

Practical Mechanisms for Location Privacy

The continuously increasing use of location-based services poses an important threat to the privacy of users. A natural defense is to employ an obfuscation mechanism, such as those providing geo-indistinguishability, a framework for obtaining formal privacy guarantees that has become popular in recent years.

Ideally, one would like to employ an optimal obfuscation mechanism, providing the best utility among those satisfying the required privacy level. In theory optimal mechanisms can be constructed via linear programming. In practice, however, this is only feasible for a radically small number of locations. As a consequence, all known applications of geo-indistinguishability simply use noise drawn from a planar Laplace distribution.

In this work, we study methods for substantially improving the utility of location obfuscation, while having practical applicability as a central constraint. We provide such solutions for both infinite (continuous or discrete) as well as large but finite domains of locations, using a Bayesian remapping procedure as a key ingredient. We evaluate our techniques in two real world complete datasets, without any restriction on the evaluation area, and show important utility improvements wrt the standard planar Laplace approach.

1 Introduction

In recent years, the increasing availability of location information about individuals has led to a growing use of systems that record and process location data. Examples include Location Based Services (LBSs), location-data mining algorithms to determine points of interest, and location-based machine learning algorithms to predict traffic patterns.

While these systems have demonstrated to provide enormous benefits to individuals and society, the growing exposure of users' location information raises impor-

tant privacy issues. On one hand, location information itself may be considered as sensitive. For instance, there have been reported episodes of men tracking women with GPS or specific applications [1, 2], and in California location records have been used in divorce proceedings to prove claims about suspicious displacements of spouses [3]. Furthermore, location data can be easily linked to a variety of other information that an individual usually wishes to protect: by collecting and processing such data on a regular basis, it is possible to infer an individual's home or work location, sexual preferences, political and religious inclinations, etc. [4].

It is not surprising, therefore, that a lot of effort has been dedicated to design and implement methods for protecting the user's privacy, while preserving the utility and the dependability of location data for their use in location-based systems. In general, all computational methods for privacy protection are based on degrading the precision of information. In the case of location privacy, one typical way to reduce the precision is by spatial obfuscation, which has the advantage of requiring no trusted third party, since the user himself can run the obfuscation mechanism locally on his device. In this paper we focus on the *randomized* approaches to spatial obfuscation, which in recent years have become increasingly more popular. Much of this success is due to their properties of robustness with respect to attackers that may combine the observation of the user's activity with any side-information they have about the user or, more generally, about the habits of the population, the characteristics of the territory, etc.

The most influential proposals of this kind have been, to the best of our knowledge, those by Shokri et al. [5] and by Andrès et al. [6]. Both their frameworks are built on rigorous and natural notions of privacy, and both are based on the idea of confusing the adversary by reporting a noisy location, generated from the true location according to some probability distribution.

Specifically, the authors of [5] considered a Bayesian model of adversary. They focussed on optimizing the trade-off between privacy and utility and proposed a method to compute the optimal noise-generating mechanism, which consists in formalizing the utility constraints and the privacy target as a linear optimization problem, and then applying linear programming techniques. The approach of computing an optimal noisy mechanism via linear programming was also adopted in

Konstantinos Chatzikokolakis: CNRS, Inria and LIX, École Polytechnique, France

Ehab ElSalamouny: Inria, France and Faculty of Computer and Information Science, Suez Canal University, Egypt

Catuscia Palamidessi: Inria and LIX, École Polytechnique, France

[7] and [8], although they considered the reverse problem of optimizing the utility for a given level of privacy.

On the other hand, the authors of [6] proposed the notion of *geo-indistinguishability*, based on (an extension of) differential privacy [9]. Like differential privacy, geo-indistinguishability is independent from the prior of the adversary, and is robust with respect to composition. Furthermore, it can be implemented in a simple and efficient way using a planar Laplacian as the noise function. Indeed, probably thanks to its efficiency, geo-indistinguishability via the Laplacian mechanism has been adopted as the basis or as a component of several tools and frameworks for location privacy, including: Location Guard [10], LP-Guardian [11], LP-Doctor [12], the system for secure nearby-friends discovery in [13], SpatialVision QGIS plugin [14], and it is one of the possible input methods in STAC [15]. Furthermore, the PIM mechanism [16] can be considered an extension of the planar Laplacian to the case of traces (temporally correlated sequences of points): The authors of [16] attack the problem of the degradation of privacy due to correlation by adding Laplacian noise directly to the convex hull of the trace.

