On the Effect of Update Frequency on Geo-Indistinguishability of Mobility Traces

Ricardo Mendes
CISUC, Department of Informatics Engineering
Coimbra, Portugal
rscmendes@dei.uc.pt

João Vilela
CISUC, Department of Informatics Engineering
Coimbra, Portugal
jpvilela@dei.uc.pt

ABSTRACT

Sharing location data is becoming more popular as mobile devices become ubiquitous. Location-based service providers use this type of data to provide geographically contextualized services to their users. However, sharing exact locations with possibly untrustworthy entities poses a thread to privacy. Geo-indistinguishability has been recently proposed as a formal notion based on the concept of differential privacy to design location privacy-preserving mechanisms in the context of sporadic release of location data. While adaptations for the case of continuous location updates have been proposed, the study on how the frequency of updates impacts the privacy and utility level is yet to be made. In this paper we address this issue, by analyzing the effect of frequency updates on the privacy and utility levels of four mechanisms: the standard planar Laplacian mechanism suitable for sparse locations, and three variants of an adaptive mechanism that is an adaptation of the standard mechanism for continuous location updates. Results show that the frequency of updates largely impacts the correlation between points. As the frequency of updates decreases, the correlation also decreases. The adaptive mechanism is able to adjust the privacy and utility levels accordingly to the correlation between past positions and current position. However, the estimator function that is used to predict the current location has a great influence in the obtained results.

CCS CONCEPTS

• **Security and privacy → Privacy protections**; *Usability in security and privacy*;

KEYWORDS

Location-Based Services, Geo-indistinguishability, Update Frequency, Correlation, Differential Privacy

ACM Reference Format:
Ricardo Mendes and João Vilela. 2018. On the Effect of Update Frequency on Geo-Indistinguishability of Mobility Traces. In *WiSec '18: Proceedings of the 11th ACM Conference on Security & Privacy in Wireless and Mobile Networks, June 18–20, 2018, Stockholm, Sweden.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3212480.3212490

WiSec '18, June 18–20, 2018, Stockholm, Sweden
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5731-9/18/06…$15.00
https://doi.org/10.1145/3212480.3212490

1 INTRODUCTION

The pervasiveness of smart devices and the always on and always connected paradigm has fostered applications that benefit from sensing the environment to provide contextualized services to its users. One category that has recently seen enormous growth in this space is the Location-Based Services (LBS) [5], in where mobile devices, such as smartphones, share their current position in order to obtain related information (e.g. finding the nearest restaurants). However, LBS providers may not be trustworthy as data may be shared with third-parties for business or financial advantages, thus posing a threat to users' privacy.

In this scenario, user's privacy can only be preserved by applying privacy-preserving mechanisms at collection time [9] (i.e. before the data leaves the mobile device), such that the true locations are obfuscated before reaching LBS providers. This transformation of the original data is therefore used in order to retain a certain level of privacy at the expense of data utility, i.e. at the expense of a degraded service due to the degradation in the quality of the reported data.

Geo-indistinguishability [2] has been introduced as a formal notion based on the concept of differential privacy [6] to design user-centric location privacy-preserving mechanisms. Geo-indistinguishability guarantees that any two points within a radius r around the user are statistically indistinguishable, that is, the reported (obfuscated) point is generated with (almost) the same probability for any point within the circle with radius r, thus concealing the exact location of the user.

In the seminal work of geo-indistinguishability [2], a mechanism named Planar Laplace (PL) achieving this privacy notion is presented, in where noise is added independently to each true location. However, the authors note on how privacy is degraded with the number of queries due to the geo-temporal correlation of subsequent points [4] and conclude that such system is satisfactory as long as the number of queries remains fairly low. This may be valid for sporadic LBS, i.e. applications where users expose their location sparsely over time [12]. However it is not satisfactory for continuous collection of location data, i.e. location traces.

