 0.75 of the original data. In this dataset, we observe that *PROM* has a better *AC* than EDEN . The reason for this is that EDEN has better results in terms of privacy in this dataset. Finally, *Geoi* and *TRL* with their different configurations generate on average 89 % and near 98 % of mobility data with an *AC* ¡ 0.75. Only 26 % of the generated data by *Geoi-0.01* has an *AC* ¿ 0.75.

In the MDC Dataset, Unlike *PROM* with its different configurations which causes 32 %, 41 % and 50 % of data loss, EDEN's policies reduce these amounts to 7 %, 8 % and 10 % with EDEN-opt, EDEN-balanced and EDEN-pes, respectively. EDEN's heuristics outperform all other LPPMs in terms of *AC*. It protects an average of 67 % of all data with *AC* ¿ 0.75 in comparison to an average of 0 %, 14 % and 39 % of the data obfuscated by *TRL*, *Geoi* and *PROM*, respectively.

5.5.2 Macro-benchmark on Air Pollution Measurements

We study the impact of LPPMs on quality of air pollution data using a dataset [6]. In addition to mobility data, two application-specific measurements are collected : the concentration of NO_2 and CO. This dataset involves 13 metropolitan bikes which have been equipped with pollution monitoring sensors for a duration of 112 days.

We compute the average CO measures over the duration when the NO_2 value is above $40\mu g/m^3$ which is the toxicity threshold as defined by the WHO[5].

In Figure 9, we show the average CO measurement results when using the raw data (NOBF baseline), and after the application of EDEN and its LPPM competitors. We can observe that *Geoi* and *TRL* spread the measurements and, as the noise or the range increases, they create additional hotspots, i.e., areas where the CO is high. The application of *PROM-50m* creates around 6 new hotspots where the level of CO is now above the toxicity threshold. Thus, although the good privacy vs. utility tradeoff provided by *PROM-50m* on synthetic mobility data (shown in Section 5.6), including the CO measurements yields to a poor performance. The use of *PROM-100m* can be harmful to public health because it eliminates existing high pollution hotspots from the original data. However, the application of EDEN is able to closely mirror the CO measurements of NOBF; only one hotspot is missed.

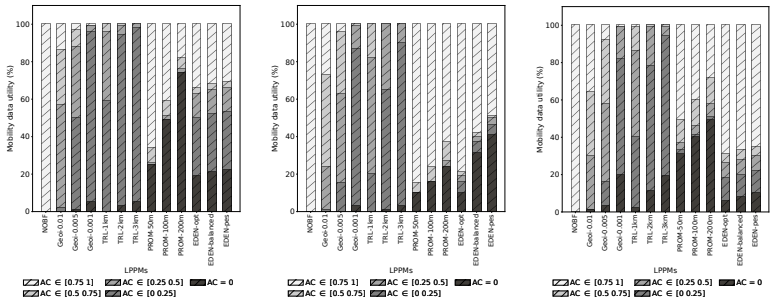

(a) PrivaMov (b) Geolife (c) MDC

Fig. 8: Impact of EDEN vs. LPPM competitors on the utility of the mobility data using AC metric.

5.5.3 Macro-benchmark on the number of cab drivers in San Francisco city

We study the impact of EDEN and LPPM competitors on the mobility of 50 cab drivers from the Cabspotting dataset [49]. The objective, for example is to find bottleneck locations where cab drivers go through in the city of San Francisco. For that purpose, we use range queries metric, previously defined in Section 5.3.1. It counts how many unique users cross an area during a time window. We choose time windows ranging from 2 hours to 8 hours and circle areas whose radius range from 500m to 5,000m. We report about the average query distortion in Table 2, which is the average distortion over 1,000 randomly generated queries. The results show that EDEN provides the smallest average distortion with 0.55% in comparison to, respectively,

[5]https://www.who.int/fr/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health

PROM which can reach 1.39%, TRL which can reach 16.45% and Geoi which can reach 19.88%.

