Ricardo Mendes*, Mariana Cunha, and João P. Vilela

Impact of Frequency of Location Reports on the Privacy Level of Geo-indistinguishability

Abstract: Location privacy has became an emerging topic due to the pervasiveness of Location-Based Services (LBSs). When sharing location, a certain degree of privacy can be achieved through the use of Location Privacy-Preserving Mechanisms (LPPMs), in where an obfuscated version of the exact user location is reported instead. However, even obfuscated location reports disclose information which poses a risk to privacy. Based on the formal notion of differential privacy, Geo-indistinguishability has been proposed to design LPPMs that limit the amount of information that is disclosed to a potential adversary observing the reports. While promising, this notion considers reports to be independent from each other, thus discarding the potential threat that arises from exploring the correlation between reports. This assumption might hold for the sporadic release of data, however, there is still no formal nor quantitative boundary between sporadic and continuous reports and thus we argue that the consideration of independence is valid depending on the frequency of reports made by the user. This work intends to fill this research gap through a quantitative evaluation of the impact on the privacy level of Geo-indistinguishability under different frequency of reports. Towards this end, state-of-the-art localization attacks and a tracking attack are implemented against a Geo-indistinguishable LPPM under several values of privacy budget and the privacy level is measured along different frequencies of updates using real mobility data.

Keywords: Location-Based Services, Location Privacy, Location Privacy-Preserving Mechanisms, Geo-indistinguishability

DOI 10.2478/popets-2020-0032
Received 2019-08-31; revised 2019-12-15; accepted 2019-12-16.

***Corresponding Author: Ricardo Mendes:** CISUC, Department of Informatics Engineering, University of Coimbra, E-mail: rscmendes@dei.uc.pt
Mariana Cunha: CISUC, Department of Informatics Engineering, University of Coimbra, E-mail: mccunha@dei.uc.pt
João P. Vilela: CISUC, Department of Informatics Engineering, University of Coimbra, E-mail: jpvilela@dei.uc.pt

1 Introduction

Mobile devices and ubiquitous connectivity fostered services that take into consideration users' contextual information. One emergent category of these services is the Location-Based Services (LBSs), in which users share their location to obtain geographically and temporally related information (e.g. finding the nearest open restaurant). While beneficial to the user, sharing location data poses a threat to privacy that goes beyond physical safety. In fact, visited locations can reveal users' identity, habits, addictions, health conditions and even social connections [1, 2].

Untrustworthy LBS providers, that may share or publish the data, passive eavesdroppers and security breaches can cause disclosure of location data thus putting at risk the privacy of its users. Preserving privacy against this range of attack vectors requires Location Privacy-Preserving Mechanisms (LPPMs) at collection time, i.e. mechanisms that run in-device in an online scenario [3]. LPPMs report an obfuscated version of the exact user location as to preserve a certain level of privacy at the expense of a degraded quality of service.

Geo-indistinguishability [4], a recently proposed formal notion based on differential privacy [5] has seen increasing research interest due to its simplicity of implementation, efficiency and effectiveness [6–8]. Geo-indistinguishability guarantees that any two points within a given radius around the user are statistically indistinguishable independently of an adversary's background information. Specifically, the reported (obfuscated) point is generated with (almost) the same probability for any point within this circle, consequently concealing the exact location of the user.

Depending on the LBS, location data can be reported either continuously or rather sporadically [9, 10]. This frequency of reports directly impacts the temporal correlation between subsequent reports which in turn can be used by an adversary to track users over time and even predict future locations [1, 6, 11]. While geo-indistinguishability bounds the amount of disclosure, it considers reports to be independent between each other. In fact, in the context of sporadic release

of data this consideration has been assumed when designing LPPMs [9, 12]. However, there is no formal nor quantitative distinction between sporadic and continuous reports and thus, the distinction is often based on the type of LBS application [10]. In this work we argue that the consideration of independence depends on the frequency of updates, even in the context of sporadic reports. Therefore and to evaluate our premise, we quantitatively study the impact of the frequency of reports on the achieved privacy level through geo-indistinguishability. The contributions of this work are as follows.

- We evaluate the effect of the frequency of updates in the privacy level of the Planar Laplace [7], a geo-indistinguishable LPPM, using state-of-the-art localization attacks and a tracking attack on real datasets. The variation of the frequency of updates is made such that typical values for both continuous and sporadic are considered as well as values in between both ends. Results showed that the privacy level when considering localization attacks is roughly constant over the range of tested frequencies of updates, while the effectiveness of tracking attacks decays as the frequencies of updates lowers. These results suggest that the consideration of independence between reports can be effectively assumed in the sporadic scenario.
- We evaluate the effectiveness of several values of ϵ, the privacy budget, in the privacy level of the Planar Laplace against the state-of-the-art localization attacks. The choice of a privacy budget in differential privacy, and consequently based approaches such as geo-indistinguishability, is still an open problem as it strongly depends on the application [8]. In fact, it has been discussed that the definition of ϵ in geo-indistinguishability may be misleading in terms of the privacy level [13]. In contrast with [12], our results showed that the relation between the average quality loss and average adversary error is only linear after a non-negligible threshold. That is, there exists an upper bound on the value of the privacy budget necessary to guarantee relevant privacy protection, which in our setting was $\epsilon = 4$ km^{-1}.
- We assess the effects of the grid resolution in the effectiveness of the implemented localization attacks. These results show a linear correlation between the cell width and the average adversary error, and thus suggest that a powerful adversary (with infinite computational power) could potentially defeat obfuscation. However, increasing the obfuscation (by decreasing ϵ) decreases the slope of the linear correlation. Consequently, by increasing the obfuscation, a higher decrease in cell width, and consequently an increase in computational complexity, is required for the same reduction in the average adversary error.

A previous work [14] has shown that the correlation between subsequent reports can be explored by an adversary using simple regression models as estimators. From such results it was concluded that not only does the frequency of updates greatly impacts the temporal correlation but also that the estimation function affects the results significantly. However, the privacy level evaluation in that work was limited due to the use of simple regressions as attacks. This work greatly expands those results by providing a quantitative privacy evaluation with state-of-the-art attacks under both continuous and sporadic release of location data. While map-aware LPPMs have been proposed in the literature (e.g. [15]) and map knowledge has been used to reduce obfuscation areas (e.g. [16]), to the best of our knowledge, we are the first to consider road network map-matching as a tracking attack.

The remainder of this document is structured as follows. Section 2 presents an overview of the location privacy paradigm and details the implemented localization attacks and LPPM. Section 3 describes the empirical methodology whose results are displayed and discussed in Section 4. Section 5 presents the related work and this document concludes in Section 6.

2 Location Privacy Overview

This section provides an overview of the location privacy paradigm. Section 2.1 presents the notation used throughout this work and formalizes the problem to be tackled. The implemented geo-indistinguishable LPPM is described in Section 2.2 and Section 2.3 presents the localization attacks carried by a possible adversary.

2.1 Problem Definition

As in previous relevant works [7, 10, 12, 14, 17], we shall consider a user of an LBS which reports his location to the LBS provider to obtain information. We consider as adversary any entity with access to the location reports attempting to infer private information [2, 16], including the LBS provider or any passive eavesdropper. Furthermore, the adversary can have arbitrary background information (prior) and computational power. In order

Table 1. Summary of notation

Symbol	Description
x^r	r^{th} location from $\mathcal{X}$, with $r \in \{1, \ldots, \lvert\mathcal{X}\rvert\}$
x_i	Exact user location at timestamp i.
z_i	Obfuscated location at timestamp i.
$\hat{x}_i$	Adversary's estimated location at timestamp i.
t_i	Time at timestamp i.
$\mathbf{x}$, $\mathbf{z}$, $\hat{\mathbf{x}}$	Vector of all real, obfuscated or estimated locations, respectively.
$\mathbf{x}_i$, $\mathbf{z}_i$, $\hat{\mathbf{x}}_i$	Vector of real, obfuscated or estimated locations up to timestamp i.
$\mathcal{X}$, $\mathcal{Z}$, $\hat{\mathcal{X}}$	Set of all possible real/obfuscated/estimated locations.
Δ_t	Minimum interval between consecutive reports.
f, $p(z_i \lvert x_i)$	Location Privacy-Preserving Mechanism (LPPM).
ϵ	Geo-indistinguishability privacy parameter.
h, $p(\hat{x}_i \lvert z_i)$	Adversary's attack.
$P_{AE}(f, h, \mathbf{x}, \mathbf{z})$	Mean adversary error of $\hat{\mathbf{x}}$ given $\mathbf{z}$ and h.
$Q(f, \mathbf{x}, \mathbf{z})$	Mean quality loss given the LPPM f and locations $\mathbf{x}$.
$d(\cdot)$	Euclidean distance metric.
$g(\cdot)$	Great-circle distance.
o_i	Noisy GPS reading at timestamp i.
$s_{i,k}$	k^{th} candidate location for o_i at timestamp i.
$p(o_i \lvert s_{i,k})$	Map-matching emission probability.
$p(s_{i,k} \lvert s_{i-1,j})$	Map-matching transition probability
σ	Standard deviation of the (GPS) measurement error.
λ_y	Parameter for the exponential of the measure of circuitousness.
λ_z	Parameter for the exponential of the measure of temporal plausibility.

to protect his privacy, the user uses an LPPM to report an obfuscated version of his exact location, consequently trading the quality of the LBS response for privacy.

Formally, let $x_i \in \mathcal{X}$ denote the exact user's location at the report with timestamp $i \in \{1, 2, \ldots, T\}$ and $z_i \in \mathcal{Z}$ the reported obfuscated location at the same i computed using the LPPM f. For convenience, we use t_i to express the real time of timestamp i. The adversary has access to z_i and it is assumed to know f and possibly have some a priori knowledge and thus computes $\hat{x}_i \in \hat{\mathcal{X}}$, an estimation of x_i at each timestamp i, using an attack h. We shall denote $\mathbf{x}_i$ and $\mathbf{z}_i$ the vectors of real and obfuscated locations up to timestamp i, respectively, that is, $\mathbf{x}_i = \{x_1, \ldots, x_i\}$ and $\mathbf{z}_i = \{z_1, \ldots, z_i\}$. We assume the exact location to be placed in a finite grid $\mathcal{X}$ but let $\mathcal{Z}$ and $\hat{\mathcal{X}}$ to be $\mathbb{R}^2$, i.e., the obfuscated report and the adversary estimation can be any point in the map. In the context of frequency of updates, we define Δ_t in seconds as the minimum interval between any two consecutive location reports. Formally, $\Delta_t = \text{argmin}_i\ (t_{i+1} - t_i)$.