For the case of location traces, correlation between geographical points can be explored to track users over time and space and even predict future locations [8]. In this scenario, the tracking success will largely depend on the obfuscation mechanism used by users and the referred correlation between points, which will in turn depend on the frequency of location updates. At the same time, adaptations of the PL mechanism have been proposed to explore this correlation to increase privacy and/or utility [1, 4, 13]. However, the study on how the frequency of updates affects the correlation between points, which in turn affect the privacy level and utility of the data under a geo-indistinguishable privacy-preserving mechanism is yet to be made.

In this work, we address this gap by exploring the effects of varying the frequency of updates in the privacy level and data utility of Location-Based Services. We perform an analysis on the effect of the frequency of updates with a real fine-grained dataset of mobility traces, that is comprehensive enough to allow us to fine-tune the frequency of updates by periodically suppressing points. This approach allows us to assess the impact on geo-indistinguishability of varying frequency of location updates, from more continuous to more sparse datasets. Furthermore, we evaluate both data utility and the privacy level under this frequency tuning using two geo-indistinguishability mechanisms: the PL mechanism and an existing adaptation proposed in [1] for location traces by taking into account the correlation between locations. We also extend the study of [1] by using two other prediction methods aside from the simple linear regression used in the original work.

The remainder of this paper is organized as follows. Section 2 gives the formal definition on geo-indistinguishability and presents the Planar Laplace mechanism and an adaptation for the case of continuous traces proposed in [1]. Section 3 details the problem to be tackle and our experimental setup, Section 4 reports on the results and Section 5 concludes this work.

2 BACKGROUND

With the increase in popularity of applications where location is shared, researchers have studied the privacy implications and proposed several privacy-preserving mechanisms [7]. These mechanisms largely defer on the objective to be protected, which can either be protecting users' identity, location(s), points of interest, social connections or even habits. And for each goal, applications may have different requirements, which impose different constraints.

Recently, geo-indistinguishability [2] has been proposed as a formal notion of location privacy based on the concept of differential privacy [6]. 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. In fact, information disclosure in differential privacy is quantitatively measured as the difference between the prior knowledge and the posterior knowledge which is bounded to 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.

Geo-indistinguishability is formally defined as follows. Consider a location privacy mechanism as a probabilistic function $K(\cdot)$ that assigns to each location $x \in \mathcal{X}$ a probability distribution on $\mathcal{Z}$, the set of all possible obfuscated locations, where $\mathcal{X}$ and $\mathcal{Z}$ are assumed to be discrete to simplify notation. A mechanism K satisfies ϵ-geo-indistinguishability iff:

$$d_{\mathcal{P}}\left(K(x), K(x')\right) \leq \epsilon d_{\mathrm{X}}(x, x') \quad \forall x, x' \in \mathcal{X} \tag{1}$$

where $d_{\mathrm{X}}(\cdot)$ 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, equation 1 states that the probability of reporting location z while standing in location x is similar to that of standing in any location x'. In fact, both probabilities differ at most by the distance between x and x' factored by a small constant ϵ, where ϵ may be used to tune geo-indistinguishability. Commonly, and as specified in the seminal work [2], this constant is set to $\epsilon = l/r$, such that for any x, x' s.t. $d_{\mathrm{X}}(x, x') \leq r$, $d_{\mathcal{P}}\left(K(x), K(x')\right) \leq l$, where d_{X} is an arbitrary metric. 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, thus preserving some degree of utility.

The seminal work on geo-indistinguishability [2] briefly discussed on how privacy is degraded linearly with the number of queries sent by a user. In fact, the authors show that a user performing n queries through a ϵ-geo-indistinguishable mechanism enjoys $n\epsilon$-geo-indistinguishability, which is only acceptable for a small n.

A natural extension for the case of continuous traces was proposed by Chatzikokolakis et al. [4] consisting in using the distinguishability metric as $d_{\mathrm{X}} = d_{\infty}$, where, for two mobility traces x and x', $d_{\infty}(x, x') = \max_i d_{\mathrm{X}}(x[i], x'[i])$, i.e., two traces are as distinguishable as their two most distinguishable points [4]. Using this notion, traces can be protected by applying noise to each location independently as long as the noise mechanism satisfies ϵ-differential privacy. In this case, the trace will be $n\epsilon d_{\infty}$-private [4], which still has the privacy linear degradation w.r.t. the number of queries.