The Laplacian is not the only way to implement geo-indistinguishability: on a discrete map one can also use a planar variant of the geometric mechanism [17], and if the map is bounded, the tight-constraints mechanism [18] is applicable as well. Like the Laplacian, these mechanisms can be computed efficiently.

The advantages and disadvantages of these two approaches could be, at a first thought, resumed as follows: Generating the noise via a direct method (Laplacian, geometric, tight-constraints) is efficient, but there is no guarantee of optimality. Generating the noise via linear programming techniques, on the other hand, is computationally expensive, but it gives the optimal trade-off between privacy and utility.

In this paper we challenge this simplistic judgement, on the basis of the following consideration: The variables of the linear program are the probabilities of reporting location z when the true location is x . Hence, their number is the square of the number of locations in the problem. Consequently the linear programming technique can only be applied the number of locations is relatively small: In our experiments, we found that the optimization program would already take several days to produce the result when such number was approaching 100. This means that, if we want to apply the technique in an area of, for instance, 100 Km², we need to divide the map in “cells” of 1 Km² (for reference, 100 Km² corresponds approximatively to the area

of a city like Paris). On the other hand, the scale of typical utility functions is much smaller. For instance, if we intend to use an LBS to find point of interests in walking distance, it seems natural to measure the distance in terms of meters, not kilometers. Now, ignoring distances up to a kilometer in general induces a fictitiously higher utility, because when we discretize a map into cells, the distances are measured as if all the points in a cell were identified with its center. The utility computed using this discretized distance is only an approximation of the real utility, and the coarser is the granularity of the discretization, the less precise is this approximation. In order to make a fair comparison with the direct methods, where the granularity can be much finer, we should compute the utility on the same finer granularity. The experiments conducted in this paper show that for a city like Paris the real utility of the “optimal” methods is much smaller than what it would appear with the discretized distance.

Furthermore, a coarse granularity generates problems not only about utility, but also about privacy, because the privacy constraints usually depend on the distance between points as well. Assume, for instance, the privacy constraints are those of geo-indistinguishability, and consider two points x and x' that are very close but after the discretization belong to different cells. Since they are close they should be highly undistinguishable, but since they are in different cells they will be considered as being located in the centers of their respective cells, and they will therefore become distinguishable at the level allowed by the distance between the two centers.

On the other hand, it is possible to increase the utility of the direct methods (Laplacian, geometric, tight-constraints) by using one of the principles of the “optimal” methods. In general the construction of the optimal mechanism (for the case in which we fix the privacy constraints and we optimize the utility) is based on a Bayesian technique: given a prior, each reported location is remapped into the best possible location, according to the prior and the loss function. Such process can be applied also in the case of the direct methods: Given the prior, one can compute the best remapping with respect to the prior and apply it to the method, thus increasing the utility. While in the “optimal” methods to find the best remapping is part of the linear optimization problem and therefore it adds significantly to the complexity, in the case of the direct methods the best remapping is computed separately from the mechanism, and after we know the result of the mechanism. Therefore, it can be computed much more efficiently.

Contribution

The contributions of this paper are as follows:

- We consider the main methods for ge-indistinguishability (Laplacian, geometric, tight-constraints) on various kinds of domains: continuous, discrete unbounded, and discrete bounded. We show that their utility increases considerably when we apply the best remapping (w.r.t. a given prior). The comparison is done using the Brightkite and the Gowalla datasets.
- We compare the utility of the above methods and that of the “optimal” method generated by the linear programming approach. We show that if one takes into account the loss of precision due to the necessarily coarser granularity of the latter, then the former give much better results. The evaluation is done again using the Brightkite and the Gowalla datasets.
- In order to construct the prior, we take a common machine learning approach separating between the training and testing data. From the first part, we construct a global prior used to optimally remap the mechanisms. From the other part of the dataset, we construct a user-specific prior used to measure the utility of the mechanism.

Related work

The most closely-related works have already been discussed in the first part of the introduction. Here we mention some of the other relevant work in the area of location privacy.

As already stated, in general, all computational methods for privacy protection are based on degrading the precision of information. In the particular case of location data, this is obtained essentially in two ways: spatial cloaking and spatial obfuscation.