Generically [12], online user-centric obfuscation mechanisms can be described as a probability distribution in the form of equation (1).

$$p(z_i | \mathbf{z}_{i-1}, \mathbf{x}_i) \tag{1}$$

Intuitively, an LPPM maps the real location $x_i \in \mathcal{X}$ with the knowledge of past locations $\mathbf{x}_{i-1}$ and past reports $\mathbf{z}_{i-1}$ to a new report $z_i \in \mathcal{Z}$. In the context of sporadic location privacy, existing LPPMs consider location reports to be independent, and consequently, each obfuscated report z_i is made only with respect to the exact position x_i at the same timestamp i. Therefore, equation (1) is reduced to the form:

$$p(z_i | \mathbf{z}_{i-1}, \mathbf{x}_i) = p(z_i | x_i) \tag{2}$$

LPPMs of this form are referred to as memoryless [12].

In the context of localization attacks, the primary privacy metric is the correctness of an adversary measured by the expected estimation error [18, 19] and modeled through a distance metric between the exact locations and the adversary's estimations. Given an LPPM f, an attack h and observations $\mathbf{z}$, the expected adversary estimation (AE) error is defined by the following equation:

$$P_{AE}(f, h, \mathbf{x}, \mathbf{z}) = E\{d(x_i, \hat{x}_i)\} \tag{3}$$

where the expected value is taken over x_i and $\hat{x}_i$ and $d(\cdot)$ is a distance metric which is typically the Euclidean distance [18].

From the user perspective, the LPPM f introduces a quality loss due to reporting the obfuscated location instead of the exact position [10, 19]. The average quality loss is therefore given by:

$$Q(f, \mathbf{x}, \mathbf{z}) = E\{d(x_i, z_i)\} \tag{4}$$

The objective of this work is to evaluate the impact of the interval between reports (Δ_t) in the achieved privacy level. We challenge the consideration of independence between reports taken by previous sporadic location privacy approaches through empirical experimentation. We focus on a particular geo-indistinguishable LPPM, the Planar Laplace [4, 7], and implement state-of-the-art localization attacks. We measure the adversary's correctness for the different attacks and under different values of Δ_t. The following sections detail this LPPM and the implemented attacks, respectively.

2.2 Planar Laplace

The Planar Laplace (PL) mechanism was the first proposed mechanism to achieve the notion of geo-indistinguishability which itself is a variant of differential privacy applied to location based systems [4]. While other mechanisms have been proposed, including optimal approaches w.r.t. utility [7, 17, 20], the PL is efficient and therefore practical to be used in an online setting such as the one described above. Remapping techniques have been proposed for this LPPM to increase the utility of the queries without degrading the privacy level [7]. In fact, the PL mechanism with optimal remapping is considered the state-of-the-art of geo-indistinguishability in sporadic location privacy [12].

The optimal remapping techniques only uses the current obfuscated LPPM output (z_i) and the mobility profile of the user [7]. Consequently, an adversary with knowledge of the mobility profile, which is typically assumed (c.f. Section 2.3.1) and with access to the remapped obfuscated report is able to reverse this mapping. To reduce computational complexity, we assume, without loss of generality, the adversary to have the data after reversing the remapping, that is, the obfuscated reports as computed by the LPPM.

Prior to explaining the PL, one must first introduce geo-indistinguishability and differential privacy. In the context of statistical databases, differential privacy guarantees that the presence or absence of a single individual in a database does not considerably impact the disclosure of information [5]. In fact, information disclosure in differential privacy is quantitatively measured as the difference between the prior knowledge and the posterior (after the data is added/removed) knowledge which is bounded by the privacy budget, which is a small pre-defined constant. Geo-indistinguishability on the other hand guarantees that the disclosed location is indistinguishable from any other point within a variable radius, thus concealing the exact location, while allowing for enough information release [4]. Formally and following [4], an LPPM f satisfies ϵ-geo-indistinguishability iff:

$$d_{\mathcal{P}}\left(f(x), f(x')\right) \leq \epsilon d_{\mathrm{x}}(x, x') \quad \forall x, x' \in \mathcal{X} \tag{5}$$

where ϵ is the pre-defined constant referred to as privacy budget, $d_{\mathrm{x}}(\cdot)$ is any distance function and $d_{\mathcal{P}}(\cdot)$ is the multiplicative distance between two distributions, defined as $d_{\mathcal{P}}(\sigma_1, \sigma_2) = \sup_{S \in \mathcal{S}} \left|\log \frac{\sigma_1(S)}{\sigma_2(S)}\right|$, where σ_1 and σ_2 are two distributions on some set S, with the convention that $\mathcal{L} = \left|\log \frac{\sigma_1(S)}{\sigma_2(S)}\right| = 0$ if $\sigma_1(S) = \sigma_2(S) = 0$ and $\mathcal{L} = \infty$ if one of the two is 0.

Intuitively, from equation (5), the probability of generating z using an ϵ-geo-indistinguishable LPPM from either x or x' is bounded by the distance between these two points factored by the privacy budget ϵ. Commonly [4, 7], ϵ is set to $\epsilon = l/r$, where r and l are a user specified radius and privacy level, such that for any x, x' s.t. $d_{\mathrm{x}}(x, x') \leq r$, $d_{\mathcal{P}}\left(f(x), f(x')\right) \leq l$. This enforces that closer x and x' locations will have similar probability functions, thus better concealing the true location while allowing higher dissimilarity for distant locations to preserve a certain degree of utility.

The PL mechanism consists in adding 2-dimensional Laplacian noise centered at the exact user location x and with pdf:

$$p(z|x) = D_x(z) = \frac{\epsilon^2}{2\pi} e^{-\epsilon d_{\mathrm{x}}(x,z)} \tag{6}$$

Obtaining z from x using equation (6) can be efficiently done by adding a randomly drawn vector expressed as a radius r and angle Θ. Θ is uniformly chosen from $[0, 2\pi)$ and r is computed by drawing p uniformly from $[0, 1)$ and feeding it to the inverse planar Laplacian cumulative distribution function defined as $C^{-1}(p) = -\frac{1}{\epsilon}\left(W_{-1}\left(\frac{p-1}{e}\right) + 1\right)$, where W_{-1} is the negative branch of the Lambert W function. Finally, $z = x + \langle r\cos\Theta, r\sin\Theta\rangle$.

2.3 Location Privacy Attacks

Location privacy attacks are diverse with respect to both the objective and the applied methods [21]. In this work we focused on the objective of locating the user at each timestamp. This objective is general in the sense that it allows for the reconstruction of the true mobility of the user and consequently, for posterior inference attacks, that is, attacks which produce additional knowledge from the geolocation data [2] (e.g. extraction of user's points-of-interest). Furthermore, since we consider varying the frequency of updates, this requires considering both localization and tracking attacks. Recall that tracking techniques consist in following a user over time and space, whereas localization techniques have as objective to localize the user at certain points in time [9].

For localization attacks, we focused on the state-of-the-art by considering the optimal attack given a mobility profile [18] and an heuristic which learns the mobility profile as locations are shared [12]. Section 2.3.1 and 2.3.2 detail these attacks, respectively.

In tracking attacks, one can consider regression analysis, Kalman filtering, particle filters and map-matching [1, 21]. In a previous work [14], regression analysis has been used to produce simple estimators (such as linear and polynomial) as a tracking attack. However, results showed that such solution generates a non-negligible amount of outliers due to time-gaps in reports, which occur due to failures in the GPS/communications. Kalman filters have been used effectively in navigation to reduce uncertainties arising from the noisy measurements. Particle filters can be used for the same purpose incurring in higher computational complexity. However, these two techniques are oblivious of the underlying map and consequently generate positions that are not physically possible (e.g. inside a building if the user is driving). A knowledgeable

adversary can make use of the map to reduce this kind of uncertainties and thus locate the user with higher precision [21]. This process is known as map-matching and it is typically used to locate vehicles on road-networks [22]. Therefore, in this work we have selected a state-of-the-art road network map-matching attack which is detailed in Section 2.3.3. It should be noted that the considered localization attacks can also be used for trajectories [23]. However, these attacks require the discretization of the space (and possibly time), which becomes computationally infeasible for finer resolutions.

Even though map-matching has been used as an attack, for instance, against area obfuscation [21], to the best of our knowledge, this is the first work to consider road-network map-matching as a tracking attack. We also note that this choice was further supported by the fact that hidden Markov chains, which are used in map-matching, have been shown effective in modelling the temporal correlations of location traces [11, 24]. A natural extension of this work is to consider other types of both localization and tracking attack, or even inference attacks, such as the extraction of sensitive semantic locations.