Tab. 2: Average query distortion of EDEN and its competitors.

LPPM	**Average Query Distortion**
EDEN	0.55%
PROM-50	1.1%
PROM-100	1.28%
PROM-200	1.39%
TRL-1	6.96%
TRL-2	12.13%
TRL-3	16.45%
Geoi-01	1.07%
Geoi-005	2.73%
Geoi-001	19.88%

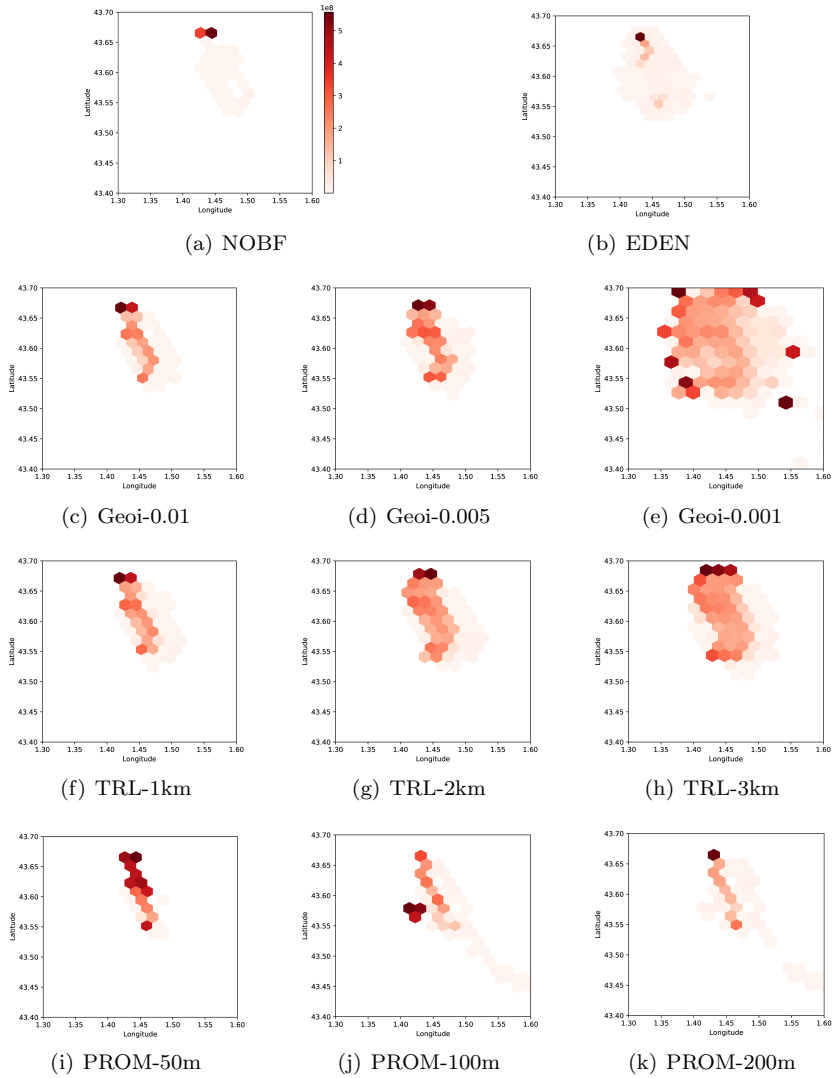

(a) NOBF (b) EDEN

(c) Geoi-0.01 (d) Geoi-0.005 (e) Geoi-0.001

(f) TRL-1km (g) TRL-2km (h) TRL-3km

(i) PROM-50m (j) PROM-100m (k) PROM-200m

Fig. 9: Macro-benchmark on air pollution dataset (CO gas).

