

Is Geo-Indistinguishability What You Are Looking for?

Simon Oya
University of Vigo
simonoya@gts.uvigo.es

Carmela Troncoso
IMDEA Software Institute
carmela.troncoso@imdea.org

Fernando Pérez-González
University of Vigo
fperez@gts.uvigo.es

ABSTRACT

Since its proposal in 2013, geo-indistinguishability has been consolidated as a formal notion of location privacy, generating a rich body of literature building on this idea. A problem with most of these follow-up works is that they blindly rely on geo-indistinguishability to provide location privacy, ignoring the numerical interpretation of this privacy guarantee. In this paper, we provide an alternative formulation of geo-indistinguishability as an adversary error, and use it to show that the privacy vs. utility trade-off that can be obtained is not as appealing as implied by the literature. We also show that although geo-indistinguishability guarantees a lower bound on the adversary's error, this comes at the cost of achieving poorer performance than other noise generation mechanisms in terms of average error, and enabling the possibility of exposing obfuscated locations that are useless from the quality of service point of view.

KEYWORDS

Location Privacy; Privacy Metrics; Geo-Indistinguishability

1 INTRODUCTION

Geo-indistinguishability (GeoInd), a formal notion of location privacy introduced in [2], builds on the concept of differential privacy [7] to design user-centric location privacy-preserving mechanisms. To gain privacy while preserving some utility, in these mechanisms users report to service providers obfuscated versions of their actual locations. GeoInd guarantees that obfuscated locations are statistically indistinguishable from other locations within a radius around the users' real location. One of the most appealing features of GeoInd, inherited from differential privacy, is that it guarantees that, regardless of any side-information about the user she might have, the adversary learns little additional information about the real location from observing the obfuscated version.

Since its proposal [2], GeoInd has drawn a lot of attention from the research community. A first research line extends this notion to consider mobility traces instead of single locations [5, 10], or to consider semantic and geographic privacy [6]. Some works focus on how to use GeoInd, or on integrating GeoInd with other privacy metrics [9, 11] to design optimal location privacy-preserving

mechanisms, either in simplified [3] or realistic [4] scenarios. Finally, GeoInd has been also used to implement plugins to sanitize locations for its use by other mobile applications [8] or browsers [1].

A common issue in these works is that they choose GeoInd based on its core *qualitative* advantage, namely that it provides protection for the users in a region around their real location regardless of the adversary's side-information. However, they do not evaluate and reason *quantitatively* about how much protection the mechanisms provide, i.e., if the level of privacy they achieve is meaningful.

In this work, we illustrate that GeoInd can be misleading both in terms of privacy and utility. We propose an alternative definition of this privacy notion as an adversary's error, and study numerically the privacy level provided by the state-of-the-art mechanisms that guarantee this property. We also examine the trade-off between privacy and utility, showing that even though GeoInd mechanisms ensure a minimum privacy protection, this comes at the expense of performing poorly in terms of average protection, and possibly generating an obfuscated location very far away from the user.

2 GEO-INDISTINGUISHABILITY

We first describe the operation of user-centric perturbation-based sporadic location privacy mechanisms. Consider a user, Alice, that wants to get some service from a service provider from her real location $x \in \mathcal{X}$. Before exposing her location to the provider, Alice uses a location privacy mechanism f to generate an obfuscated location $z \in \mathcal{Z}$, with probability $f(z|x)$. \mathcal{X} and \mathcal{Z} are sets of locations that we assume discrete for notational simplicity, although we note that all the results in this paper are applicable to the continuous scenario. By using mechanism f , Alice trades in utility for privacy. For example, if Alice's query is "give me the bars in a radius of 100 meters from my location", releasing an obfuscated location z away from x might result in bars that are far away from her, but also protects her location since the probabilistic nature of the mechanism f prevents the adversary from learning her true location x .