2.3.1 Optimal Localization Attack

As aforementioned, the adversary observes $\mathbf{z}$, knows the used LPPM f and has some priory knowledge in the form $p(\mathbf{x})$. Consequently, it computes $\hat{\mathbf{x}}$ by means of an attack h. We focus on the case that the adversary estimates x_i using only observed reports up to i, that is, $\mathbf{z}_i$. This case can be generalized to the estimation of x_i using z_k with $i \leq k$ [12], however this is rarely the case in tracking approaches. Following [12, 18], the optimal localization attack minimizes the estimation error defined by equation (3). Formally:

$$\hat{x}_i = \operatorname*{argmin}_{\hat{x}_i} \sum_{x_i \in \mathcal{X}} p(x_i|\mathbf{z}_i) \cdot d_P(x_i, \hat{x}_i) \tag{7}$$

where $p(x_i|\mathbf{z}_i)$ is the posterior probability of x_i given all reports up to i:

$$p(x_i|\mathbf{z}_i) = \frac{p(\mathbf{z}_i|x_i) \cdot p(x_i)}{p(\mathbf{z}_i)} = \frac{\prod_{l=1}^{i} p(z_l|\mathbf{z}_{l-1}, x_i) \cdot p(x_i)}{p(\mathbf{z}_i)} \tag{8}$$

Note that since z_l is conditionally independent of x_i for $l \neq i$ and since we are considering only memoryless LPPMs, we have:

$$\begin{cases} p(z_l|\mathbf{z}_{l-1}, x_i) = p(z_l|\mathbf{z}_{l-1}) & \text{if } l \neq i \\ p(z_l|\mathbf{z}_{l-1}, x_i) = p(z_i|\mathbf{z}_{i-1}, x_i) = p(z_i|x_i) & \text{if } l = i \end{cases}$$

Furthermore, since equation (7) is a minimization, we can ignore the denominator and thus reach the attackers objective function as:

$$\hat{x}_i = \operatorname*{argmin}_{\hat{x}_i} \sum_{x_i \in \mathcal{X}} p(z_i|x_i) \cdot p(x_i) \cdot d_P(x_i, \hat{x}_i) \tag{9}$$

The final consideration of an attacker is the characterization of $p(\mathbf{x})$. Traditionally [10, 18], $p(\mathbf{x})$ is described by a mobility profile π which is a probabilistic representation of the user mobility, where each user location is considered an i.i.d. sample of π. Formally, let $\pi(x)$ denote the probability that the user is at $x \in \mathcal{X}$ given the mobility profile π, then $p(\mathbf{x}) = \prod_i \pi(x_i)$. Therefore, and in practice [7], a realistic adversary would use a mobility profile built with training data, π^{train}. An omniscient adversary is sometimes considered as one who has access to the test data and thus, builds the mobility profile from this data, π^{test}. This latter adversarial consideration gives a lower bound for the expected privacy. We refer to the optimal attack using π^{train} as **optHW** and using π^{test} as **omniHW**.

Recently, Oya et al. [12] observed that building the mobility model a priori with the training data might fail to capture the true mobility of the users. The closer the model is to the real mobility, the better performant is the attack[1]. Consequently, the authors propose a new approach towards building mobility profiles which considers the true mobility to be unknown, and therefore learned based on real user behavior in an a posteriori fashion. An attack using this approach was proposed in [12] and results showed to have better performance than the optimal attack using the a priori model. The attack is denominated Profile-Estimation Based Attack (PEBA) and described in the following section.

2.3.2 Profile-Estimation Based Attack (PEBA)

The Profile-Estimation Based Attack (PEBA) [12] is based on the idea that the real mobility profile is unknown and consequently has to be learned/adapted after each query. Formally, let $p(\pi)$ be the probability of being assigned a profile $\pi \in \mathcal{F}_\pi$, then the real locations are i.i.d samples of the distribution given by π, such that:

$$p(\mathbf{x}) = \sum_{\pi \in \mathcal{F}_\pi} p(\pi) p(\mathbf{x}|\pi) = \sum_{\pi \in \mathcal{F}_\pi} p(\pi) \prod_i \pi(x_i) \tag{10}$$

1 Note that the mobility profile might not only be used by an adversary in the attack but also by the user in the LPPM [23].

This consideration creates a dependency between exact locations due to the fact that a previous location gives information on the unknown profile π which in turn affects the probability of the following locations. Therefore, the real locations and obfuscated locations will also be dependent as a location at x_j affects distribution of a location at x_i with $i > j$ which in turn affects z_i. Consequently, it becomes mathematically intractable to find the optimal attack considering equation (10) [12]. Thus, PEBA is a sub-optimal attack.

Following [12], PEBA is decomposed in two sequential steps: 1) estimation of the mobility profile using the observed obfuscated reports $\mathbf{z}_i$ up to the current timestamp, i. In the original proposal the Maximum Likelihood (ML) estimator is used and thus, this mobility profile is denoted by $\hat{\pi}_i^{ML}$; 2) estimate the real location $\hat{x}_i$ using $\mathbf{z}_i$ and assuming that x_i follows the estimated mobility profile $\hat{\pi}_i^{ML}$. We skip the foundational details of the method and focus on the implementation steps. The interested reader should refer to [12].

The procedure of the steps is as follows. From the training data an initial average mobility profile π^{avg} is built from all the users. Then, this initial profile is used to estimate $\hat{\pi}_i^{ML}$, through an iterative Expectation-Maximization method, following equation (11).

$$\pi^{r,t+1} = \frac{1}{i}\sum_{l=1}^{i} p(x_l^r|\mathbf{z}_l, \pi^t) =$$

$$= \frac{1}{i}\sum_{l=1}^{i} \frac{\pi^{r,t} \cdot f(z_l|\mathbf{z}_{l-1}, x_l^r)}{\sum_{k=1}^{|\mathcal{X}|} \pi^{k,t} \cdot f(z_l|\mathbf{z}_{l-1}, x_l^k)} \tag{11}$$

where t is an iteration counter and $\pi^r \equiv p(x = x^r)$ with $x^r \in \mathcal{X}$ and $r \in \{1, \ldots, |\mathcal{X}|\}$ denotes the probability mass function defined by π. Furthermore, $\pi^0 = \pi^{avg}$. This step is repeated while the change from π^t to π^{t+1} is significant. Then, a normalization of the profile is made following equation (12). This latter equation holds that for the initial queries, the initial mobility profile π^{avg} is dominant, and then fading out as the number of queries increase in favor of the ML estimator.

$$\hat{\pi}_i = \frac{1}{i^{0.5}} \cdot \pi^{avg} + \left(1 - \frac{1}{i^{0.5}}\right) \cdot \hat{\pi}_i^{ML} \tag{12}$$

The posterior is then computed as:

$$p(x_i|\mathbf{z}_i, \hat{\pi}_i) = p(\mathbf{z}_i|x_i, \hat{\pi}_i) \cdot \hat{\pi}_i(x_i)/p(\mathbf{z}_i) \tag{13}$$

$$= \prod_{l=1}^{i} p(z_l|\mathbf{z}_{l-1}, x_i, \hat{\pi}_i)\hat{\pi}_i(x_i)/p(\mathbf{z}_i) \tag{14}$$

And finally, using the posterior, the PEBA estimation of the exact location of the user is calculated as:

$$\hat{x}_i = \underset{\hat{x}_i}{\operatorname{argmin}} \sum_{x_i \in \mathcal{X}} p(x_i|\mathbf{z}_i, \hat{\pi}_i) \cdot d_P(x_i, \hat{x}_i) \tag{15}$$

2.3.3 Map-Matching

The previous sections described attacks against sporadic reports, referred to as localization attacks. This section focus on a tracking problem, known as map-matching. Map-matching (MM) is the process of continuously identifying the position of a vehicle on the road network given noisy location readings [22]. However, map-matching can also be used as an adversary tracking/locating a user as detailed in this section.

In the context of MM, it is typically considered high frequency of updates when reports are made up to every 1 minute. Any value above this interval is considered low frequency of updates, and commonly, low frequency MM techniques are evaluated up to a maximum of 5-6 minutes [25]. In the context of LBSs however, 5 to 6 minutes is still considered continuous reports. Nevertheless, using a MM technique allows to evaluate the impact of frequency in highly continuous updates and consequently, assess the privacy level under the full range of frequencies.

The criteria for selecting the implemented MM technique was the effectiveness over noise, as the Planar Laplace applies additive noise, and the effectiveness over low frequency of updates, which results in sparse data. A seminal work fulfilling these criteria is found in [26], where their method is evaluated over frequency of updates (referred to as sampling period) varying from 1 second to 600 seconds and over the addition random Gaussian noise to the GPS readings with multiple standard deviation values. A follow up on this work was made by Jagadeesh and Srikanthan [27], where locations were measured with cellular network positioning instead of the GPS. The measurement error from the former positioning system is higher by almost 2 orders of magnitude and therefore the MM technique was adapted to be more robust against noise.

Comparative results between [27] and the seminal work from [26] showed the former technique to be more robust to both low frequency of updates and noisy measurements. Consequently, we have implemented the MM technique from [27], which we describe next. We refer the reader to [27] for a more detailed explanation of the original method.

Let us denote $o_i \in \mathbb{R}^2$ as the location report (referred to as observation in [27]) at timestamp i. This report is not obfuscated but it is assumed to be noisy due to measurement imprecision. The road network is a direct graph $G = (V, E)$, where V is a set of nodes representing intersections and endpoints of road segments and E is the set of these segments. A path p between

nodes u and v is a sequence of edges $e_1, \ldots, e_n$ such that u is the tail of e_1 and v is the head of e_n. The objective of a MM algorithm is to find a path p that corresponds to a sequence of T locations given noisy observations $o_1, \ldots, o_T$. Towards this goal, an Hidden Markov Model (HMM) is used in [27].

At each noisy observation o_i, the HMM's hidden states at time step i correspond to potential locations on the road where the user can be. We denote the k^{th} potential location at time step i by $s_{i,k}$ and the hidden true state by $s_i^* = x_i$. Given that the location measurement error can be assumed effectively to follow a Gaussian distribution with zero mean [26, 27], the probability that the observation o_i was generated from state $s_{i,k}$, referred to as emission probability, is given by:

$$p(o_i|s_{i,k}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{g(o_i,s_{i,k})^2}{2\sigma^2}} \tag{16}$$

where σ is the standard deviation of the measurement error and $g(o_i, s_{i,k})$ is the great-circle distance, that is, the shortest distance along the surface of the earth, between the observation o_i and the state $s_{i,k}$. Note that from equation (16), it is clear that closer states to the observation will have a higher probability than farther states, as the denominator increases exponentially with the increase of the distance $g(\cdot)$.

The transition probability, that is, the probability that the vehicle moved from state $s_{i-1,j}$ to $s_{i,k}$ depends on both the circuitousness of the path and on the temporal plausibility, that is, if the travelled distance is plausible given the time interval between timestamps $(t_i - t_{i-1})$. To measure the circuitousness of the path, the authors of [27] defined the following equation:

$$y(s_{i-1,j}, s_{i,k}) = \frac{d(s_{i-1,j}, s_{i,k}) - g(s_{i-1,j}, s_{i,k})}{(t_i - t_{i-1})} \tag{17}$$

where $g(s_{i-1,j}, s_{i,k})$ is the great circle distance between the states and $d(s_{i-1,j}, s_{i,k})$ the driving distance, calculated using Dijkstra's shortest path algorithm [28]. For the temporal plausibility, the equation is given as:

$$z(s_{i-1,j}, s_{i,k}) = \frac{\max\big(f(s_{i-1,j}, s_{i,k}) - (t_i - t_{i-1}), 0\big)}{(t_i - t_{i-1})} \tag{18}$$

where $f(s_{i-1,j}, s_{i,k})$ is the free-flow travel time, in seconds, of the optimal path between the states $s_{i-1,j}$ and $s_{i,k}$. Finally, the transition probability comes in the form:

$$p(s_{i,k}|s_{i-1,j}) = \lambda_y e^{-\lambda_y y(s_{i-1,j}, s_{i,k})} \lambda_z e^{-\lambda_z z(s_{i-1,j}, s_{i,k})} \tag{19}$$

where λ_y and λ_z are empirically determined parameters from equations (17) and (18), respectively.