5.6 Privacy vs. Utility Tradeoff

In Figure 10, we evaluate the privacy vs. utility tradeoff of EDEN compared to state-of-the-art LPPMs. To this end, we use a scatter plot where a point corresponds to an LPPM configuration defined by two coordinates x and y; x represents an aggregate value of the AC utility metric computed as in Section 2.3; y represents the percentage of protected data (i.e., data that is not re-identified by the MCS-side attacker). For instance, if we consider that an LPPM reaches a good privacy vs. utility tradeoff if it belongs to the top right gray rectangle ***"R"***, i.e., a rectangle where only LPPMs that

have a protection rate greater than 80 % with an AC greater than 40 % are considered. We represented the *Privacy Oracle* in order to show the maximum utility that can be attained by a specific dataset when all mobility traces are protected. In the same vein, *NOBF* represents the maximum privacy that can be attained by a specific dataset when no mobility trace is obfuscated.

In the Privamov Dataset, Figure 10(a), EDEN 's heuristics (illustrated by purple stars) provide the best privacy vs. utility tradeoff. All three EDEN heuristics belong to ***"R"*** whereas only *PROM-100* (illustrated by a green square) belongs to the low border of ***"R"***. In the Geolife Dataset, Figure 10(b), only EDEN-pes and EDEN-balanced belong to ***"R"***. However, the rest of the LPPMs have a privacy value concentrated around 62 % and an utility varying from 26 % to 88 %. Finally, in the MDC Dataset, Figure 10(c), EDEN-pes and *PROM* with its different configurations are inside ***"R"***. However, all EDEN's heuristics achieve a better privacy vs. utility tradeoff with a privacy close to *PROM* but an utility close to *Privacy Oracle*. To conclude, in the three considered datasets, EDEN achieves a better privacy vs. utility tradeoff than any other individual LPPMs with at least one of EDEN's heuristics belonging to the target rectangle ***"R"***.

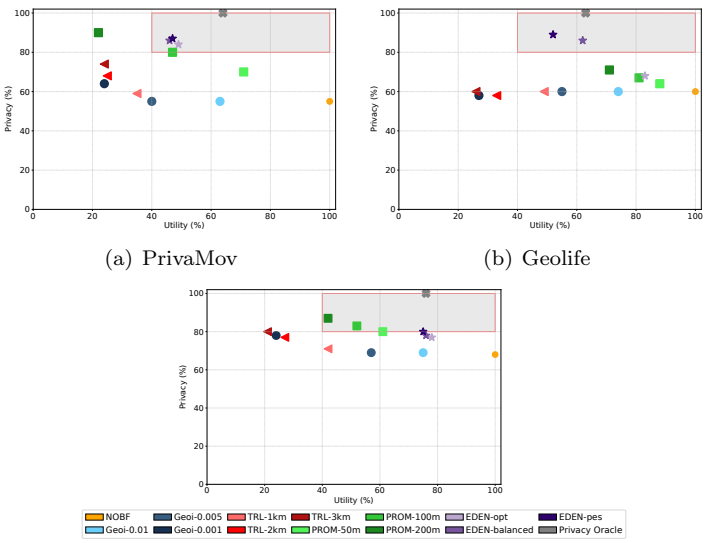

(a) PrivaMov (b) Geolife

(c) MDC

Fig. 10: Privacy vs. Utility tradeoff.

5.7 Fine-Grained Analysis of EDEN

In this section, we perform a fine-grained analysis of EDEN and its heuristics. Unlike the other LPPMs which are applied individually on the whole mobility data, EDEN protects users' mobility traces in a fine-grained way by choosing the most appropriate LPPM for each trace. In Figure 11, we represent EDEN 's heuristics with multi-color bars where the hashed part is the re-identified proportion of the mobility data and the plain part is the proportion of the data protected against the *Mv-Attack*. The color black represents the mobility data that FURIA is always able to re-identify regardless

the chosen LPPM. This data is deleted and considered as data loss from the MCS' point of view. Each LPPM is represented with a different color and its intensity expresses the strength of the LPPM's configuration: darker colors are used for stronger LPPM configurations. From this figure, we observe that the dominant color is orange which refers to NOBF. Indeed, when FURIA is not able to identify a raw mobility trace, EDEN prioritizes this choice because it provides the maximum utility.