We define the *multiplicative distance* between two distributions $\sigma_1(s)$ and $\sigma_2(s)$ on a set \mathcal{S} as $d_{\mathcal{P}}(\sigma_1(s), \sigma_2(s)) \doteq \sup_{s \in \mathcal{S}} \left| \log \frac{\sigma_1(s)}{\sigma_2(s)} \right|$ with the convention that $\left| \log \frac{\sigma_1(s)}{\sigma_2(s)} \right|$ is 0 if $\sigma_1(s) = \sigma_2(s) = 0$ and ∞ if only one of the two is 0.

In this scenario, geo-indistinguishability is defined as [2]:

Definition 2.1 (ϵ -Geo-Indistinguishability). A mechanism f provides ϵ -geo-indistinguishability if and only if, for all input locations $x, x' \in \mathcal{X}$, the following holds

$$d_{\mathcal{P}}(f(z|x), f(z|x')) \leq \epsilon \cdot d(x, x'), \quad (1)$$

where $d(x, x')$ is the Euclidean distance between x and x' .

The rationale behind this privacy notion is the following: by bounding the multiplicative distance, we ensure that the probability that Alice reports z when she is in x is similar to the probability

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WPES'17, October 30, 2017, Dallas, TX, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5175-1/17/10...\$15.00

<https://doi.org/10.1145/3139550.3139555>

that she reports z when she is in x' (up to a multiplicative factor of $e^{\epsilon \cdot d(x, x')}$). Therefore, an adversary observing z cannot statistically distinguish between x and x' as Alice's real location. The upper bound in (1) depends on $d(x, x')$ and ϵ . The former dependence is very intuitive: given an obfuscated location z , two locations $x, x' \in \mathcal{X}$ that are very close result harder to distinguish (i.e., $f(z|x)$ is close to $f(z|x')$) than if they were further apart. The role of ϵ , on the other hand, is to tune the degree of GeoInd. Smaller values of this parameter ensure that $f(z|x)$ and $f(z|x')$ are closer, and therefore provide a higher degree of privacy than larger values.

Prior-Agnostic Protection. GeoInd provides a privacy guarantee independent of any side information about x the adversary might have. Let $\pi(x)$ be a probability mass function over $x \in \mathcal{X}$ representing the *prior* adversary's side information about Alice's real location x . After observing z , the adversary can update her knowledge by computing the posterior probability mass function

$$p(x|z) = \frac{f(z|x) \cdot \pi(x)}{\sum_{x' \in \mathcal{X}} f(z|x') \cdot \pi(x')} . \quad (2)$$

By using (1) and (2), it is easy to show that GeoInd implies

$$d_{\mathcal{P}}(p(x|z), \pi(x)) \leq \epsilon \cdot d(\pi) , \quad (3)$$

where $d(\pi)$ is the maximum distance between two locations x and x' such that $\pi(x) > 0$ and $\pi(x') > 0$. In other words, GeoInd ensures a certain degree of similarity between the adversary's prior and posterior information about Alice's real location, for any prior π . Note that GeoInd is not an absolute privacy guarantee, but only ensures that given z the adversary gets no significant *extra* accuracy with respect to the prior. However, if given this prior the adversary can pinpoint a user's location to a small region in the map (small $d(\pi)$), then even though z does reveal little information about x , the adversary's estimation of x will still be accurate.

Choosing the Privacy Parameter. The general approach to selecting a proper value for the parameter ϵ is to pick a *privacy level* ϵ^* and a *privacy radius* r^* , and set $\epsilon = \epsilon^*/r^*$. This ensures that, when Alice is in x and releases z , her location is statistically indistinguishable from all the other locations x' within a radius of r^* around her, i.e., $d_{\mathcal{P}}(f(z|x), f(z|x')) \leq \epsilon^*$ as in (1).

Quantitatively, however, it is hard to determine if a bound on the multiplicative distance ϵ^* gives "enough privacy". This is reflected in the literature, where there is no consensus about which value of the bound in (1) denotes a high degree of indistinguishability. In the seminal paper [2], Andrés et. al choose $\epsilon^* = \log 2$ in a radius of $r = 200$ meters as the highest privacy level. This bound is used by some follow-up works [1, 6], while others take different values of ϵ^* , ranging from $\epsilon^* = \log 10$ [5] to $\epsilon^* = \log 1.4$ [4].