Spatial cloaking, first proposed in [19], is based on the idea of concealing the user’s exact coordinates by reporting a cloaked area, so to meet certain anonymity constraints. Often, the cloaking is not only spatial, but also temporal, so to conceal also the time in which the user was in that position. The anonymity constraints that have been mostly considered in the case of location privacy are: k -anonymity [20–23], l -diversity [24], t -closeness [25], and p -sensitivity [26]. In addition, in order to reduce the linkability between identity and trajectories, [27] proposed the so-called mix-zones. This idea assumes that people will only report their location

in certain regions, called “application zones”, where a location-based service is offered, e.g. an airport, bank, or coffee shop. In the mix zones, outside the application zones, users will receive new, unused pseudonyms. This helps prevent an attacker from linking pseudonyms, because the new pseudonym could have been assigned to anyone else in the mix zone.

In order to protect *aggregate* location information, randomized methods, and in particular differential privacy, have also been used. For instance, [28] presents a way to statistically simulate the location data from a database while providing privacy guarantees. In [29], a quad tree spatial decomposition technique is used to achieve differential privacy in a database with location pattern mining capabilities. On the other hand, Dewri [30] proposes a combination of differential privacy and k -anonymity for the purposes of hiding the location of a single individual.

Spatial obfuscation, on the other hand, works by reducing the precision of the position sent from the user to the server. In contrast to the above approaches, that require a trusted third party, spatial obfuscation mechanisms can be run by user itself [31]. Spatial obfuscation approaches can be classified into two main kinds: deterministic and randomized. We already discussed the randomized ones in the first part of this introduction, but we want to mention here also the approach of [32], which exploits the symmetry of circular noise functions to define a mechanism that is very easy to implement, that captures differential privacy and ge-indistinguishability, and that under certain conditions provides the same privacy and utility levels as more complex noise functions.

As for the deterministic methods, a typical one is that proposed by Ardagna et al. [33], which consists in blurring the user’s location by reporting a larger region (containing the position). Other similar proposals are presented in [34–37].

Unfortunately, deterministic methods are not robust with respect to composition. For instance, if the user reports one zone and then, soon afterwards, an adjacent zone, it can easily be inferred that he must be near the border between these two zones. Furthermore, these methods offer little guarantees against an adversary who already has some side knowledge (aka *prior knowledge*) about the habits of the user or of the population.

Plan of the paper

The next section recalls the necessary preliminary notions. Section 3 introduces the notion of remapping. Section 4 defines the various direct mechanisms and the result of applying the remapping to them. Section 5 compares the various mechanisms (with and without remapping), and the optimal one obtained from linear programming. Section 6 concludes.

2 Preliminaries

In this section we recall the notion of geo-indistinguishability from [6], and discuss the planar Laplace mechanism, as well as the techniques for constructing mechanisms of optimal utility of [7, 8].

2.1 Geo-indistinguishability

The notion of geo-indistinguishability is a variant of differential privacy for location based systems introduced in [6]. Let $\mathcal{X} \subseteq \mathbb{R}^2$ be the set of *possible locations* of a user, and let \mathcal{Z} be a set of reported values, often assumed to be equal to \mathcal{X} . Let $\mathcal{P}(\mathcal{Z})$ denote the set of probability distributions over \mathcal{Z} and let \mathbf{d}_{\otimes} be the *multiplicative distance* between two distributions $\sigma_1, \sigma_2 \in \mathcal{P}(\mathcal{Z})$, defined as $\mathbf{d}_{\otimes}(\sigma_1, \sigma_2) = \sup_{Z \subseteq \mathcal{Z}} \left| \ln \frac{\sigma_1(Z)}{\sigma_2(Z)} \right|$, with the convention that $\left| \ln \frac{\sigma_1(Z)}{\sigma_2(Z)} \right| = 0$ if both $\sigma_1(Z), \sigma_2(Z)$ are zero and ∞ if only one of them is zero.

An *obfuscation mechanism* K probabilistically selects a reported location $z \in \mathcal{Z}$ starting from a real location $x \in \mathcal{X}$, and is modelled by a function $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ mapping each real location to a probability distribution over the reported ones. K satisfies ϵ -geo-indistinguishability iff

$$\mathbf{d}_{\otimes}(K(x), K(x')) \leq \epsilon \mathbf{d}(x, x') \quad \forall x, x' \in \mathcal{X}$$

where \mathbf{d} denotes the Euclidean distance. This definition can be equivalently stated as $K(x)(Z) \leq e^{\epsilon \mathbf{d}(x, x')} K(x')(Z)$ for all $x, x' \in \mathcal{X}, Z \subseteq \mathcal{Z}$.