By focusing on the Geolife Dataset, we observe that the data loss of *EDEN-pes* is about 31% (i.e., mobility traces that are always re-identified by FURIA). On the other extreme, *EDEN-opt* sends this portion of data without any protection or with a default LPPM. Roughly, 10 % of mobility data is protected with the chosen LPPM while 20 % of them are re-identified. However, we highlight that *EDEN-opt* ensures zero data loss. Finally, the balanced solution shows a tradeoff between data suppression and data publishing with a gain of +7 % of protection and a raise of +3 % of re-identification while the data loss is reduced to 21 %.

Concerning the MDC Dataset, the effect of EDEN's policies is not visible in this dataset because the amount of data that FURIA is always able to re-identify is negligible (4 %). In the Privamov Dataset, the colors distribution is more balanced: around 20 %, 11 %, 8 %, 3 % of data are naturally protected, or protected with *PROM*, *TRL* and *Geoi*, respectively. However the effect of EDEN's policies is as the same as in the MDC dataset, with around 3 % of re-identified mobility data. A similar evaluation is done where the choice of NOBF is excluded and the default LPPM used to send data in case FURIA is always able to re-identify the mobility trace is *PROM-50m*. Results are depicted in Figure 12. In PrivaMov and MDC datasets, LPPMs with low configurations (*TRL-1km, Geoi-0,01* and *PROM-50m*) replaces NOBF choice over EDEN's variants and the choice of *PROM-50m* as a default LPPM in EDEN-opt and EDEN-balanced does not impact the re-identification rate. However, in Geolife dataset, the choice of *PROM-50m* to protect the 41 % of deleted data in EDEN-pes increases the re-identification rate with +20 % and the protection rate with +21 %. EDEN-balanced improves these values with only +5 % of the former and +15 % of the latter.

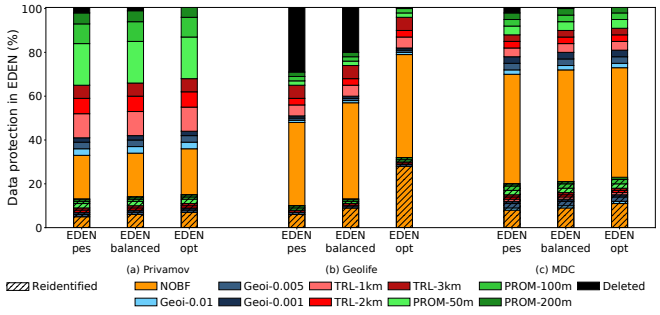

Fig. 11: Fine-grained EDEN with various heuristics.

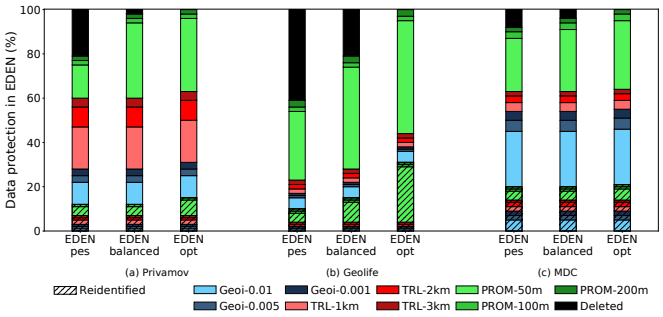

Fig. 12: Fine-grained EDEN with various heuristics without the NOBF choice.