3 GEOIND AS AN ADVERSARY ERROR

In this section, we introduce an alternative characterization of GeoInd as an adversary error. This characterization helps us in providing more intuition behind the privacy level obtained for a specific value of the privacy parameter ϵ , and in understanding the protection it provides beyond the upper bound expressed in (1).

Consider that the adversary's side information is that Alice is equally likely in either of two locations x and x' , i.e., $\pi(x) = \pi(x') = 0.5$. After observing z , the adversary has to decide between x and x' .

We refer to this adversary as the *decision adversary*. Assume, without loss of generality, that $f(z|x) \geq f(z|x')$, and thus the optimal decision in terms of minimizing the adversary's probability of error is deciding that Alice's location is x . In this case, the adversary's probability of error is

$$p_e(x, x', z) = \frac{f(z|x')}{f(z|x) + f(z|x')} . \quad (4)$$

Then, Geo-indistinguishability can be defined as follows:

LEMMA 3.1 (ϵ -GEO-INDISTINGUISHABILITY AS ERROR). *A mechanism f guarantees ϵ -geo-indistinguishability if and only if, for any pair of input locations $x, x' \in \mathcal{X}$ and any output location $z \in \mathcal{Z}$, it ensures that the minimum probability of error p_e^* of the decision adversary described above is*

$$p_e(x, x', z) \geq p_e^* = \frac{1}{1 + e^{\epsilon \cdot d(x, x')}} . \quad (5)$$

It is easy to see that the GeoInd definitions in (1) and (5) are equivalent by substituting (4) in (5) and operating. We have chosen $\pi(x) = \pi(x') = 0.5$ as prior knowledge for the decision adversary, as this ensures the guarantee in (1), but we note that when $\pi(x) \neq \pi(x')$ GeoInd does not guarantee a minimum probability of error against this adversary.

This alternative definition of GeoInd allows us to intuitively interpret the privacy guarantee achieved by f . For example, given $\epsilon = 2\text{km}^{-1}$ and two locations x and x' separated $d(x, x') = 0.5\text{km}$, according to (1) the multiplicative distance between $f(z|x)$ and $f(z|x')$ is bounded by 1. However, whether an upper bound on the multiplicative distance of 1 is a reasonable level of protection is not clear. In terms of probability of error, (5) bounds the adversary's error to be $p_e \geq p_e^* = 0.27$: before observing z the decision adversary has a probability of correctly guessing Alice's input location of 0.5, and after the release her probability of success is in the worst case 0.73. It is difficult to consider this worst case probability as "high indistinguishability", which contradicts the GeoInd idea of achieving indistinguishability for every pair of input locations. We explore the implications of this interpretation in the next section.

4 GEOIND IN NUMBERS

In this section, we quantitatively evaluate the privacy and utility achieved by GeoInd mechanisms. For privacy, given a fixed distance between two locations x and x' , we measure both the upper bound on the multiplicative distance in (1), $\epsilon^* = \epsilon \cdot d(x, x')$, and the lower bound on the probability of error (4) of the decision adversary, p_e^* . For utility we consider two metrics: the average loss \bar{r} , measured as the average Euclidean distance between the actual location of the user x and the obfuscated location z [3, 4, 6, 9–11], and the radius of the circular region centered around x where z is with probability 0.95, denoted by r_{95} [1].

We evaluate two GeoInd mechanisms. First, the planar Laplace mechanism, proposed in the seminal work [2], implemented in [1, 8], and used as a baseline for comparison in [3, 4]. This mechanism generates obfuscated locations z from the actual location x by adding 2-dimensional Laplace noise to the latter. Second, the planar Laplace with remapping [4], current state-of-the-art. This mechanism first generates a temporary location z' by adding 2-dimensional Laplace noise to x , and then performs a deterministic remapping from z' to




Figure 1: Performance of the Planar Laplace mechanism.

z that is designed to minimize the average loss \bar{r} of the scheme while providing the same privacy guarantees as the original version. This remapping is computed using a dataset with information about the popularity of each input location $x \in \mathcal{X}$, and therefore the mechanism is tied to a particular dataset. We leave optimal mechanisms that can only be implemented in simple discrete scenarios [3, 9, 11] out of our evaluation, as the low-resolution quantization needed to implement them in a real scenario makes them suboptimal (cf. [4]).