Note that $\epsilon \mathbf{d}$ can be viewed as a *distinguishability metric*: the closest x, x' are wrt this metric, the more similar the distributions $K(x), K(x')$ are required to be. Following [6], we choose ϵ by selecting a radius r in which we want to enjoy strong privacy (eg. $r = 0.1$ km), some small privacy level l for that radius (eg. $l = \ln(1.4)$) and then set $\epsilon = l/r$. This ensures that locations within distance r will have distinguishability level at most l .

Note also that, although the Euclidean distance is a natural choice of distinguishability metric for location privacy, the same definition could be used with a different distance metric capturing a different notion of privacy. Privacy under arbitrary metrics is studied in [38], while [39] proposes a method to construct a distinguishability metric taking into account the semantics of each location. In this paper we mostly assume \mathbf{d} to be Euclidean, but also discuss mechanisms that allow the use of an arbitrary privacy metric.

2.2 Utility

The utility provided by an obfuscation mechanism K depends on the specific application it is used for. Certain applications (eg. whether report) can tolerate a high amount of noise without affecting the provided service, while others (eg. GPS navigation) require finer information. Following several works on location privacy [5, 7], we define utility as the expected *quality loss* of K wrt some prior distribution π :¹

$$\text{QL}(K, \pi, \mathbf{d}_Q) = \sum_{x \in \mathcal{X}, z \in \mathcal{Z}} \pi(x) K(x)(z) \mathbf{d}_Q(x, z)$$

where \mathbf{d}_Q is a *quality loss metric*, measuring how much the service is degraded when z is reported instead of x . When evaluating a mechanism without a specific application at hand, we generally use the Euclidean metric for \mathbf{d}_Q , since the quality of all LBSs degrades when z moves far away from x .

2.3 The planar Laplace mechanism

The planar Laplace (PL) mechanism is a simple and efficient continuous mechanism satisfying ϵ -geo-indistinguishability. It can be used when \mathcal{X} is any subset of \mathbb{R}^2 and the reported set \mathcal{Z} is the full \mathbb{R}^2 . The mechanism consists at drawing from a 2-dimensional Laplace distribution centered at the real location x , having pdf:

$$D_x(z) = \frac{\epsilon^2}{2\pi} e^{-\epsilon \mathbf{d}(x, z)}$$

Drawing from this distribution can be easily and efficiently performed in polar coordinates, by adding to x a randomly drawn vector expressed as a radius r and angle θ , as follows:

¹ For simplicity, the expression for QL assumes discrete \mathcal{X}, \mathcal{Z} .

- draw θ uniformly in $[0, 2\pi)$,
- draw p uniformly in $[0, 1)$ and set $r = C^{-1}(p)$,
- report $z = x + \langle r \cos(\theta), r \sin(\theta) \rangle$,

where $C^{-1}(p)$ is the inverse of PL’s cumulative distribution function for r , given by $C^{-1}(p) = -\frac{1}{\epsilon} (W_{-1}(\frac{p-1}{\epsilon}) + 1)$ and W_{-1} is the Lambert W function (-1 branch) implemented in several numerical libraries.

Note that, as with any continuous mechanism, using the above algorithm in a machine with floating point arithmetic essentially corresponds to using a discretized version of the mechanism, where the obtained location z corresponds to a subset of \mathbb{R}^2 around z . In that sense we can think of z as a “single point” drawn from a continuous distribution.

Concerning utility, it can be shown that, due to the symmetry of \mathbb{R}^2 , the Euclidean quality loss of PL is independent from π , namely:

$$\text{QL}(\text{PL}, \pi, \mathbf{d}) = 2/\epsilon \quad \forall \pi \in \mathcal{P}(\mathcal{X}) \quad (1)$$

2.4 Constructing optimal mechanisms

Although the planar Laplace mechanism is simple and efficient, its utility is not guaranteed to be optimal among all mechanisms satisfying ϵ -geo-indistinguishability. If \mathcal{X} is finite and small in size, given π and \mathbf{d}_Q an optimal mechanism $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{X})$ can be constructed by solving a linear optimization problem [7]. In the finite case the mechanism can be represented by a $|\mathcal{X}| \times |\mathcal{X}|$ stochastic matrix K , where K_{xz} is the probability to report z from x . Having elements of K as variables, the optimal mechanism is given as a solution to:

$$\begin{aligned} & \text{minimize } \text{QL}(K, \pi, \mathbf{d}_Q) \\ & \text{subject to } K_{xz} \leq e^{\mathbf{d}(x, x')} K_{x'z} \quad x, x', z \in \mathcal{X} \\ & \text{and } \sum_z K_{xz} = 1 \quad x \in \mathcal{X} \end{aligned}$$

The constructed mechanism is guaranteed to satisfy ϵ -geo-indistinguishability, and have no worse utility than any other mechanism K' also satisfying the same privacy definition. However, solving the above problem is feasible only when $|\mathcal{X}|$ is very small. In [7] an approximation technique is proposed using spanners, but even in this case the technique is limited to at most a few hundred locations.