The work in [1] explores the potential threat that arises from exploiting the correlation of subsequent reported locations under continuous (LBS) queries. This correlation can be used to degrade the privacy level of a user using geo-indistinguishability. The authors thus propose to employ a dynamic ϵ that is used to either increase the privacy level, if the correlation of the new location with past locations is high, or to increase utility if the correlation is low. The correlation is measured as the (inverse of the) error between an estimation/prediction and the exact location, where the estimation/prediction is computed using a simple linear regression. Formally, let x be the real location, $\hat{x}$ the estimation, $d_2(\cdot)$ the euclidean distance, Δ_1 and Δ_2 two thresholds, where $\Delta_2 > \Delta_1$, and $0 < \alpha < 1$ and $\beta > 1$ two constants. Then ϵ is defined as:

$$\epsilon = \begin{cases} \alpha * \epsilon, & \text{for } d_2(x, \hat{x}) < \Delta_1 \\ \epsilon, & \text{for } \Delta_1 \leq d_2(x, \hat{x}) < \Delta_2 \\ \beta * \epsilon, & \text{for } d_2(x, \hat{x}) \geq \Delta_2 \end{cases} \tag{2}$$

Equation (2) simply states that if the error between prediction and the exact location is lower than a small threshold Δ_1 (first branch), thus stating that a high correlation between past and current location exists, then privacy should be enhanced. This is done by decreasing ϵ by a factor $\alpha < 1$. However, if the distance is bigger than a larger threshold Δ_2 (third branch), then utility may be enhanced as the correlation between points is low. Thus, ϵ is increased by the factor $\beta > 1$. Otherwise, the error is at an acceptable interval $[\Delta_1, \Delta_2[$ and thus ϵ remains unchanged.

Finally, it should be noted that mechanisms achieving optimal privacy or utility for sporadic updates only have been proposed [3, 11]. Such techniques can only be implemented for discrete scenarios, that is, scenarios where the set of locations is countable and finite.

3 IMPACT OF FREQUENCY ON PRIVACY AND UTILITY

The objective of this work is to evaluate the impact of the frequency of updates on the privacy and utility of LBS users using geo-indistinguishability privacy mechanisms. We focus on two mechanisms that are suitable to be ran on mobile devices due to their computational efficiency: the PL mechanism (henceforth also referred to as standard mechanism) [2] and the adaptive mechanism [1] as described in Section 2. For the adaptive mechanism, we consider different prediction/estimation mechanisms: apart from the simple linear regression, we also consider the parrot function [4], a simple method in which the estimation/prediction corresponds to the last observation, and a polynomial regression that allows for a non-linear estimation.

The following subsections will formally define the problem, describe the dataset used in this work and detail the supra-cited privacy mechanisms and prediction methods.

3.1 Problem Definition

Consider users that report their location to a LBS provider at different update rates, while attempting to preserve their location privacy. To avoid unwanted disclosure the protection mechanim must act at collection time, that is, at the mobile device and in an online fashion. Thus, for each location update at time i, a user instead of reporting his true location, x_i, reports z_i, an ϵ_i-geo-indistinguishable location with respect to x_i. Note that ϵ_i refers to both the standard mechanism, where $\forall i \; \epsilon_i = \epsilon$, and to the adaptive method, where ϵ_i is set as defined in (2). Similarly to other works [2], we consider that the LBS knows the identity of users, and thus the focus is on the protection of the location and not of the identity of the users themselves.

The privacy level will be measured as the estimation error $d_2(x_i, \hat{x}_i)$, where $d_2(\cdot)$ is the euclidean distance and $\hat{x}_i$ is an estimation of the true location x_i. The utility level will be measured as $d_2(x_i, z_i)$, that is, the euclidean distance between the true location and the reported value.

Although the expected estimation error depends on the prediction method, we chose this metric to quantify the privacy level as it better transmits the privacy concerns of the users. In fact, ϵ only gives a quantitative approach to the amount of information that is released, which may be misleading in terms of the privacy level [10].