Planar Laplace Mechanism. Given a privacy value ϵ , the average loss of the Laplace mechanism is $\bar{r} = 2/\epsilon$, and r_{95} can be computed analytically using the Lambert W function (cf. [2]). The value of p_e^* can be computed from ϵ following (5). We show p_e^* and ϵ^* for this mechanism, when locations are separated $d(x, x')$ meters, in Figure 1 for different utility levels. As expected, as we add more noise (larger \bar{r} or r_{95}), protection improves (larger p_e^* or smaller ϵ^*).

To better understand the trade-off between privacy and utility let us consider as reference a privacy level $p_e^* = 0.4$ (i.e., the decision adversary succeeds at most 60% of the times). To obtain such protection level in a radius of r^* one needs to add Laplacian noise with average loss of $\bar{r} \approx 5r^*$. This results 5% of the time on an obfuscated location z further than $r_{95} \approx 12r^*$ from the real location x . Consider that we want $p_e^* = 0.4$ in locations within a radius of $r^* = 200\text{m}$. In this case, the obfuscated location would be on average $\bar{r} = 1\text{km}$ away from the real location, and 5% of the time it would be further than 2.3km away (yellow line in Fig. 1). In applications that are not sensitive to large amounts of noise (e.g., weather forecast) this might be reasonable. However, in other applications where one would require a utility in the same order of magnitude as the privacy protection (e.g., finding nearby points of interest), the Laplace mechanism and, up to some extent, GeoInd, are not desirable.

Planar Laplace with Optimal Remapping. Since this mechanism cannot be evaluated analytically, we follow the empirical approach in [4]: we use 80% of the users from Gowalla dataset¹ to design the remapping function, and use the remaining 12 112 users as a testing set to evaluate the utility of the mechanism after remapping. We generate an output z for 20 000 random checkins from testing set users, for values of ϵ from the previous experiment ($\epsilon = \{6.67, 4, 2, 1\}$, in km^{-1}), and use them to compute \bar{r} and r_{95} .

¹<https://snap.stanford.edu/data/loc-gowalla.html>

The results in terms of p_e^* and ϵ^* vs. $d(x, x')$ coincide with the ones in Fig. 1, but we obtain much better quality: $\bar{r} = 159, 266, 578$ and 1271 meters, i.e., 37 – 47% smaller than plain Laplace. The 95% loss percentile in each case is $r_{95} = 565, 999, 2146$, and 4162 meters, which is only a 10 – 21% reduction from the planar Laplace without remapping. To obtain a protection of $p_e^* = 0.4$ in a radius of r^* around the real location, in this scenario one needs to add noise with a loss of roughly $\bar{r} \approx 3r^*$ and $r_{95} \approx 10r^*$. Although the average loss reduction is considerable, the utility cost is still large compared to the radius of the privacy region this mechanism ensures. This highlights the importance of analyzing GeoInd numerically to understand the actual privacy vs. utility trade-off it provides.

5 OTHER PROPERTIES OF GEOIND

So far we have studied the lower bound (p_e^*) GeoInd mechanisms provide on the probability of error of the decision adversary (p_e). We now study other properties of GeoInd mechanisms against this adversary. For this purpose, we evaluate three mechanisms: the planar Laplace mechanism, described above, and the Gaussian and uniform circular mechanisms. The latter mechanisms generate z by adding to the real location x , respectively, 2-dimensional Gaussian noise and uniform noise in a circle. We choose not to use remapping, as its improvement would be similar for all mechanisms and thus does not influence the comparison. For each experiment, we consider two locations x and x' separated a distance $d(x, x')$ and generate z using the three mechanisms. Then, we measure the probability of error of the decision adversary p_e in that realization, and repeat this 20 000 times for different values of $d(x, x')$ and \bar{r} .