To compute the most likely path from the HMM, a Viterbi algorithm is used as follows:

$$\begin{aligned} V_{1,k} &= p(o_1|s_{1,k}) \\ V_{i,k} &= p(o_i|s_{i,k}) \max_j \big(V_{i-1,j} p(s_{i,k}|s_{i-1,j})\big) \end{aligned} \tag{20}$$

where $V_{i,k}$ is the joint probability of the most likely state sequence ending at state $s_{i,k}$ based on the observations $o_1, \ldots, o_i$. The index j that maximizes $V_{i,k}$ is stored for each potential location k as it points to the predecessor state $s_{i-1,j}$ that most likely lead to $s_{i,k}$. Consequently, the most likely sequence for observations $o_1, \ldots, o_T$ is obtained by saving the indices j at each timestamp that maximize $V_{i,k}$, starting in $\max_w V_{T,w}$. The path p is then obtained by concatenating the optimal (shortest) paths between successive states in the most likely sequence.

Using the shortest segments to connect the states might not be the optimal solution. Therefore, in [27] is also presented an heuristic that uses features to take into consideration drivers' preferences and thus increase the likelihood of getting the right segment between states. However, this additional heuristic achieves only marginal improvements (c.f. [27]) at the expenses of computational power. Since we will be computing map-matching under several configurations (see Section 3.2), we did not implement the heuristic as to decrease execution time.

Returning to the problem defined in Section 2.1, MM is typically used as a pre-processing phase of an LBS service in which the noisy locations are mapped to the most likely position for x_i. Therefore, in our problem the user is considered to already have the real location x_i, $\forall i$. Nevertheless, an adversary can use MM to track/locate users in a road given obfuscated location/versions of x_i. In this latter scenario, the location readings (observations) are the obfuscated locations.

As for measuring privacy, we can use the adversary error from equation (3) using z_i. However, a point-by-point metric would fail to assess the effectiveness of the tracking, as the adversary error could be 0 and the estimated trajectory be different from the true trajectory. This can occur for instance when the true location is at a cross-road and the true path crosses the matched path. In such case, the true position matches the MM estimation, but the paths only overlap on that single point. Thus, we further consider a trajectory metric from the original authors of the MM technique [27], the F_1 score computed as:

$$precision = \frac{L_{correct}}{L_{matched}}$$
$$recall = \frac{L_{correct}}{L_{truth}} \quad (21)$$
$$F_1 = 2 * \frac{precision * recall}{precision + recall}$$

where $L_{matched}$ is the length of the output path, L_{truth} is the length of the corresponding ground truth and $L_{correct}$ is the length of the portions of the output path that overlap with the ground truth path. Intuitively, the precision and recall measure the length of the segments that were correctly matched as a fraction of the map-matching output and the true path, respectively. The F_1 score is then the harmonic mean between both metrics.

3 Impact of Frequency

The main objective of this work is to evaluate the impact of the frequency of location reports on the privacy level of a Geo-indistinguishable LPPM, namely the Planar Laplace (PL) [7] described in Section 2.2. Towards this goal, we obfuscate the location reports using the PL mechanism to several sub-samples of real datasets, where each sub-sample corresponds to a different frequency of reports. Subsequently, we apply state-of-the-art localization attacks as to measure the privacy level obtained through the PL mechanism against possible adversaries. The following sections will describe the datasets used in this work and detail the carried out methodology.

3.1 Datasets Characterization

To evaluate the impact of frequency one must consider both continuous and sporadic release of data. As mentioned in Section 1 there is no formal nor quantitative boundary for the frequency of updates that defines what intervals belong to the continuous or sporadic scenarios. In fact, this distinction is made based on the type of LBS application [9]. Therefore, and to allow for tuning the frequency of updates from highly frequent to "sporadic" reports, we selected three highly continuous datasets: the **Cabspotting** [29] and **Portocabs** [30] datasets, which are composed of taxi trajectories from the city of San Fracisco, USA, and Porto, Portugal, respectively; and the **Geolife** dataset [31], a dataset of GPS data captured by handheld devices.

The Cabspotting dataset [29] contains trajectories from over 500 taxis navigating in San Francisco Bay Area in a period of 30 days. It contains not only geo-location collected through a GPS at an average rate of 10 seconds, but also whether the cab is occupied or not. The Portocabs dataset is composed of trajectories belonging to 441 taxis in the city of Porto, Portugal, collected over a full year (from 2013/07/01 to 2014/06/30) with a sampling rate of 15 seconds [30]. The Geolife dataset [31] is a well known repository of GPS traces collected from 182 worldwide users in the period from April 2007 to August 2012. It contains a total of 18670 trajectories reflecting the movements under a variety of transportation means, where 91% of these have a sampling rate of 1 to 5 seconds or 5 to 10 meters per point. The majority of the trajectories lie in Beijing.

While the datasets of taxi mobility are highly continuous, these movements often have a limited timespan. In fact, most of these trajectories present a timespan under 1 hour. On the other hand, [14] shows that the Geolife dataset contains a significant amount of time-gaps between reports, that is, discontinuities in the frequency of reports. Furthermore, since our tracking attack is a road network map-matching technique [27], only vehicular trajectories can be considered. Consequently, we use the Cabspotting and Portocabs datasets to evaluate highly continuous reports and the Geolife dataset in a more sporadic scenario. It should be noted that while the Geolife is not a sporadic dataset, the continuity of reports allows to fine-tune the frequency of updates by periodically suppressing points to cover the full spectrum. Intuitively, this subsampling can be perceived as users in their quotidian trajectories making sporadic accesses to a LBS.

Our pre-processing for each dataset is as follows:

Geolife – since in a sporadic scenario there are no trajectories, we first append all traces of each user as a single array of locations and subsequently sort by date. We then filter out locations that fall outside a bounding box containing the 5^{th} ring road of Beijing as illustrated in Figure 1a. This filtering reduces the space of possible user locations ($\mathcal{X}$), which in turn allows for a finer grid for the localization attacks. A total of 65.4% of points belonging to 179 of the 182 initial users remained after this pre-processing.

Cabspotting – we first limit the trajectories to a bounding box within the San Francisco peninsula as specified in Figure 1b. Then we consider only trajectories with passenger as to remove cases where the taxi is stopped waiting for a client. Finally, we select trajectories with a duration of at least one hour, with intervals

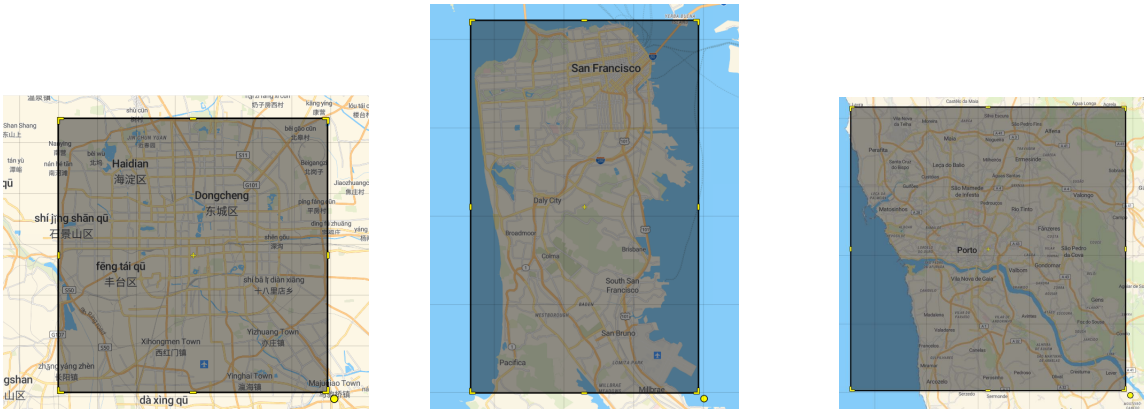

(a) Bounding box over 5^{th} ring road of Beijing used for the Geolife dataset. Approximately defined from South and North by the latitudes 39.753, 40.026, and from West and East by longitudes 116.199, 116.547.

(b) Bounding box over the peninsula of San Francisco used for the Cabspotting dataset. Approximately defined from South and North by the latitudes 37.600, 37.811, and from West and East by longitudes -122.517, -122.354.

(c) Bounding box over the city of Porto, Portugal, used for the Portocabs dataset. Approximately defined from South and North by the latitudes 41.0524, 41.257, and from West and East by longitudes -8.727, -8.456.

Fig. 1. Bounding boxes used in this work for each of the three datasets.

between reports of at most (approximately) 2 minutes as to avoid temporal discontinuities between reports. After this pre-processing, 85 trajectories remained.

After manual inspection of some of these trajectories in the map we were able to observe that the dataset contains noisy readings. For example, some GPS locations are reported in the ocean instead of in a bridge. Thus, to improve the original (noisy) data so as to build our ground-truth, we apply the MM technique described in section 2.3.3 to the original dataset. This way, we obtain a set of locations in the road network that serves as our ground-truth to compare against the locations after obfuscation and being subject to adversary attacks, as illustrated in the diagram of Figure 2. For that, we use the parameters from [32], which uses GPS data and is the work that served as baseline to the development of [27]. In [32] the estimated standard deviation was $\sigma = 6.86$m and they limited the potential locations $s_{i,k}$ to a circular radius of 50m from o_i. This discards candidate locations with low emission probability (c.f. equation (16)) and speeds up the map-matching process. For the remaining parameters we used the original values from [27]: $\lambda_y = 0.69$ and $\lambda_z = 13.35$. The restriction of the 50m radius around o_i produced observations without candidate points in some trajectories due to both the considered road network (explained in the following section) and to the noisy dataset. For these observations, we considered the nearest road network node as candidate. Furthermore, after manual inspection of the 85 trajectories, we observed that in some the taxi stays roughly in the same place to which we attribute to heavy traffic. Consequently, we removed those trajectories and ran our tests for the 63 remaining trajectories.