5.8 EDEN Performance Overhead

EDEN is a user-side protection approach that operates directly on edge devices. In this section, we measure the run-time induced by EDEN on different sizes of mobility data ranging from 1 mobility record, i.e., equivalent to real-time crowd sensing applications to longer mobility traces (up to 1,600 records). Table 3 provides statistics about the considered datasets. The run-time overhead of EDEN mainly includes the run-time for (1) the protection process where a set of LPPMs with different configurations are executed (i.e., *Geoi, TRL, PROM*) and (2) the computation of AC metric to choose the best LPPM. In addition, the run-time for data processing (e.g., data formatting to feature vectors) and for the FURIA risk test are insignificant (nanosecond range). Results in Figure 13(a) with a logarithmic scale on the y-axis (for readability of small values), show that the longer a mobility trace, the longer it takes to protect it. Precisely, a mobility trace length ranging from 200 to 1,600 of mobility records can take from 3 to 9 seconds to protect it. This takes 15 milliseconds to protect a single mobility record which is acceptable in the context of real-time mobile crowd sensing applications. Moreover, we measure the execution time of the training phase in FURIA. Figure 13(b) presents the average training time per user in each learning round over all the datasets. Over all the datasets, we record a training time between 2 and 4 seconds per user. Specifically, in Geolife dataset, the training time is slightly higher than PrivaMov and MDC, this is due to a higher number of participant in comparison to Privamov (i.e., 19 users vs. 13 users) and denser mobility traces in comparison to MDC (i.e., 323 records vs. 46 records in average). Our experiment is conducted on a desktop machine (see Section 5.1) and it is still practical while using a smartphone. For instance, the average time of applying an LPPM on an edge device is in order of milliseconds [10]. Also, the authors in [17] evaluate the computation time of a learning round on different smartphones. The latter is in order of seconds.

Tab. 3: Mobility dataset statistics.

Dataset	Average records per user	Standard deviation	Minimum records	Maximum records	Average #Users per round
Geolife	323	288	1	1,800	19
PrivaMov	117	60	1	180	13
MDC	46	43	1	412	95

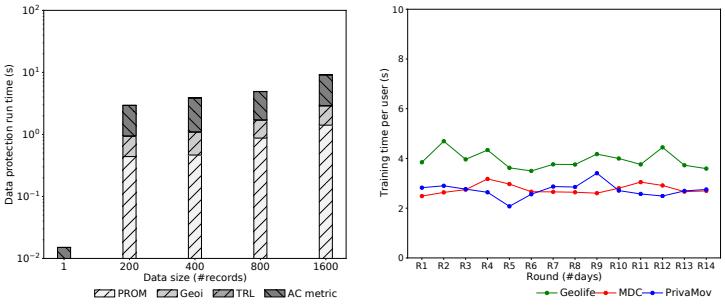

(a) EDEN protection run time (b) FURIA's model training time

Fig. 13: EDEN run-time overhead.

6 Conclusion

In this paper we presented EDEN a user-side approach for protecting mobility data by choosing the best among a set of LPPMs and without relying on a trusted proxy. EDEN relies on federated learning to train a decentralized user re-identification risk assessment model called *FURIA*. This model is then used on the user's device in order to locally compare LPPMs for each mobility trace to protect and select the one which is resilient against *FURIA*. EDEN also relies on an utility metric to ensure a good quality of the resulting data.

We evaluated EDEN by performing a set of experiments on real-world mobility traces. Results show that EDEN outperforms individual LPPMs both in terms of privacy measured in terms of the resilience against a strong attack combining state-of-the-art re-identification attacks and in terms of data utility measured using AC metric. In addition, EDEN was also evaluated on a crowd sensed air pollution dataset [6]. The results show that EDEN better preserves the distribution of gaseous pollutant in comparison with the other individual LPPMs. Our solution is promising as the re-identification risk assessment actively and incrementally learns new discriminative mobility patterns over time with incoming users updates without accessing raw data. This model can be improved by investigating new features such as the regularity of users mobility, speed and direction of travel, social links or the semantic of visited places, which we shall investigate in our future work. In conclusion, this work opens interesting perspectives where EDEN could be combined with other state-of-the-art orthogonal techniques in order to face untrusted users that aim to introduce a backdoor in the model [4, 58, 7], or to counter model information leakage attacks [9, 26, 43]. This work benefited from the support of the French National Research Agency (ANR), through the SIBIL-Lab project (ANR-17-LCV2-0014), and the PRIMaTE project (ANR-17-CE25-0017).