The average probability of error of these mechanisms, denoted by \bar{p}_e and computed by averaging the 20 000 samples of p_e , is shown in Figure 2. Given an average loss \bar{r} , both Gaussian and circular mechanisms achieve a larger average error than the Laplace mechanism, up to a certain distance $d(x, x')$ marked with \bullet in the figure. We do not show these marks for $\bar{r} = 1000$ and $\bar{r} = 2000$, but they also lay close to $\bar{p}_e = 0.1$. At these points, the Laplace mechanism achieves $p_e^* \approx 0.01$ and $\bar{p}_e \approx 0.1$ for all tested values of \bar{r} . The fact that the Laplace mechanism performs better from \bullet onwards is not significant: in these scenarios, regardless of the mechanism, the adversary guesses the right location with an average probability larger than 0.9, i.e., no mechanism provides privacy. We conclude that, in all relevant scenarios (i.e., reasonable privacy levels), the Gaussian and circular mechanisms achieve a larger average error than the Laplace mechanism. This means that using GeoInd as a way of providing an *average protection level* against an adversary with unknown side-information [9, 11] is not recommended, as it is not the guarantee that this notion provides.

Figure 3 shows the percentage out of the 20 000 realizations where the Gaussian and circular mechanisms achieve a larger p_e than the Laplace mechanism. We see that these mechanisms achieve a larger probability of error more often than the Laplace mechanism when x and x' are separated up to a distance $d(x, x')$ corresponding to the points marked with \blacksquare (Gaussian) and \blacktriangle (circular). The figure also shows the performance of the Laplace mechanism in terms of p_e^* and \bar{p}_e at these points (these values remain almost constant when changing \bar{r}). Similarly to the previous case, these values represent a very low privacy regime, i.e., for all reasonable privacy levels,

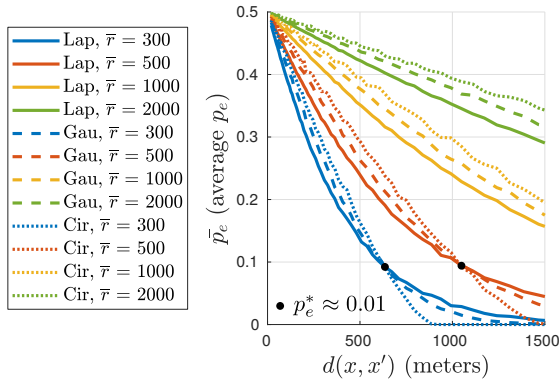


Figure 2: Performance in terms of the average probability of error of the decision adversary.




Figure 3: Percentage of times the Gaussian and Circular mechanisms outperform the Laplace mechanism.




Figure 4: Normalized histogram of p_e , for $d(x, x') \approx 100$ and $\bar{r} = 500$ meters ($p_e^* = 0.4$).

the Gaussian and circular mechanisms are more likely to achieve a larger p_e than the Laplace mechanism. This is better illustrated in Figure 4, which shows the normalized histogram of the probability of error p_e provided by the Gaussian and Laplace mechanisms for $d(x, x') = 100\text{m}$ and $\bar{r} = 500\text{m}$. As expected, the Laplace mechanism ensures a minimum probability of error (as often as the Gaussian mechanism). These experiments reinforce the idea that Geolnd is not a “cure-all” privacy guarantee against a prior-agnostic adversary.

All the above experiments compare mechanisms offering the same average loss \bar{r} . In terms of r_{95} , the Laplace mechanism ($r_{95} \approx 2.37\bar{r}$) performs worse than the circular ($r_{95} \approx 1.46\bar{r}$) and the Gaussian ($r_{95} \approx 1.95\bar{r}$) alternatives. Thus, compared with a fixed r_{95} , the Laplace mechanism would perform *even worse* than the others.