Portocabs – following a similar procedure to the one taken for the Cabspotting dataset, we limit the selected trajectories to a bounding-box containing the city of Porto, as illustrated in Figure 1c. From these trajectories, we select only the ones that present no missing data, that is, there is a location report every 15 seconds. Finally, we select the trajectories with a duration of 1 hour and 1 hour and 15 seconds, as to increase the number of trajectories. This resulted in 134 trajectories, which after some manual inspection as performed for the Cabspotting dataset, reduced to 123. To these final trajectories, we perform the same procedure as in the Cabspotting to obtain the ground-truth data.

The Geolife, Cabspotting and Portocabs datasets can be found in [33], [34] and [35], respectively.

3.2 Experimental Setup

The methodology for the experiments consists in subsampling the datasets, applying the Planar Laplace mechanism described in Section 2.2 and subsequently apply the localization and tracking attacks from Section 2.3. As explained in the dataset characterization, the Cabspotting and Portocabs datasets are more suitable for the map-matching attack due to being highly continuous, present no temporal discontinuities between reports and for being vehicular trajectories. Consequently, we only apply the localization attacks to the Geolife dataset, while executing both localization and map-matching attacks to the Cabspotting dataset. We use the Portocabs dataset to further validate the map-matching results.

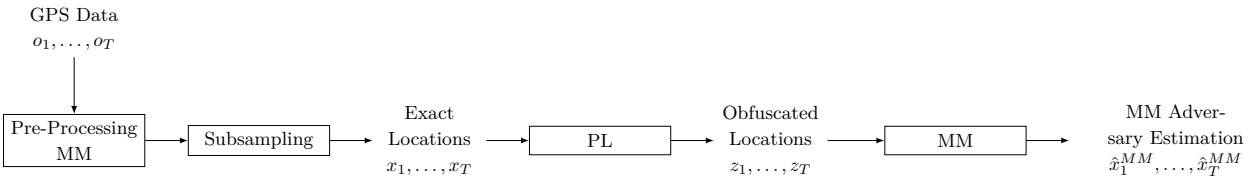

Fig. 2. Diagram of the methodology conducted for the Map-Matching attack.

3.2.1 Subsampling

To vary the frequency of reports we subsample the datasets by suppressing reports such that the interval between consecutive points is at least Δ_t. To contemplate both continuous and sporadic scenarios, several values of Δ_t are considered. For the Cabspotting and Portocabs datasets, we set Δ_t = [60, 120, 180, 240, 300, 360, 420, 480, 540, 600] seconds as our highly continuous reports. Note that in the context of map-matching, the previous values of Δ_t are already considered low sampling rate [22, 26]. For the Geolife dataset we consider a larger range of frequencies and thus set Δ_t = [480, 540, 600, 1800, 5400, 16200, 48600, 145800, 437400, 1312200] seconds. This interval goes from 8 minutes up to 15 days, and thus is comprehensive enough to encompass both continuous and sporadic scenarios.

3.2.2 LPPM

To each dataset subsample we apply the Planar Laplace described in Section 2.2 under multiple values of ϵ. Since map-matching is computationally expensive, we have used fewer ϵ values for the Cabspotting dataset. Specifically, for the Cabspotting and Portocabs datasets we have used $\epsilon = [16, 32, 64, 128]$ km^{-1} and for the Geolife dataset
$\epsilon = [1, 1.5, 2, 3, 4, 8, 12, 16, 24, 32, 48, 64]$ km^{-1}. The average quality loss is measured using equation (4).

3.2.3 Localization Attacks

Following [7, 12], we use part of the dataset for training and the remainder for testing. Thus, and as described in Section 2.3, we consider three types of attacks: **optHW**, the optimal attack using the training dataset to build the mobility profile π^{train}; **omniHW**, the optimal attack using the test dataset to build the mobility profile π^{test}, which corresponds to an omniscient adversary; and **PEBA** as described in Section 2.3.2 and using the parameters from its original work [12], with $\pi^{avg} = \pi^{train}$. The adversary error defined in equation (3) is used to measure the privacy level against these attacks.

The considered localization attacks assume the space of exact user locations $\mathcal{X}$ to be discrete. Therefore, and similarly to previous works [7, 10, 24], we have discretized the space for both datasets in a grid of equally spaced cells, where the center of the cell corresponds to a locationstamp that is common to any GPS observation within the cell. For the Geolife dataset, the 5^{th} ring road of Beijing was partitioned in cells of 2000×2000 meters for a total of 17×16 cells. For the Cabspotting dataset, and for a fair comparison between MM and the localization attacks we measure the adversary error not as the distance from the estimation $\hat{x}_i$ to the center of the grid x_i (as in equation (3)), but instead from $\hat{x}_i$ to the ground-truth point, as the tracking attack would naturally consider it. Therefore, we also evaluate the effect of the grid resolution in the adversary error. This evaluation is done for the Cabspotting dataset using the subsample corresponding to $\Delta_t = 300$s, to decrease execution time, and several grid sizes composed by squared cells of [80, 90, 100, 125, 150, 175, 200, 250] m.

The selection of the train/test data partition for the Geolife dataset was done as follows. We select the users with at least 20 points for $\Delta_t = 1312200$ seconds, our highest Δ_t. The test data for each Δ_t is then the locations of these selected users. Using these users ensures that the training data does not contain data pertaining the victims of the attacks, the same users are present in all subsamples of the dataset, and that enough test data is present to allow for profile tuning in the PEBA attack, even for the sparsest subsample (highest Δ_t value). The training data corresponds to using the locations of all users that were not selected as testing data for $\Delta_t = 480$ seconds, the lowest Δ_t. That is, the training data is the same for all Δ_t values. This avoids having poorer results for higher Δ_t due to the sparseness of the dataset. For the same reason, in the OmniHW attack the mobility profile π^{test} is also constant for all values of Δ_t and is built with the testing data with $\Delta_t = 480$. The mobility profiles π^{train} and π^{test} are therefore built using

respectively 73.4% and 26.6% of the $\Delta_t = 480$ subsampled dataset.

For the Cabspotting dataset we use the 63 trajectories as test set and all remaining trajectories contained within the bounding box from Figure 1b as training data. To be precise, we use 905255 trajectories as training data. However, it should be noted that, contrary to training a classifier, using all this data as training data does not lead to overfitting. In fact, this corresponds to an adversary which has a very precise statistic model of the average mobility profile, or in other words, a model of how a "normal" individual moves in this area.

3.2.4 Map-Matching

The diagram from Figure 2 illustrates the methodology taken when using the MM technique. The "Pre-Processing MM" computes a ground-truth from the noisy dataset as explained in Section 3.1 to which is then applied the subsampling considering the aforementioned values of Δ_t. To the subsampled locations is applied the Planar Laplace (PL) using the described values of ϵ to obtain the obfuscated reports. Finally, MM is executed on the obfuscated locations to obtain the adversary's estimations. To assess the privacy level, we compare the ground-truth against the adversary estimations using the adversary error from equation (3) and the F_1 score from equation (21). The parameters σ, λ_y and λ_z for the MM attack were estimated following the original proposal [27]. For the Cabspotting data we used trajectories within the bounding-box from Figure 1b with duration between 1 and 5 minutes with at least 2km of travelled distance (a total of 6003 trajectories). Equivalently, for the Portocabs dataset we selected trajectories within the bounding-box from Figure 1c with a duration of 5 minutes and with at least 2.5 km travelled distance, resulting in 4598 trajectories. For efficiency, and similarly to [32], we only consider candidates points within a radius r which we calculate using the inverse cumulative distribution function of the Gaussian distribution. The radius r is computed such that the circle centered at the observation contains the exact location with 90% probability. When this circle contains no candidates, which can happen due to the use of the LPPM and selected road network, the nearest road network node is used as candidate. The road network was obtained from OpenStreetMap using the OSMnx tool [36] over the area defined by the respective bounding boxes.

4 Results

This section details the obtained results. A separation based on the dataset is made, such that Section 4.1 details the results using the Geolife dataset, which focuses the sporadic scenario, and Section 4.2 describes the results using the Cabspotting and Portocabs datasets, the continuous case.

4.1 Geolife Results

For the Geolife dataset, only the localization attacks were executed. Figure 3 shows the average adversary error per Δ_t for all ϵ values and for each of the three attacks. The first thing we can observe is that the adversary error is roughly similar for any Δ_t. This allows to conclude that the frequency of updates has no significant impact on the privacy level. This is to be expected since in contrast with the tracking attack, the selected localization attacks do not take into account the temporal correlation. Consequently, the consideration of independence between reports is valid for the sporadic case. We note that while there are localization attacks which take into account the correlation between reports, such as [37], and thus our results with such attacks could differ, the reported performance in [37] is significantly lower to the attacks we consider.

Figure 3 also shows that omniHW performed better than the optHW attack, which was to be expected as the test mobility profile is used in the former. At the same time, the PEBA attack was even better than the omniHW for most values of Δ_t, thus confirming the results of the original work [12]. For the two highest values of Δ_t this was not the case, which we justify with the fact that not enough test data was present for PEBA to learn the mobility profile. Consequently, the PEBA results for these higher Δ_t are closer to the results of the optHW, which is in accordance with equation (12).

The last observation from Figure 3 is the amount of values of privacy budget (ϵ) that resolve in near zero average adversary error. Only the lowest 5 of the 12 experimented values of ϵ produced a non-negligible adversary error. For the setup we considered, values of $\epsilon \geq 8\ \mathrm{km}^{-1}$ lead to basically no privacy protection. Our results indicate that for this setup a maximum value of $\epsilon = 4\ \mathrm{km}^{-1}$ is needed for relevant privacy protection. As future work we intend to formulate a relation between the effectiveness of the optimal attack (measured by the adversary error) and the value of ϵ.