3.2 Experimental Setup

To evaluate the impact of frequency on the privacy and utility levels, we have subsampled a dataset (see Subsection 3.3 for details) for various minimum time interval between sequential points, Δ_t, in seconds. Specifically, two sequential points in a trace must have a time difference of at least Δ_t. Thus, increasing Δ_t corresponds to a decrease in the frequency of updates.

For each subsampled dataset, we have the true location x_i, and compute z_i, the geo-indistinguishable version of x_i using the standard PL method and each of the adaptive variants. However, before obtaining z_i in the adaptive variants, through the considered prediction methods we compute an estimation/prediction $\hat{x}_i$ of the true location, that is then used in equation (2). For that, we consider a time-window specified by ws historical points that is used to approximate the parameters of the regression model. Since there are no guarantees on which ws is best, for a given subsampled dataset we compute the results for an array of possible ws values and select the result with the lowest median estimation error over all estimated points. We use the median of the estimation error as base line for our comparisons, as the median rectifies the skewness of the results due to the existence of outliers in the estimations, as illustrated in Subsection 3.3. Selecting the lowest median estimation error translates in a better evaluation of the privacy level, and in fact, these values correspond to the lower bound in the privacy level within the considered set of ws values.

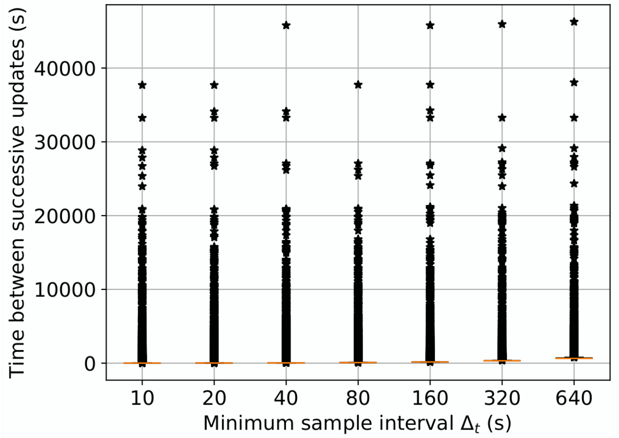

Figure 1: Boxplot of the time interval between two successive points under each of the Δ_t subsampling. Outliers represented as black stars can be perceived as time gaps between successive updates.

For the adaptive geo-indistinguishability, the following three estimator/predictor methods were used: simple linear regression, a parrot function and polynomial regression. The parrot function simply returns the previous value as the prediction. This estimator was used in [4] with the reasoning that for highly continuous updates, using the previous value as an estimator incurs in a small error. In both types of regression (linear and polynomial), the cost function used was the linear least squares function and the prediction was made for a single independent variable, that is, the prediction for the longitude was made independently of the prediction for the latitude. For the polynomial degree d, there is no obvious choice as there is high uncertainty in the velocity and movement patterns, specially when tuning the frequency of updates. Therefore, in each point we compute the estimation error for each Δ_t, each ws and for each polynomial order in the set $d \in [2, 3, 5, 7, 9, 11]$. From all these, we select only the one with lowest estimation error. Thus, and similarly to the number of numerical points (ws) considered, the polynomial degree d may vary for each point. Finally note that using a linear regression with only two past points corresponds to