References

[1] Osman Abul, Francesco Bonchi, and Mirco Nanni. Never walk alone: Uncertainty for anonymity in moving objects databases. In *2008 IEEE 24th international conference on data engineering*, pages 376–385. Ieee, 2008.

[2] Osman Abul, Francesco Bonchi, and Mirco Nanni. Anonymization of moving objects databases by clustering and perturbation. *Information systems*, 35(8):884–910, 2010.

[3] Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 901–914, 2013.

[4] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2938–2948. PMLR, 2020.

[5] Noam Bardin. *https://www.waze.com*, 2008.

[6] Christophe Bertero. *Perception of the urban environment with the help of a fleet of sensors on bicycles. Application to air pollution*. PhD thesis, Toulouse University, Toulouse University, France, 2020.

[7] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pages 634–643. PMLR, 2019.

[8] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečnỳ, Stefano Mazzocchi, H Brendan McMahan, et al. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046*, 2019.

[9] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191, 2017.

[10] Ioannis Boutsis and Vana Kalogeraki. Location privacy for crowdsourcing applications. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 694–705, 2016.

[11] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

[12] Azby Brown, Pieter Franken, Sean Bonner, Nick Dolezal, and Joe Moross. Safecast: successful citizen-science for radiation measurement and communication after fukushima. *Journal of Radiological Protection*, 36(2):S82, 2016.

[13] Andrea Capponi, Claudio Fiandrino, Burak Kantarci, Luca Foschini, Dzmitry Kliazovich, and Pascal Bouvry. A survey on mobile crowdsensing systems: Challenges, solutions, and opportunities. *IEEE Commun. Surv. Tutorials*, 21(3):2419–2465, 2019.

[14] Sophie Cerf, Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, Robert Birke, Sara Bouchenak, Lydia Y Chen, Nicolas Marchand, and Bogdan Robu. Pulp: achieving privacy and utility trade-off in user mobility data. In *2017 IEEE 36th Symposium on Reliable Distributed Systems (SRDS)*, pages 164–173. IEEE, 2017.

[15] Sung-Hyuk Cha. Comprehensive survey on distance/similarity measures between probability density functions. *City*, 1(2):1, 2007.

[16] Mingqing Chen, Rajiv Mathews, Tom Ouyang, and Françoise Beaufays. Federated learning of out-of-vocabulary words. *arXiv preprint arXiv:1903.10635*, 2019.

[17] Georgios Damaskinos, Rachid Guerraoui, Anne-Marie Kermarrec, Vlad Nitu, Rhicheek Patra, and François Taïani. Fleet: Online federated learning via staleness awareness and performance prediction. In Dilma Da Silva and Rüdiger Kapitza, editors, *Middleware '20: 21st International Middleware Conference, Delft, The Netherlands, December 7-11, 2020*, pages 163–177. ACM, 2020.

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

[19] Srinivas Devarakonda, Parveen Sevusu, Hongzhang Liu, Ruilin Liu, Liviu Iftode, and Badri Nath. Real-time air quality monitoring through mobile sensing in metropolitan areas. In *Proceedings of the 2nd ACM SIGKDD international workshop on urban computing*, pages 1–8, 2013.

[20] Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.

[21] Lorenzo Franceschi-Bicchierai. Redditor cracks anonymous data trove to pinpoint muslim cab drivers. *Online at: http://mashable. com/2015/01/28/redditor-muslim-cab-drivers*, 2015.

[22] Clement Fung, Chris J. M. Yoon, and Ivan Beschastnikh. Mitigating sybils in federated learning poisoning. *CoRR*, abs/1808.04866, 2018.

[23] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. De-anonymization attack on geolocated data. *Journal of Computer and System Sciences*, 80(8):1597–1614, 2014.

[24] Raghu K Ganti, Fan Ye, and Hui Lei. Mobile crowdsensing: current state and future challenges. *IEEE communications Magazine*, 49(11):32–39, 2011.

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

[26] Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017.