Due to this scalability issue, it might seem a natural idea to employ this technique on a coarse grid, consisting of only a few large cells, and use the resulting mechanism as if all users are located at the center of these cells.

This solution, however, has important privacy and utility consequences, discussed in detail in Section 4.3.1.

Note that this technique can be employed to an arbitrary privacy metric \mathbf{d} , not necessarily the Euclidean one. Moreover, [8] proposes similar constructions taking also into account inference error constraints.

3 Remapping locations

Remapping the output of a mechanism is a simple yet effective technique for improving its utility while satisfying the same level of privacy. In this section we discuss the technique in general, the optimal way of performing it and an efficient way to remap points on a continuous plane. In Section 4 we further discuss remapping for various types of mechanisms, and in Section 5 we show that remapping can be effective in practice for substantially improving the mechanism’s utility.

Given a mechanism $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$, and starting from a location x , we can first use K to draw a noisy location z , then remap z to a new location $z^* = R(z)$ using a remapping function² $R : \mathcal{Z} \rightarrow \mathcal{Z}$, and finally report z^* . Let KR denote the composed mechanism (if K, R are written as stochastic matrices, the composition is their product KR). It is well known that if K satisfies ϵ -geo-indistinguishability so does KR , hence remapping is safe from the point of view of privacy.

On the other hand, remapping can improve the mechanism’s utility. Geo-indistinguishability requires nearby locations x, x' to produce z with the same probability, but z does not have to be far away from x . Consider, for instance, two locations close to the sea, say in San Francisco (Fig 6). A mechanism such as the planar Laplace adds noise uniformly in all directions, leading to a high chance of z falling into the sea. Knowing that the user is unlikely to be in the sea, we can remap it back into the edge of the city, bringing it closer to the true location. Of course, such a remap will decrease the utility if a user is indeed in the sea, but this only happens with small probability; on average utility is improved.

Clearly, one wishes to employ the remap providing the best utility, captured by the following optimality property.

² In general R could also be probabilistic, although in this paper we only consider deterministic ones, since an optimal deterministic remap always exists.

Definition 3.1 (Optimal remap). *A remap R is optimal for K wrt π, \mathbf{d}_Q iff $\text{QL}(KR, \pi, \mathbf{d}_Q) \leq \text{QL}(KR', \pi, \mathbf{d}_Q)$ for all remaps R' .*

Optimality can be achieved by a *Bayesian* remap, choosing the location z^* that minimizes the expected loss wrt the *posterior* distribution $\sigma \in \mathcal{P}(\mathcal{X})$ obtained from π, K after observing z by applying Bayes' rule. Let $\sigma = \mathbf{Post}(\pi, K, z)$ denote this distribution, given by

$$\sigma(x) = \frac{\pi(x)K(x)(z)}{\sum_{x'} \pi(x')K(x')(z)} \quad x \in \mathcal{X}$$

Proposition 3.2 (Bayesian remap). *Given π, K, \mathbf{d}_Q , the remap defined by:*

$$R(z) = \arg \min_{z^* \in \mathcal{Z}} \sum_{x \in \mathcal{X}} \sigma(x) \mathbf{d}_Q(x, z^*) \quad \text{where}$$

$$\sigma = \mathbf{Post}(\pi, K, z)$$

is optimal for K wrt π, \mathbf{d}_Q .

3.1 Efficiency of the Bayesian remap

The Bayesian remap consists of computing the posterior $\mathbf{Post}(\pi, K, z)$ and then finding the point $z^* \in \mathcal{Z}$ that minimizes the expected loss. If \mathcal{X}, \mathcal{Z} are finite, the remap can be computed directly using its definition (Prop 3.2). Note that computing the remap for an existing mechanism K is far more efficient than constructing an optimal mechanism (Section 2.4); the latter is feasible only for sizes of the order of 10^2 , while a direct computation of R can be performed for sizes two to three orders of magnitude larger.