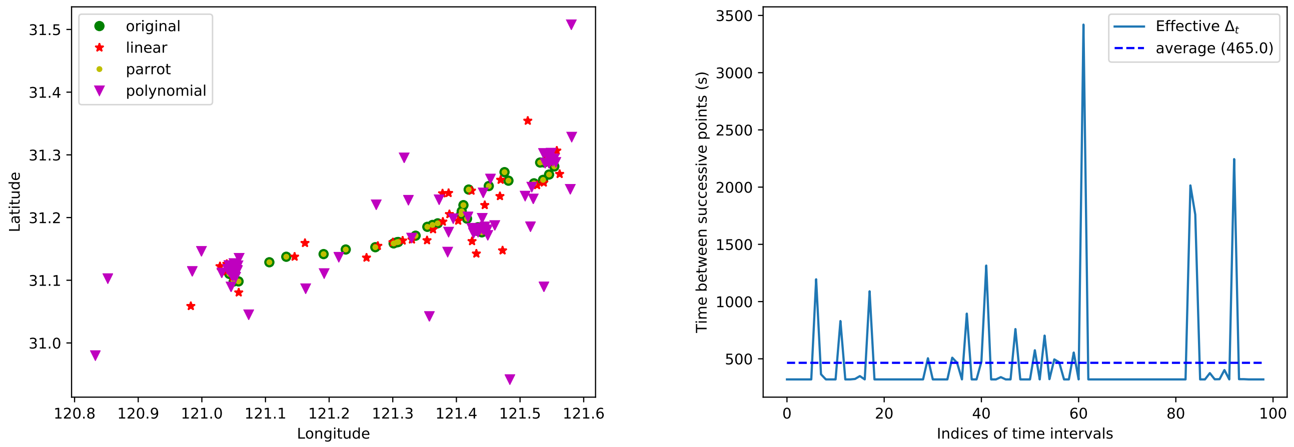

(a) Representation of the original trace points and respective estimations by each estimator function.

(b) Plot of the time between two successive updates (effective Δ_t) in the trace and respective average value.

Figure 2: Example of the effect of the time gaps between successive updates in the estimations. This example corresponds to trace 2 of user 1 using $ws = 2$ and $\Delta_t = 320$s.

Table 1: Experiment's parameters

Parameter	Designation and Values
Δ_t	Minimum time between successive points: 10, 20, 40, 80, 160, 320, 640
ϵ	Geo-indistinguishable privacy parameter: 0.001, 0.01, 0.1, 1
Δ_1	Estimation error lower threshold for equation (2): $0.96/\epsilon$
Δ_2	Estimation error higher threshold for equation (2): $2.7/\epsilon$
α	Constant to decrease ϵ in equation (2): 0.1
β	Constant to increase ϵ in equation (2): 5
ws	Number of past points to consider in the prediction 2, 3, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50
d	Degree for the polynomial regression: 2, 3, 5, 7, 9, 11

using the average velocity between these two previous points to estimate the next position.

While one can argue that better methods for estimating the position can be used, such as using a probabilistic inference based on past positions [12], we note that our estimators make use of the real locations in the estimation, which in turn are not disclosed to the provider. Therefore, while the estimation may not be optimal, the data used for the estimation is, thus representing a worst-case scenario as a measure of privacy.

Table 1 summarizes the parameters used in the experiments. The parameters for the adaptive mechanism including ϵ are the same used in the original proposal work [1] with the exception of ws where the value 100 was removed, since very few traces presented more than 100 points for $\Delta_t = 640$s.

3.3 Dataset Characterization

The dataset used in this work is the Geolife dataset [14], 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.

Since we are varying the frequency of updates by subsampling points in traces, pre-processing was applied to the dataset as to keep only traces with a sufficiently large time-span. Specifically, we have removed traces that for the sparsest condition (i.e. time interval between points of at least $\Delta_t = 640s$) had a number of points lower or equal to the maximum used window size (see table 1). Note that for smaller Δ_t, the removed traces could present more points than the maximum window size, however this data cleaning guarantees that all remaining traces were evaluated for all Δ_t and all window sizes used in this work.

After this pre-processing, the dataset contains 312 trajectories from 79 users spawning from April 2007 to May 2012 and a total of 3953369 location reports with an average update rate of $2.64s$. This represents an average number of points per trajectory of approximately 12671.

While the dataset has been subsampled with different frequency of updates Δ_t, we only ensured that each successive update had a time interval of *at least* Δ_t. However, the Geolife dataset presents a high quantity of discontinuities or gaps between successive updates. Figure 1 presents a boxplot of the time interval between two successive updates under each Δ_t, where the outliers represent these temporal gaps. These time-gaps often correspond to long periods where the user stayed in the same place, which critically impacts the quality of the regression (see also Figure 2a), not only on that specific point, but also when these points are present in the historical window of ws points.