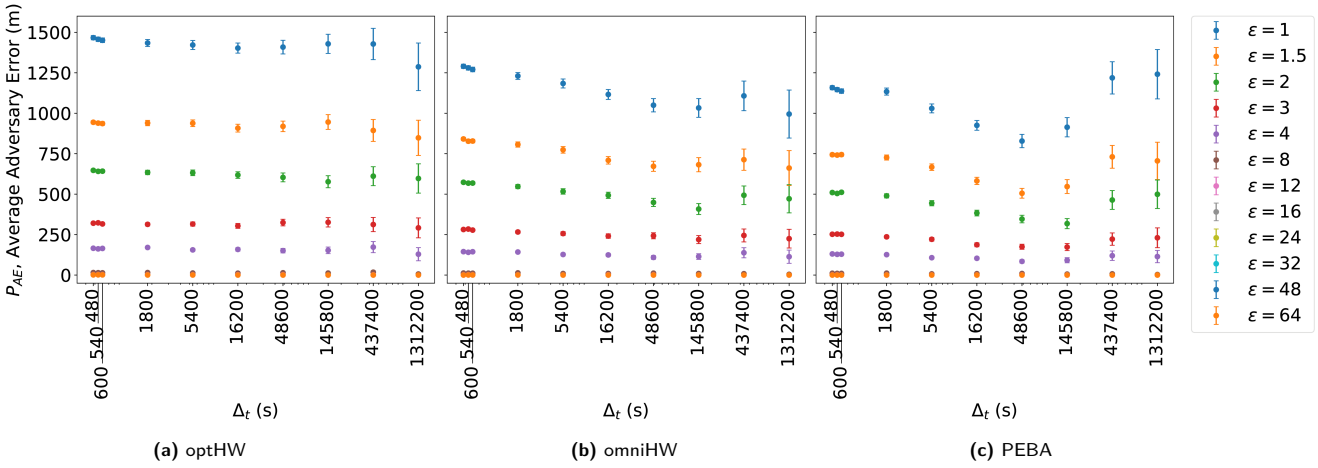

(a) optHW **(b)** omniHW **(c)** PEBA

Fig. 3. Geolife average adversary error and respective 95% confidence intervals per Δ_t for all values of ϵ for the three localization attacks. The x axis is logarithmic and the y axis and legend are shared between the three plots.

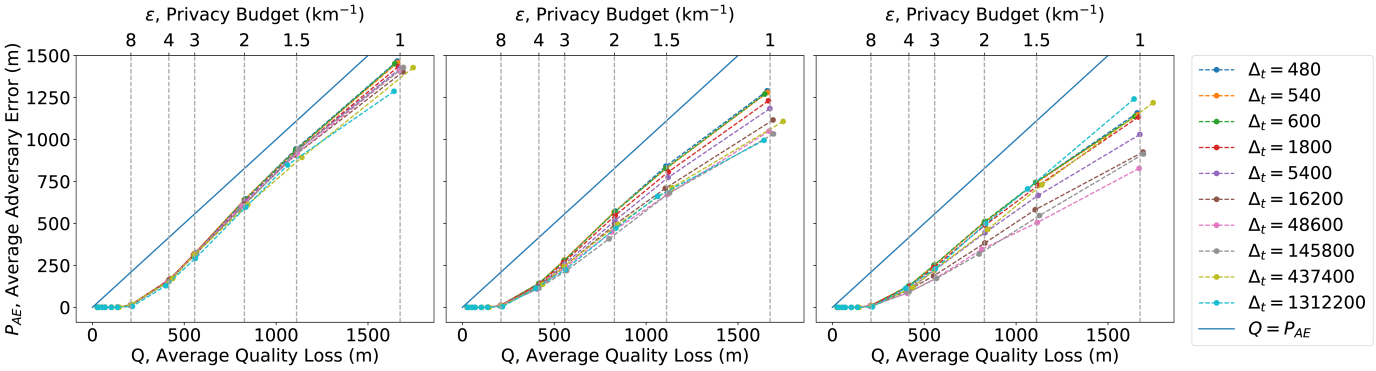

(a) optHW **(b)** omniHW **(c)** PEBA

Fig. 4. Geolife privacy versus utility for all values of Δ_t for the three localization attacks. Each color represents a Δ_t value, where the points are the pair (P_{AE}, Q), which is obtained for a particular value of ϵ. Dashed vertical lines indicate the epsilon at the empirical quality loss averaged over all values of Δ_t. The solid line represents an adversary using the report as the estimation, for reference. The y axis and legend are shared between the three plots.

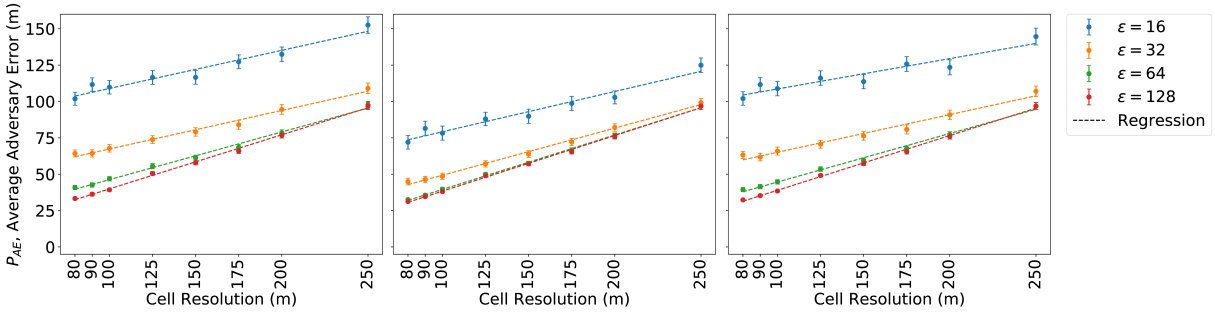

(a) optHW **(b)** omniHW **(c)** peba

Fig. 5. Effect of the grid resolution on the average adversary error (and respective confidence intervals) for each localization attack using the Cabspotting dataset with $\Delta_t = 300$ (to decrease execution time).

The last results for the Geolife dataset are displayed in Figure 4. These results show the average adversary error P_{AE} as a function of the average quality loss Q, which corresponds to the performance of an LPPM, for all values of Δ_t. Each color represents a Δ_t value, where the points are the pair (P_{AE}, Q), which is obtained for

a particular ϵ. The dashed dark lines illustrate average quality loss averaged over all values of Δ_t for each specific ϵ. The results obtained in [12] showed that the relation between P_{AE} and Q is highly linear. Looking at Figure 4, we observe this to be the case only when $P_{AE} > 0$, which as we have seen from Figure 3 occurs for $\epsilon < 8$ km^{-1}.

The second result observable from Figure 4 is the similarity of the curve for the different Δ_t, which proves again that the frequency of updates has no major effect on the privacy level using these localization attacks. In fact, it is not possible to identify a specific Δ_t that has highest average adversary error for all values of ϵ.

4.2 Cabspotting and Portocabs Results

The Cabspotting dataset is employed to assess the effect of attacks (both localization attacks as well as MM) on the continuous scenario. Since the effectiveness of the localization attacks is highly dependent on the grid resolution employed, we start by evaluating the effect of the cell size on localization attacks, as depicted in Figure 5. We can observe that for any epsilon and for any attack, there is a linear correlation between the cell resolution and the adversary error. As the cells get smaller, so does the average adversary error. Given these results, a resourceful adversary can potentially defeat obfuscation by using a very small cell resolution. However, it should be noted that as the privacy budget ϵ decreases (i.e. the obfuscation increases) the slope of the linear regression diminishes. For example, for the omniHW (Figure 5b), a grid of 100m squared cells is required to get an adversary error of around 75m for $\epsilon = 16$ km$^{-}1$. For the remaining values of ϵ ($\epsilon = [32, 64, 128]$) however, a similar adversary error is achieved using a cell resolution of 250m. That is, increasing the obfuscation also increases the computational complexity required for an attack. From the user point of view, the privacy budget ϵ thus additionally relates (with inverse proportionality) to the computational power that an adversary must employ to compromise user privacy. While the smallest average adversary error is achieved using the smallest grid resolution (80m), to decrease execution time we opt to use squared cells of 125 meters for the remainder of the results. This corresponds to a total of 189×115 cells over the peninsula of San Francisco.

Figure 6 shows the average adversary error per Δ_t and for all ϵ values for the map-matching attack (MM) and the localization attacks (optHW, omniHW and PEBA). Similarly to the results obtained for the Geolife dataset, we can observe that the average adversary error is similar for any Δ_t, which does not reveal the effect of the frequency of reports. Another relevant result from Figure 6 is that the adversary error in the map-matching is lower than the localization attacks in all epsilon values. However, as the obfuscation increases the difference in the adversary error between MM and the localization attacks diminishes. This is due to the fact that the localization attacks take into consideration the use of the LPPM and hence, the localization attacks surpass the MM performance for higher obfuscation or for a smaller grid resolution. Notwithstanding note that the adversary error is not an effective privacy metric for tracking attacks. In fact, the adversary error can be close to or even zero and the F_1 score can also be zero. This extreme case occurs, for instance, when between two exact locations the matched trajectory and the true trajectory only overlap in those two points, that is, the trajectories are disjoint except in the end-points.

To assess the impact of the frequency of updates in the privacy level of geo-indistinguishability, Figure 7 presents the effect of the privacy budget ϵ in the F_1 score. It is visible that varying the value of ϵ has more effect when higher sampling rates (i.e. lower values of Δ_t) are employed. As the frequency becomes smaller (larger Δ_t values), there is fewer correlation between points, which naturally harms the efficacy of MM, irrespectively of the ϵ value employed. This indicates a relevant trade-off between the value of the privacy budget ϵ of geo-indistinguishability and the sampling frequency, in where lower values of ϵ can cause more obfuscation, thus possibly compensating higher frequency rates.

Comparing our results with those of the proposal of the MM technique [27], it is clear that our F_1 scores are significantly lower. The two main differences that can be the source for this disparity are the dataset and the road network. Our dataset is from San Francisco and therefore requires the road network from San Francisco, which is significantly denser than Singapure's road network and, more importantly, highly symmetric. Consequently, multiple optimal (shortest) paths might exist between states of the map-matching leading to a F_1 score of zero for these segments.