[27] Philippe Golle and Kurt Partridge. On the anonymity of home/work location pairs. In *International Conference on Pervasive Computing*, pages 390–397. Springer, 2009.

[28] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018.

[29] Yan Huang, Zhipeng Cai, and Anu G Bourgeois. Search locations safely and accurately: A location privacy protection algorithm with accurate service. *Journal of Network and Computer Applications*, 103:146–156, 2018.

[30] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.

[31] Besma Khalfoun, Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. 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*, pages 136–148. ACM, 2019.

[32] John Krumm. Inference attacks on location tracks. In *International Conference on Pervasive Computing*, pages 127–143. Springer, 2007.

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

[34] Juha K Laurila, Daniel Gatica-Perez, Imad Aad, Olivier Bornet, Trinh-Minh-Tri Do, Olivier Dousse, Julien Eberle, Markus Miettinen, et al. The mobile data challenge: Big data for mobile computing research. Technical report, 2012.

[35] Huaxin Li, Haojin Zhu, Suguo Du, Xiaohui Liang, and Xuemin Shen. Privacy leakage of location sharing in mobile social networks: Attacks and defense. *IEEE Transactions on Dependable and Secure Computing*, 15(4):646–660, 2016.

[36] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubramaniam. l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):3–es, 2007.

[37] Nicolas Maisonneuve, Matthias Stevens, Maria E Niessen, and Luc Steels. Noisetube: Measuring and mapping noise pollution with mobile phones. In *Information technologies in environmental engineering*, pages 215–228. Springer, 2009.

[38] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. Hmc: Robust privacy protection of mobility data against multiple re-identification attacks. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–25, 2018.

[39] Mohamed Maouche, Sonia Ben Mokhtar, and Sara Bouchenak. 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*, pages 48–57, 2017.

[40] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.

[41] Lakhdar Meftah, Romain Rouvoy, and Isabelle Chrisment. Empowering mobile crowdsourcing apps with user privacy control. *Journal of Parallel and Distributed Computing*, 147:1–15, 2020.

[42] Scott Menard. *Applied logistic regression analysis*, volume 106. Sage, 2002.

[43] Fan Mo and Hamed Haddadi. Efficient and private federated learning using tee. In *EuroSys*, 2019.

[44] Mohamed F Mokbel, Chi-Yin Chow, and Walid G Aref. The new casper: Query processing for location services without compromising privacy. In *Proceedings of the 32nd international conference on Very large data bases*, pages 763–774, 2006.

[45] Sonia Ben Mokhtar, Antoine Boutet, Louafi Bouzouina, Patrick Bonnel, Olivier Brette, Lionel Brunie, Mathieu Cunche, Stephane D'Alu, Vincent Primault, Patrice Raveneau, et al. Priva'mov: Analysing human mobility through multi-sensor datasets. In *NetMob 2017*, 2017.

[46] Farid M Naini, Jayakrishnan Unnikrishnan, Patrick Thiran, and Martin Vetterli. Where you are is who you are: User identification by matching statistics. *IEEE Transactions on Information Forensics and Security*, 11(2):358–372, 2015.

[47] Sankar K. Pal and Sushmita Mitra. Multilayer perceptron, fuzzy sets, and classification. *IEEE Trans. Neural Networks*, 3(5):683–697, 1992.

[48] Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar Erlingsson. Scalable private learning with pate. *arXiv preprint arXiv:1802.08908*, 2018.

[49] Michal Piorkowski, Natasa Sarafijanovic-Djukic, and Matthias Grossglauser. Crawdad data set epfl/mobility (v. 2009-02-24), 2009.

[50] Layla Pournajaf, Daniel A Garcia-Ulloa, Li Xiong, and Vaidy Sunderam. Participant privacy in mobile crowd sensing task management: A survey of methods and challenges. *ACM Sigmod Record*, 44(4):23–34, 2016.

[51] Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. Adaptive location privacy with alp. In *2016 IEEE 35th Symposium on Reliable Distributed Systems (SRDS)*, pages 269–278. IEEE, 2016.