The plots in Figure 2 show an example of this impact for a particular trace. Figure 2a presents the original trajectory and the estimations for each of the estimator functions, where one can see that both the linear and the polynomial regression have some outliers. Figure 2b shows the time interval between each successive point in this trace, where it is clear that some intervals are greatly

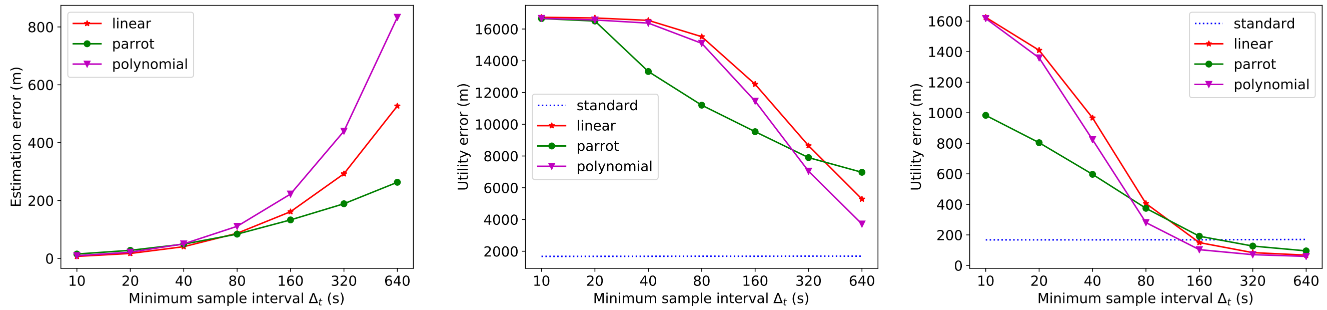

(a) Median of estimation errors per Δ_t for all adaptive geo-indistinguishable variants.

(b) Median of utility errors per Δ_t for all geo-indistinguishable variants with $\epsilon = 0.001$.

(c) Median of utility errors per Δ_t for all geo-indistinguishable variants with $\epsilon = 0.01$.

Figure 3: Plots of the estimation errors and utility errors per Δ_t.

larger than the defined minimum time between intervals Δ_t, which was 320s in this example. These peaks represent the referred temporal gaps and are the cause for the major outliers in the estimation. In fact, these gaps affect not only a certain point being estimated, but all subsequent points that include this point in the historical window defined by *ws*.

While this example illustrates a particular trajectory, this occurs in most traces and in all considered Δ_t (as depicted on Figure 1). Finally note that while these temporal outliers may impact our estimation of the correlation between points, it has limited impact on our findings because the dataset is rich enough so that by increasing increasing Δ_t we are effectively reducing the overall frequency of updates. Furthermore, a worse prediction due to outliers means that there is low correlation between points, which will be used by the adaptive geo-indistinguishable mechanism as a metric to say that the point has an acceptable level of privacy and thus utility can be increased (c.f. equation 2)

4 RESULTS AND ANALYSIS

This section presents our results. We limit our analysis to the case $\epsilon = 0.001$ and to the median values. However, we note that the conclusions hold for the remaining values of ϵ.

Figure 3a presents the median estimation errors as a function of Δ_t for each geo-indistinguishable mechanism considered. It can be seen clearly from this figure that the parrot estimator had the best performance, even for the lowest frequencies (e.g. $\Delta_t \geq 160$s). This was somehow unexpected, since by using a single point for prediction, the parrot should be affected greatly by the decrease in frequency, which is expected to happen if we consider even larger Δ_t values. However, we found that the cause for the higher estimation errors in the linear and polynomial regression are related to discontinuities in the frequency update as explained in Subsection 3.3.

Figure 3a also shows that the estimation degrades as the interval between successive points (Δ_t) increases. That is, as the frequency of updates decreases, the correlation between points also decreases. This result is also visible in Figure 3b, which shows the median of utility errors for each privacy mechanism as a function of Δ_t. Note that in the standard mechanism, since ϵ is constant and each point is treated independently, the median utility error is constant. In fact, it can be show that the standard mechanism presents an average utility error equal to $2/\epsilon$ [3], which is in line with our results.