As aforementioned, to further validate our map-matching results we considered Portocabs as an additional dataset of highly continuous location reports. Figure 8 presents the results obtained for this dataset. Comparing with the results obtained for the Cabspotting dataset and illustrated in Figure 7, it is clear that the same conclusions can be drawn. Specifically, the degradation of the F_1 score with the increase in Δ_t (de-

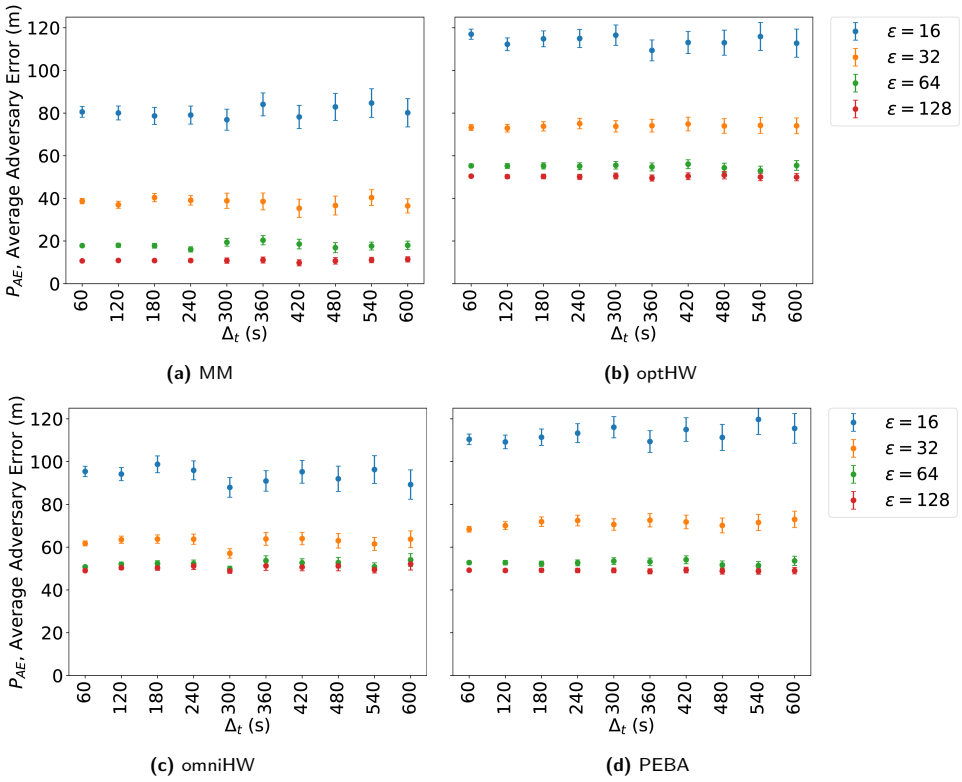

Fig. 6. Cabspotting average adversary error and respective 95% confidence intervals per Δ_t for all values of ϵ for the MM technique and the three localization attacks.

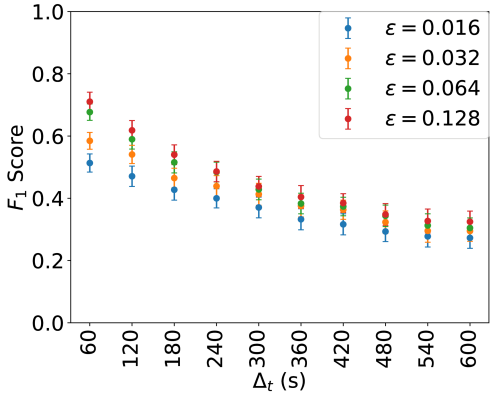

Fig. 7. Effect of the epsilon and frequency of reports (Δ_t) in the F_1 score of the MM technique for the Cabspotting dataset. 95% confidence intervals are represented as the vertical lines.

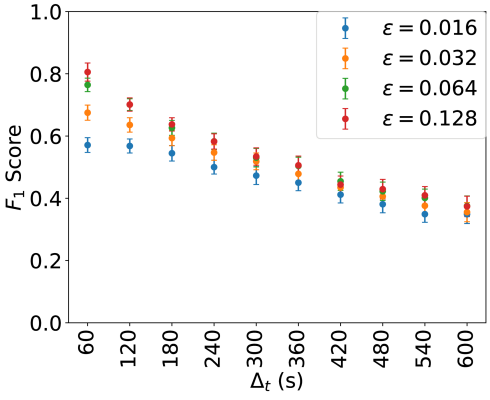

Fig. 8. Effect of the epsilon and frequency of reports (Δ_t) in the F_1 score of the MM technique for the Portocabs dataset. 95% confidence intervals are represented as the vertical lines.

crease in frequency of reports) and with the decrease in ϵ. It is observable that the F_1 score for this third dataset is slightly higher than for the Cabspotting case. This difference can be attributed to the already raised fact that the road network in San Francisco is highly symmetric, which can produce a relevant number of equally optimal (shortest) paths between nodes.

4.3 Limitations

From the previous results it was possible to conclude that, given our setup, the frequency of updates does not have a significant impact on the adversary error. An inherent limitation of this finding is the limited number of attacks considered, and specifically, tracking attacks. However, we do note that, the choice of

this specific attack was justified by the effectiveness in modelling temporal correlations using hidden Markov models [11, 24]. As future work, we would like to expand this work by considering other tracking attacks, such as Kalman filtering. Similarly, we would like to consider the effects of the frequency of updates in LPPMs which take into consideration the correlation between reports (e.g. [11, 23, 38]). These type of LPPMs are often output-based as opposed to the memoryless considered in this work, that is, the previous reports are considered when reporting the new obfuscated location [12]. Consequently, the frequency of updates should not only impact the attack success, but also the effectiveness of the privacy-preserving mechanism.

The conducted methodology also had some limitations. For example, the datasets are not sporadic, and arguably, subsampling for different values of minimum interval between reports (Δ_t) might not necessarily resemble a sporadic dataset. Nonetheless, to allow for fine-tuning the frequency a continuous dataset is required. Additionally, and similarly to [38], we argue that such subsample can be perceived as users in their quotidian trajectories making sporadic accesses to a LBS.

5 Related Work

LBSs have been classified as either sporadic or continuous based on the frequency of location reports [9]. In turn, the amount of disclosed information shapes the possible attacks that can be carried by an adversary [21]. In fact, preserving location privacy is challenging due to the fact that human mobility traces are highly unique [39, 40], that points-of-interest (POI) act as quasi-identifiers [41, 42], that is, data that can be used in combination with other available information towards deanonymization, and that individual's traces are extremely predictable given past location history [40]. Furthermore, visited locations may reveal personal interests, habits and even social connections [1, 6].

Localization and tracking techniques are the most general type of attacks in the context of user-centric location privacy [9]. Having exact geolocation data then allows for more specific inference attacks [2, 16, 43], such as the extraction of sensitive locations, as the data becomes more precise [6]. Tracking techniques consist in following a user over time and space, whereas localization techniques have as objective to localize the user at certain points in time [10]. Therefore, tracking attacks make use of the continuous releases of data while localization attacks are commonly used in the sporadic case [9]. Surveys on location attacks can be found in [1, 6, 21].

In this work, we implemented state-of-the-art localization and tracking attacks to effectively measure the privacy level from the point-of-view of a powerful yet realistic adversary. As future work we intend to complement our experiments with specific inference attacks, such as the extraction of sensitive locations. For the tracking attack we have focused on a state-of-the-art map-matching technique [27]. Map-matching (MM) is used in navigation systems to continuously pin-point vehicles on road-networks given noisy location readings [22]. While other techniques, such as using a Kalman Filter [1], have been showed to work when tracking users, the MM technique from [27] uses a Markov chain to model temporal transitions, which has shown effective in modeling temporal correlations in location data [11, 24]. For the localization attacks, we have implemented the state-of-the-art: the optimal attack given a mobility model [10] and PEBA [12] an heuristic that learns the user profile as the user reports locations. While these localization attacks disregard temporal correlation, they have been shown to perform well even when considering a coarsed discretization of the user locations' space [10, 12].

Similarly to the attacks, LPPMs have been developed for both continuous [11, 23] and sporadic scenarios [4, 7, 12]. While in the continuous scenario, temporal correlation between subsequent reports is typically considered, in the sporadic scenario reports are considered to be independent [9, 12]. Thus, this latter type of LPPMs completely disregard the frequency of usage and temporal correlations.

Geo-indistinguishability [4, 7], a formal privacy notion based on differential privacy [5] has been proposed to design LPPM with privacy guarantees irrespectively of the background information available to an adversary. Due to these properties, geo-indistinguishability has been raising increasing research interest [7, 12, 14, 44] and we therefore focus on this particular notion.

As in standard differential privacy, the privacy disclosure of geo-indistinguishability degrades linearly with the number of queries [4]. Therefore, this LPPM is only effective while the number of queries remains low as in a sporadic scenario. In a later work [38], the same authors proposed an adaptation to the case of location traces to reduce the privacy budget (ϵ) consumption through the use of a private prediction test. This test predicts the next location based on the correlations of past reports. If the correlation is high, no obfuscation occurs and the

predicted location is report instead. This test consumes a privacy budget that is smaller than generating a new obfuscation, and therefore improves on the linear privacy degradation.

The work in [44] also proposes an adaptation to geo-indistinguishability in which the privacy budget (ϵ) is adjusted for each location report according to the correlation between past reports. If the correlation is high, measured through a simple linear regression, then the mechanism adjusts to increase privacy (utility decreased) by decreasing ϵ, and vice-versa for low correlation.

A novel approach to differential privacy in continuous location reports was studied in [11], in where the proposed mechanism achieves the optimal lower privacy bound. Specifically, instead of a linear (w.r.t. the number of queries) privacy degradation, a logarithmic degradation is achieved instead. In fact, this work is considered state-of-the-art for differentially private LPPM for the continuous release of data in an online scenario (c.f. [7]).

From the above-mentioned LPPMs it is clear that for effective privacy protection in location traces, correlation between reports must be effectively captured. However, and to the best of our knowledge, the evaluation on how the frequency of updates (or equivalently, the correlation) impacts the privacy level was yet to be made. Differential privacy provides formal guarantees on the privacy disclosure of data. Nevertheless, in the context of location data, this privacy metric may be misleading when compared to the adversarial error [13]. This work attempts to fill this gap through an empirical analysis and consequently to challenge/validate the consideration of independence between reports under varying frequencies of reports. [14] and [44] showed that regression analysis can be used by an adversary to estimate the user position due to the temporal correlations. However, [14] showed that these results varied greatly depending on the estimator function and that additionally, discontinuities in the traces produced non-negligible outliers. In contrast, this work uses state-of-the-art localization attacks and a tracking attack against geo-indistinguishability under multiple privacy and adversary configurations.