Moreover, even when iterating over all $z^* \in \mathcal{Z}$ is infeasible, the remap can be approximated by considering only locations within a certain \mathbf{d}_Q -ball $B_r(z)$ centered at z . The posterior σ can be assumed to be 0 outside the ball and $\arg \min_{z^* \in \mathcal{Z}}$ can be replaced by $\arg \min_{z^* \in B_r(z)}$. Intuitively, when moving away from z the posterior σ (for realistic mechanisms) decreases while the loss \mathbf{d}_Q increases, hence the optimal z^* is not expected to be too far away from z . The complexity in this case depends only on r , not on $|\mathcal{X}|, |\mathcal{Z}|$ and this technique can give good results in practice.

Note that any remap function preserves geo-indistinguishability, so from the point of view of privacy we are safe to use any heuristic might make it faster.

On the other hand, if \mathcal{Z} is continuous the problem becomes harder, since, even at a restricted radius r , we would still need to consider uncountably many points. This case is discussed in the next section.

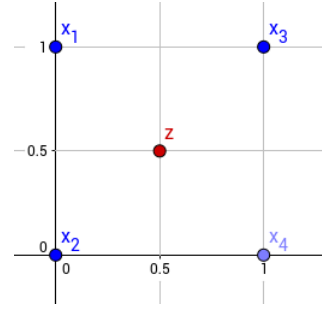


Fig. 1. The uniformly distributed corners of a unit square should be remapped to the center.

3.2 Efficient remap on the continuous plane

For many applications, we would like to have the flexibility to position a user anywhere on a plane, as well as to report arbitrary locations, hence it is convenient to take \mathcal{X}, \mathcal{Z} to be the full \mathbb{R}^2 .

The first problem in this case is computing the posterior. Although users could be potentially located anywhere in the plane, in practice the prior is constructed from some finite dataset of past service uses, points of interests, etc. Hence, we can assume that a posterior $\sigma \in \mathcal{P}(\mathbb{R}^2)$ of *finite support*, can be constructed for the remap (see Section 4.1.1 for a possible way of constructing σ).

Now, assuming a posterior σ of finite support, denoted by $[\sigma]$, the problem is that we still have an uncountable set of possible locations for z^* . The reason is that, although $[\sigma]$ is finite, the z^* minimizing the expected loss is not necessarily an element of the support. Consider, for instance, a posterior σ assigning uniform probability to the four corners of the unit square, displayed in Fig 1. Selecting any of those points for z^* gives an expected Euclidean loss of $\frac{2+\sqrt{2}}{4} \approx 0.85$, while the center of the square gives an expected loss of $\frac{\sqrt{2}}{2} \approx 0.71$.

Thankfully, for the common case of the Euclidean loss metric \mathbf{d} , constructing the optimal z^* can be done geometrically. Intuitively, z^* should be “in between” the points in the support of σ , that is in the convex hull of $[\sigma]$. Moreover, to minimize the expected loss, z^* should be closer to points having large probability. Hence, we can construct the optimal z^* as a convex combination of elements of $[\sigma]$ (note that such elements are vectors), using the probability of each point as the coefficient, a construction which can be thought of as the *mean* of σ .

Proposition 3.3 (Continuous remap). *Let $\sigma \in \mathcal{P}(\mathbb{R}^2)$ be a distribution with finite support and let*

$$\bar{\sigma} = \sum_{x \in [\sigma]} \sigma(x) x$$

be the mean of σ . Then:

$$\bar{\sigma} = \arg \min_{z^* \in \mathbb{R}^2} \sum_{x \in [\sigma]} \sigma(x) \mathbf{d}(x, z^*)$$

Note that construing $\bar{\sigma}$ is practical even for large sizes (linear on the size of $[\sigma]$). In the example of Fig 1 $\bar{\sigma}$ is exactly the center of the square. Section 4.1.1 discusses this technique in the context of the planar Laplace mechanism, and the evaluation of Section 5 shows its practical applicability.

3.3 The choice of prior

A crucial element of the Bayesian remap is the prior distribution π from which the posterior σ is constructed. In theory, the prior should describe the behaviour of the user, and a good quality prior, constructed from a large dataset of past behaviour of that user, is assumed to be available. In practice, however, there are two major problems with this assumption.

First, general-purpose obfuscation mechanisms should be able to work with *new users* for which *no information* whatsoever is available. For instance, a user installing the Location Guard browser plugin [10] is expecting to use it immediately, but the plugin has no information about this user.