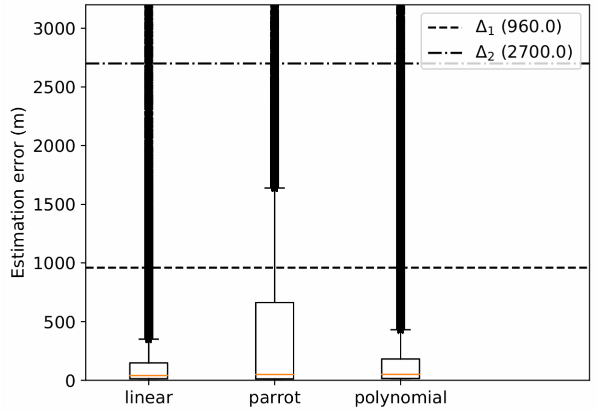

Figure 4: Boxplot of the estimation errors for $\Delta_t = 40$s and $\epsilon = 0.001$. Outliers represented as black stars, that due to the density form black solid lines, and the y axes limited to 3100.

From equation (2) we have that if the estimation error is lower than a small threshold Δ_1, which for $\epsilon = 0.001$ has the value $\Delta_1 = 960$, then the privacy level must be increased. This can be seen clearly in Figure 3b where the utility error is the highest for smaller Δ_t, as the estimation errors are the lowest (c.f. Figure 3a). As the estimation errors increase, the utility error starts to decrease. This starts occurring for $\Delta_t = 40$s for the parrot and for $\Delta_t = [80, 160]$s for the regressions as some of the estimation errors start to be higher than $\Delta_1 = 960$, as shown in Figures 4 and 5, respectively. For higher Δ_t, the estimation errors of the linear and polynomial regression are considerably greater than the parrot estimations, and in fact a great amount of errors is over $\Delta_2 = 2700$, in where ϵ is adjusted to improve utility and thus, the utility error is greatly decreased as observed in Figure 3b. Finally we note that for this value of ϵ the adaptive mechanism has a higher utility error than the standard mechanism for all Δ_t. However, this is not the case for all ϵ values as illustrated in Figure 3c. This figure presents the estimation errors per Δ_t for $\epsilon = 0.01$, where it is clear that for $\Delta_t \geq 160$ the adaptive mechanism has better utility than the standard. We omit

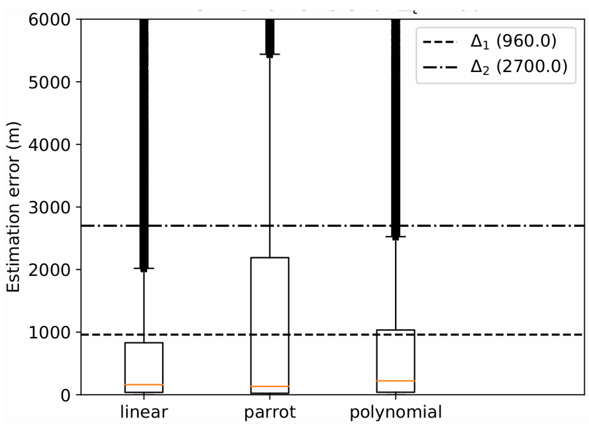

Figure 5: Boxplot of the estimation errors for $\Delta_t = 160$s and $\epsilon = 0.001$. Outliers represented as black stars, that due to the density form black solid lines, and the y axes limited to 6000.

the analysis for the remaining values of ϵ as the elated conclusions are transversal to these values.

Although the different prediction models lead to varying results from a privacy and utility perspectives, the general behavior is the same. In particular, the estimation error is bound to increase as the frequency of updates decreases, while the utility error decays with increasing frequency of updates. Naturally, the use of different prediction mechanisms can lead to different results. This calls for alternative approaches to geo-indistinguishability that take into consideration not only the correlation between points, but also the frequency of updates in a way that is oblivious to the prediction mechanisms employed.