6 Conclusion

As users report even an obfuscated variant of their location to a Location-Based Service (LBS), information is being disclosed. The amount of usage of these services, or in other words, the frequency of reports directly impacts the correlation between reports which in turn can be used by an adversary to further degrade privacy. Geo-indistinguishability has been proposed as a formal notion based on differential privacy to bound the amount of information released on independent queries. However, the analysis on how the frequency of queries impacts the level of privacy in geo-indistinguishability was yet to be made.

In this work we analyze the effects of the frequency of updates in the privacy level of geo-indistinguishability. We evaluate privacy and utility against state-of-the-art localization attacks and a tracking attack. Results show that the frequency of updates has in fact low significance in the privacy level, principally in the sporadic release of data. These results provide practical evidence that the consideration of independence between reports can effectively be assumed in the sporadic scenario. However, in the continuous scenario, the frequency of updates directly impacts the effectiveness of the attacks, with high frequencies leading to more privacy disclosure. In such case, obfuscation degraded the correlation between reports and consequently the effectiveness of the attack, thus acting as a countermeasure to high update frequencies. Our experiments with several values of the privacy budget reveal that there is an upper bound that is required for effective privacy protection, such that values above that threshold will result in no effective privacy. Moreover, our evaluation depicts a trade-off between the frequency of reports and the privacy budget of geo-indistinguishability, showing that lowering the frequency or increasing the the level of noise (i.e. decreasing the privacy budget) are effective measures that can be applied independently against continuous gathering of location data.

Acknowledgment

The work presented in this paper was carried out in the scope of project COP-MODE, that has received funding from the European Union's Horizon 2020 research and innovation programme under the NGI_TRUST grant agreement no 825618, and the MobiWise project: From mobile sensing to mobility advising (P2020 SAICTPAC/0011/2015), co-financed by COMPETE 2020, Portugal 2020 - Operational Program for Competitiveness and Internationalization

(POCI), European Union's ERDF (European Regional Development Fund), and the Portuguese Foundation for Science and Technology (FCT). Ricardo Mendes wishes to acknowledge the Portuguese funding institution FCT - Foundation for Science and Technology for supporting his research under the Ph.D. grant SFRH/BD/128599/2017.

The authors would like to thank Simon Oya for sharing the code for the localization attacks used in [12] and for being always available to explain some concepts which helped us throughout our implementation. We would also like to show gratitude to George R. Jagadeesh for enlightening us regarding the implementation of his work described in [27]. Last but not least, we thank the anonymous reviewers and our shepherd Dr. Sébastien Gambs for their insightful suggestions that helped improve this work.

References

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

[2] S. Gambs, M.-O. Killijian, and M. N. del Prado Cortez, "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, ACM, 2010.

[3] R. Mendes and J. P. Vilela, "Privacy-Preserving Data Mining: Methods, Metrics, and Applications," *IEEE Access*, vol. 5, pp. 10562–10582, 2017.

[4] M. Andrés, N. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-Indistinguishability: Differential Privacy for Location-Based Systems," in *20th ACM Conference on Computer and Communications Security*, pp. 901–914, ACM, 2013.

[5] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*, pp. 1–19, Springer, 2008.

[6] B. Liu, W. Zhou, T. Zhu, L. Gao, and Y. Xiang, "Location Privacy and Its Applications: A Systematic Study," *IEEE Access*, vol. 6, pp. 17606–17624, 2018.

[7] K. Chatzikokolakis, E. Elsalamouny, and C. Palamidessi, "Efficient utility improvement for location privacy," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 308–328, 2017.

[8] J. Hsu, M. Gaboardi, A. Haeberlen, S. Khanna, A. Narayan, B. C. Pierce, and A. Roth, "Differential privacy: An economic method for choosing epsilon," in *Computer Security Foundations Symposium (CSF), 2014 IEEE 27th*, pp. 398–410, IEEE, 2014.

[9] R. Shokri, G. Theodorakopoulos, G. Danezis, J.-P. Hubaux, and J.-Y. Le Boudec, "Quantifying location privacy: the case of sporadic location exposure," in *International Symposium on Privacy Enhancing Technologies Symposium*, pp. 57–76, Springer, 2011.

[10] R. Shokri, G. Theodorakopoulos, J.-Y. Le Boudec, and J.-P. Hubaux, "Quantifying location privacy," in *2011 IEEE symposium on security and privacy*, pp. 247–262, IEEE, 2011.

[11] Y. Xiao and L. Xiong, "Protecting locations with differential privacy under temporal correlations," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 1298–1309, ACM, 2015.

[12] S. Oya, C. Troncoso, and F. Pérez-González, "Rethinking location privacy for unknown mobility behaviors," in *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*, pp. 416–431, June 2019.

[13] S. Oya, C. Troncoso, and F. Pérez-González, "Is Geo-Indistinguishability What You Are Looking for?," in *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society*, pp. 137–140, ACM, 2017.

[14] R. Mendes and J. Vilela, "On the Effect of Update Frequency on Geo-Indistinguishability of Mobility Traces," in *Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks*, pp. 271–276, ACM, 2018.

[15] H. Liu, X. Li, H. Li, J. Ma, and X. Ma, "Spatiotemporal correlation-aware dummy-based privacy protection scheme for location-based services," in *INFOCOM 2017-IEEE Conference on Computer Communications, IEEE*, pp. 1–9, IEEE, 2017.

[16] J. Krumm, "Inference attacks on location tracks," in *International Conference on Pervasive Computing*, pp. 127–143, Springer, 2007.

[17] R. Shokri, "Privacy games: Optimal user-centric data obfuscation," *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 2, pp. 299–315, 2015.

[18] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec, "Protecting location privacy: optimal strategy against localization attacks," in *Proceedings of the 2012 ACM conference on Computer and communications security*, pp. 617–627, ACM, 2012.

[19] S. Oya, C. Troncoso, and F. Pérez-González, "Back to the drawing board: Revisiting the design of optimal location privacy-preserving mechanisms," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1959–1972, ACM, 2017.

[20] N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Optimal geo-indistinguishable mechanisms for location privacy," in *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pp. 251–262, ACM, 2014.

[21] M. Wernke, P. Skvortsov, F. Dürr, and K. Rothermel, "A classification of location privacy attacks and approaches," *Personal and ubiquitous computing*, vol. 18, no. 1, pp. 163–175, 2014.

[22] M. Kubicka, A. Cela, H. Mounier, and S.-I. Niculescu, "Comparative Study and Application-Oriented Classification of Vehicular Map-Matching Methods," *IEEE Intelligent Transportation Systems Magazine*, vol. 10, no. 2, pp. 150–166, 2018.

[23] R. Shokri, G. Theodorakopoulos, and C. Troncoso, "Privacy games along location traces: A game-theoretic framework for optimizing location privacy," *ACM Transactions on Privacy*

and Security (TOPS), vol. 19, no. 4, p. 11, 2017.

[24] T. Murakami, "Expectation-Maximization Tensor Factorization for Practical Location Privacy Attacks," *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 138–155, 2017.

[25] M. Hashemi and H. A. Karimi, "A critical review of real-time map-matching algorithms: Current issues and future directions," *Computers, Environment and Urban Systems*, vol. 48, pp. 153–165, 2014.

[26] P. Newson and J. Krumm, "Hidden Markov map matching through noise and sparseness," in *Proceedings of the 17th ACM SIGSPATIAL international conference on advances in geographic information systems*, pp. 336–343, ACM, 2009.

[27] G. R. Jagadeesh and T. Srikanthan, "Online map-matching of noisy and sparse location data with hidden markov and route choice models," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 9, pp. 2423–2434, 2017.

[28] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.

[29] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "A parsimonious model of mobile partitioned networks with clustering," in *2009 First International Communication Systems and Networks and Workshops*, pp. 1–10, IEEE, 2009.

[30] L. Moreira-Matias, J. Gama, M. Ferreira, J. Mendes-Moreira, and L. Damas, "Predicting taxi–passenger demand using streaming data," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1393–1402, 2013.

[31] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma, "Mining interesting locations and travel sequences from GPS trajectories," in *Proceedings of the 18th international conference on World wide web*, pp. 791–800, ACM, 2009.

[32] C. Y. Goh, J. Dauwels, N. Mitrovic, M. T. Asif, A. Oran, and P. Jaillet, "Online map-matching based on hidden markov model for real-time traffic sensing applications," in *Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on*, pp. 776–781, IEEE, 2012.

[33] "Geolife gps trajectories." https://www.microsoft.com/en-us/download/details.aspx?id=52367, 2012. [Online; Accessed: 2019-12-12].

[34] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "CRAWDAD dataset epfl/mobility (v. 2009-02-24)." Downloaded from https://crawdad.org/epfl/mobility/20090224, Feb. 2009. [Online; Accessed: 2019-12-12].

[35] "Taxi service trajectory prediction challenge @ ecml pkdd 2015." http://www.geolink.pt/ecmlpkdd2015-challenge/dataset.html, 2015. [Online; Accessed: 2019-12-12].

[36] G. Boeing, "OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks," *Computers, Environment and Urban Systems*, vol. 65, pp. 126–139, 2017.

[37] T. Murakami and H. Watanabe, "Localization attacks using matrix and tensor factorization," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 8, pp. 1647–1660, 2016.

[38] K. Chatzikokolakis, C. Palamidessi, and M. Stronati, "A predictive differentially-private mechanism for mobility traces," in *International Symposium on Privacy Enhancing Technologies Symposium*, pp. 21–41, Springer, 2014.

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

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

[41] C. Bettini, X. S. Wang, and S. Jajodia, "Protecting privacy against location-based personal identification," in *Workshop on Secure Data Management*, pp. 185–199, Springer, 2005.

[42] V. Primault, S. B. Mokhtar, C. Lauradoux, and L. Brunie, "Differentially Private Location Privacy in Practice," in *Third Workshop on Mobile Security Technologies (MoST) 2014*, 2014.

[43] E. Herder, P. Siehndel, and R. Kawase, "Predicting user locations and trajectories," in *International Conference on User Modeling, Adaptation, and Personalization*, pp. 86–97, Springer, 2014.

[44] R. Al-Dhubhani and J. M. Cazalas, "An adaptive geo-indistinguishability mechanism for continuous lbs queries," *Wireless Networks*, pp. 1–19, 2017.