Second, even after an extended training period and assuming that all user’s movements are recorded, the prior constructed from this data is not necessarily of good quality. For instance, such a prior would assign zero probability to all locations never visited in the past, but clearly the user might visit some new locations in the future. Moreover, such a prior would mostly contain locations that the user visits on a daily basis, such as its home or work. Remapping using such a prior is similar to revealing the prior to the provider: we might often assume that an adversary already has this prior knowledge, but if he doesn’t we shouldn’t reveal it to him ourselves.

The problem of prior knowledge is particularly crucial for the *proper evaluation* of a mechanism. Constructing a prior from a dataset and then evaluating it on the *same* dataset (a common practice in the literature of location privacy) might be misleading, especially for datasets containing only a few locations per user. In the extreme case, a prior constructed from a single visited location will lead to a remap that is perfect for *that*

location, but likely very bad for *any other* location the user might visit in the future. Such a practice is similar to training and testing a learning algorithm on the same dataset.

To deal with the above issues, we assume the remap is performed using a *generic prior* π , constructed from a large dataset containing a large number of different users, with the intention of describing the behaviour of an *average user*. In particular, in the evaluation of Section 5 we construct the prior from a training dataset containing *no data whatsoever* of the users of the testing set. This simulates a remapping applied to a new user seen for the first time, based on generic data about past uses of the specific service.

4 Practical Mechanisms

In this section we discuss several mechanisms satisfying geo-indistinguishability that can be applied to a generic LBS. Our emphasis is on practical applicability: the mechanisms should be efficient and applicable to realistic domain sizes, without any unreasonable assumptions.

We classify the mechanisms based on the restrictions they assume on the structure of the domains \mathcal{X}, \mathcal{Z} . In the more flexible case, the domains can be the full continuous plane. However, in some applications it might be reasonable to restrict to a discrete, either infinite or a finite but realistically large domain. Mechanisms for each category can be clearly applied to the more restrictive cases, but the extra restrictions allow for the construction of additional mechanisms with certain advantages.

4.1 Mechanisms for the continuous plane

In the general unrestricted case of the full continuous plane, the natural choice is to apply the planar Laplace mechanism described in Section 2.3. The PL mechanism is simple and efficient, and due to its generality and lack of restrictions, it is used in all known applications of geo-indistinguishability. However, even on a continuous plane, the remapping technique can be practically employed to improve the mechanism’s utility, as discussed in the next section.

4.1.1 Remapping the planar Laplace mechanism

The challenge for remapping the PL mechanism lies in the continuity of its domain. Using a continuous prior on \mathbb{R}^2 is not realistic, it's hard to construct such a prior, let alone to use it. Instead, we assume that prior information about the service is given in the form of a dataset of previous uses of the service, or even in the form of generic information about points of interest (POIs). Let $Q \subset \mathbb{R}^2$ be a possibly large but finite set of locations, each $q \in Q$ associated with a weight $w(q) \geq 0$. For instance Q might contain the locations from which the service has been accessed in the past, with $w(q)$ being the number of users accessing the service from that location. Or Q might be a list of POI relevant for the service, with $w(q)$ capturing the popularity of each POI.

Given a noisy location z drawn from PL, the goal is to first construct a reasonable posterior $\sigma \in \mathcal{P}(\mathbb{R}^2)$ of finite support. To do so, we need to restrict Q (which can be very large) to a limited area around z . Let $r = C^{-1}(0.99)$, where C^{-1} is the inverse of PL's cumulative distribution function for the radius (see Section 2.3), and let $Q_r = Q \cap B_r(z)$. Intuitively, PL reports a point within distance r from x with probability 99%, hence it is reasonable to remap z to a point z^* no farther than r away from z .

Note that Q_r can be computed efficiently using a spatial data-structure such as kd-trees. The 0.99 threshold can be used to trade efficiency for accuracy, smaller values will lead to a smaller Q_r but might cause the optimal point to be outside the area. If Q is dense, a pre-processing phase could be used to reduce the size of Q_r , merging points into small clusters, and setting $w(q)$ to the weight of the cluster.

The weights $w(q), q \in Q_r$ can be thought of (after normalization) as a finite prior on \mathbb{R}^2 . Applying Bayes' law to this prior, using the pdf of PL, we can construct the following posterior σ with finite support $[\sigma] = Q_r$:

$$\sigma(x) = \frac{w(x)e^{-\epsilon \mathbf{d}(x,z)}}{\sum_{q \in Q_r} w(q)e^{-\epsilon \mathbf{d}(q,z)}} \quad x \in Q_r$$

Finally, the remapped point z^* can be constructed as the mean of σ , using Prop 3.3. Note that z^* is not necessarily an element of Q_r , but always lies in the convex hull of Q_r .