6 WHERE TO GO NOW

Geo-indistinguishability, which provides differential privacy-like guarantees in the location privacy scenario, has drawn a lot of attention from the community. However, our quantitative evaluation shows that the (worst-case or average) privacy guarantees it provides are unsatisfying unless utility is sacrificed. The main reason for this poor performance is that, in the counting queries on a database scenario where differential privacy was initially proposed [7], queries have *low sensitivity*, i.e., the contribution of a single user does not significantly affect the outcome. This enables the achievement of a high privacy level (e.g., $\epsilon^* = 0.01$) without introducing much noise, thus preserving utility. In the location scenario where Geo-indistinguishability operates, each query has high sensitivity and therefore requires large noise to provide protection. For instance, to achieve Geolnd with $\epsilon^* = 0.01$ between locations in an area of 100m, the average loss is 20km. Moreover, Geolnd can only be achieved at the expense of having an *unbounded maximum quality loss* since the guarantee (1) no longer holds if the mechanism is truncated to ensure a minimum utility for the users.

This does not mean that Geolnd should be abandoned, but we argue that it should be carefully configured, understanding the type and amount of protection it provides. Our Geolnd characterization as an adversary error should assist in this task, as it helps to quantitatively interpret the degree of protection provided. We have

also shown that in some scenarios there are levels of protection that are not achievable without unreasonable utility loss. Potential solutions could be to use bandwidth as a resource to improve utility [2], or re-design location queries to have lower sensitivity (e.g., aggregating queries [7] locally, at the user level).

ACKNOWLEDGMENTS

This work is partially supported by EU H2020-ICT-10-2015 NEXTLEAP (GA n 688722), the Agencia Estatal de Investigación (Spain) and the European Regional Development Fund (ERDF) under projects WINTER (TEC2016-76409-C2-2-R) and COMONSENS (TEC2015-69648-REDC), and by the Xunta de Galicia and the European Union (ERDF) under projects Agrupación Estratégica Consolidada de Galicia accreditation 2016-2019 and Red Temática RedTEIC 2017-2018. Simon Oya is funded by the Spanish Ministry of Education, Culture and Sport under the FPU grant.

REFERENCES

- [1] 2016. Location Guard. (2016). <https://github.com/chatziko/location-guard> Accessed: 2017-06-12.
- [2] Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2013. Geo-indistinguishability: Differential privacy for location-based systems. In *ACM Conference on Computer & Communications Security*. ACM, 901–914.
- [3] Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. 2014. Optimal geo-indistinguishable mechanisms for location privacy. In *ACM Conference on Computer & Communications Security*. ACM, 251–262.
- [4] Konstantinos Chatzikokolakis, Ehab Elsamouny, and Catuscia Palamidessi. 2017. Efficient Utility Improvement for Location Privacy. *Proceedings on Privacy Enhancing Technologies* 2017, 4 (2017), 210–231.
- [5] 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.
- [6] Konstantinos Chatzikokolakis, Catuscia Palamidessi, and Marco Stronati. 2015. Constructing elastic distinguishability metrics for location privacy. *Proceedings on Privacy Enhancing Technologies* 2015, 2 (2015), 156–170.
- [7] Cynthia Dwork. 2008. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*. Springer, 1–19.
- [8] Kassem Fawaz and Kang G. Shin. 2014. Location privacy protection for smartphone users. In *ACM Conference on Computer & Communications Security*. ACM, 239–250.
- [9] Reza Shokri. 2015. Privacy games: Optimal user-centric data obfuscation. *Proceedings on Privacy Enhancing Technologies* 2015, 2 (2015), 299–315.
- [10] Yonghui Xiao and Li Xiong. 2015. Protecting locations with differential privacy under temporal correlations. In *ACM Conference on Computer & Communications Security*. ACM, 1298–1309.
- [11] Lei Yu, Ling Liu, and Calton Pu. 2017. Dynamic Differential Location Privacy with Personalized Error Bounds. (2017).