5 CONCLUSION

Geo-indistinguishability has been proposed as a notion of privacy to design privacy-preserving mechanisms for location-based services for the sporadic release of location updates. While some adaptations to the continuous release of data have been proposed, no study has evaluated the impact of the frequency of updates on the privacy and utility levels attained.

Our experiments using a real-world dataset under several frequency of updates show that frequency has a great impact on the correlation between successive points, which in turn affect the privacy and utility of privacy-preserving location mechanisms. We have shown that the adaptive geo-indistinguishable mechanism is able to adapt the privacy and utility level by adjusting ϵ according to the correlation between past and current position. However, the measure of correlation, which quantitatively measures the privacy level, largely depends on the prediction function, thus calling for geo-indistinguishability privacy mechanisms for continuous collection of data that are oblivious to the prediction mechanisms employed.

ACKNOWLEDGMENTS

This work was carried out in the scope of projects SWING2 (PTDC/EEI-TEL/3684/2014) and MobiWise (P2020 SAICTPAC/001/2015), funded by Fundos Europeus Estruturais e de Investimento (FEEI) through Programa Operacional Competitividade e Internacionalização - COMPETE 2020, by National Funds from FCT - Fundação para a Ciência e a Tecnologia, through project POCI-01-0145-FEDER-016753, and European Union's ERDF (European Regional Development Fund). 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.

REFERENCES

[1] Raed Al-Dhubhani and Jonathan M Cazalas. 2017. An adaptive geo-indistinguishability mechanism for continuous LBS queries. *Wireless Networks* (2017), 1–19. https://doi.org/10.1007/s11276-017-1534-x
[2] Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 901–914.
[3] Konstantinos Chatzikokolakis, Ehab ElSalamouny, and Catuscia Palamidessi. 2017. Efficient Utility Improvement for Location Privacy. *Proceedings on Privacy Enhancing Technologies* 2017, 4 (2017), 308–328.
[4] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. 2014. A predictive differentially-private mechanism for mobility traces. In *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, 21–41.
[5] Subhankar Dhar and Upkar Varshney. 2011. Challenges and business models for mobile location-based services and advertising. *Commun. ACM* 54, 5 (2011), 121–128.
[6] Cynthia Dwork. 2008. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*. Springer, 1–19.
[7] Sébastien Gambs, Marc-Olivier Killijian, and Miguel Núñez del Prado Cortez. 2010. 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*. ACM, 34–41.
[8] John Krumm. 2009. A survey of computational location privacy. *Personal and Ubiquitous Computing* 13, 6 (2009), 391–399.
[9] Ricardo Mendes and João P Vilela. 2017. Privacy-Preserving Data Mining: Methods, Metrics, and Applications. *IEEE Access* 5 (2017), 10562–10582.
[10] Simon Oya, Carmela Troncoso, and Fernando Pérez-González. 2017. Is Geo-Indistinguishability What You Are Looking for?. In *Proceedings of the 2017 on Workshop on Privacy in the Electronic Society (WPES '17)*. ACM, New York, NY, USA, 137–140. https://doi.org/10.1145/3139550.3139555
[11] Reza Shokri. 2015. Privacy games: Optimal user-centric data obfuscation. *Proceedings on Privacy Enhancing Technologies* 2015, 2 (2015), 299–315.
[12] Reza Shokri, George Theodorakopoulos, George Danezis, Jean-Pierre Hubaux, and Jean-Yves Le Boudec. 2011. Quantifying location privacy: the case of sporadic location exposure. In *International Symposium on Privacy Enhancing Technologies Symposium*. Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 57–76.
[13] Yonghui Xiao and Li Xiong. 2015. Protecting locations with differential privacy under temporal correlations. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 1298–1309.
[14] Yu Zheng, Lizhu Zhang, Xing Xie, and Wei-Ying Ma. 2009. Mining interesting locations and travel sequences from GPS trajectories. In *Proceedings of the 18th international conference on World wide web*. ACM, 791–800.