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

[53] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Differentially private location privacy in practice. *arXiv preprint arXiv:1410.7744*, 2014.

[54] Vincent Primault, Sonia Ben Mokhtar, Cédric Lauradoux, and Lionel Brunie. Time distortion anonymization for the publication of mobility data with high utility. In *2015 IEEE Trustcom/BigDataSE/ISPA*, volume 1, pages 539–546. IEEE, 2015.

[55] Nick Mathewson Roger Dingledine. *The TOR Project: https://www.torproject.org*, 2006.

[56] Marco Romanelli, Catuscia Palamidessi, and Konstantinos Chatzikokolakis. Generating optimal privacy-protection mechanisms via machine learning. *arXiv preprint arXiv:1904.01059*, 2019.

[57] Pierangela Samarati. Protecting respondents identities in microdata release. *IEEE transactions on Knowledge and Data Engineering*, 13(6):1010–1027, 2001.

[58] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019.

[59] Stuart A. Thompson and Charlie Warzel. *How to Track President Trump*, 2019.

[60] Jiangtao Wang, Yasha Wang, Daqing Zhang, and Sumi Helal. Energy saving techniques in mobile crowd sensing: Current state and future opportunities. *IEEE Communications Magazine*, 56(5):164–169, 2018.

[61] Marius Wernke, Pavel Skvortsov, Frank Dürr, and Kurt Rothermel. A classification of location privacy attacks and approaches. *Personal and ubiquitous computing*, 18(1):163–175, 2014.

[62] Mingqiang Xue, Panos Kalnis, and Hung Keng Pung. Location diversity: Enhanced privacy protection in location based services. In *International Symposium on Location-and Context-Awareness*, pages 70–87. Springer, 2009.

[63] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. Applied federated learning: Improving google keyboard query suggestions. *arXiv preprint arXiv:1812.02903*, 2018.

[64] Yu Zheng, Xing Xie, Wei-Ying Ma, et al. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data Eng. Bull.*, 33(2):32–39, 2010.

[65] Changqing Zhou, Dan Frankowski, Pamela Ludford, Shashi Shekhar, and Loren Terveen. Discovering personal gazetteers: an interactive clustering approach. In *Proceedings of the 12th annual ACM international workshop on Geographic information systems*, pages 266–273, 2004.

A Appendix

Algorithm 1 EDEN algorithm.

INPUT: T_U - mobility trace, $\mathbb{L}$ - a set of LPPMs, AF_i - "i"th FURIA model, AC - utility metric,
policy - {*EDEN-opt,EDEN-pes, EDEN-balanced*}, H - past user data, δ - deviation threshold.
OUTPUT: T' - a protected mobility trace.

function $\mathcal{EDEN}(T_U$,$\mathbb{L}$,AF_i,AC, *policy*, H,$\delta)$
  $Candidates \leftarrow \emptyset$
  for $\mathcal{L}$**in** $\mathbb{L}$ **do** ▷ Apply LPPMs
    $T' \leftarrow \mathcal{L}(T_U)$
    $V' \leftarrow FormatData(T')$ ▷ Transform T' to a Feature vector V'
    if $AF_i(V') \neq U$ **then** $Candidates \leftarrow Candidates \cup \{T'\}$ ▷ Re-identification Risk assessment
  end for
  if $Candidates \neq \emptyset$ **then**
    return $\{\underset{T' \in Candidates}{\arg\max} (AC(T_U, T'))[0]\}$
  else ▷ EDEN Policies
    if *policy* ="EDEN-pes" **then return** $\emptyset$
    if *policy* ="EDEN-opt" **then return** T_U.
    if *policy* ="EDEN-balanced" **then**
      $P \leftarrow Heatmap(T_U)$
      $Q \leftarrow Heatmap(H)$
      $deviation \leftarrow d_{Topsoe}(P, Q)$
      if $deviation > \delta$ **then**
        return T_U
      else
        return $\emptyset$
      end if
    end if
  end if
end function