Note that, in practice, the dataset Q might not be detailed enough to provide sufficient information for every location. A new user might access the service from a location in which no or few other users have been in the past. A remap using low quality data is likely to decrease the user's privacy, hence we use a simple heuristic for

assessing the data's quality: if the size (or alternatively the total weight) of Q_r is below a certain threshold t , we skip the remap and report z directly. The complete algorithm is displayed in Fig 1.

Algorithm 1: Planar Laplace with remap

Data: $x \in \mathbb{R}^2, \epsilon > 0, t \geq 0, Q \subset \mathbb{R}^2, w : Q \rightarrow \mathbb{R}^+$
Result: An obfuscated location z^*

- 1 draw $z \sim PL_\epsilon(x)$;
- 2 set $r = C^{-1}(0.99)$;
- 3 compute $Q_r = Q \cap B_r(z)$;
- 4 **if** $|Q_r| < t$ **then**
- 5 **return** z
- 6 **else**
- 7 compute

$$\sigma(x) = \frac{w(x)e^{-\epsilon \mathbf{d}(x,z)}}{\sum_{q \in Q_r} w(q)e^{-\epsilon \mathbf{d}(q,z)}} \quad x \in Q_r ;$$
- 8 compute $z^* = \sum_{x \in Q_r} \sigma(x) x$;
- 9 **return** z^*

4.2 Mechanisms for the discrete plane

For some applications, it is reasonable to assume that users are located on a discrete grid, and reported points should always be on that grid. However, we might still want to cover very large areas (or the whole world), hence it is convenient to think of this grid as being infinite. The PL mechanism can of course be employed in this case, with or without remap, by further projecting its output on the grid. However, the discrete nature of the domain allows us to use the Planar Geometric mechanism, a discretized version of PL.

4.2.1 The planar geometric mechanism

The planar geometric mechanism can be seen as a discretized version of the planar Laplace mechanism described in Section 2.3 (similarly to the relationship between the one-dimensional geometric and Laplace mechanisms). The two mechanisms share the idea of producing the output by adding random noise to the user's location x such that the added noise is independent of the real location of the user.

Precisely, let the Euclidean space \mathbb{R}^2 be partitioned into a grid of square cells such that the side length of every cell is $d > 0$. The centers of these cells form there-

fore an infinite grid of points \mathcal{G} , in which the horizontal (or vertical) distance between any successive two points is exactly d . We call this distance the ‘spacing’ of \mathcal{G} . Then we define the planar geometric mechanism PG as a probabilistic function from \mathcal{G} to itself as follows:

$$\text{PG}(x)(z) = \lambda e^{-\epsilon \mathbf{d}(x,z)} \quad x, z \in \mathcal{G}, \quad (2)$$

where λ is a normalization constant given by

$$\lambda = 1 / \sum_{(i,j) \in \mathbb{Z}^2} e^{-\epsilon d \sqrt{i^2 + j^2}}.$$

In general, every output location z of a mechanism can be seen as the sum of the input location x and an Euclidean ‘noise’ vector u having the magnitude $\mathbf{d}(x, z)$, and describing the added obfuscation noise. From this perspective, it is easy to see from (2) that the probability assigned by PG to every noise vector depends only on its magnitude regardless of the input of the mechanism. The normalization constant λ ensures that such probabilities sum up to 1, forming a valid distribution on the vectors drawn from \mathcal{G} . It is also easy to see that the planar geometric mechanism satisfies ϵ -geo-indistinguishability on \mathcal{G} .

Sampling from a discrete mechanism, such as PG, can be performed efficiently since we have an explicit construction of its probability mass function. First, λ can be approximated to any precision, which only needs to be done once for every choice of ϵ, d . Then, to draw a random location z , we uniformly select $p \in [0, 1]$ and we iterate \mathcal{G} , starting from x , by visiting locations in increasing order of distance from x . Once the cumulative probability of all visited points reaches p , we report the last visited location.

Finally, remapping can be applied to PG in two ways: one approach is to employ the continuous remap described in the previous section, and project the result to the grid. An alternative approach, is to directly use the definition of the Bayesian remap, but restrict to a radius around z , as discussed in Section 3.1. The second approach has the advantage of being applicable for any choice of \mathbf{d}_Q , but is less efficient, especially if the spacing d is small.

4.3 Mechanisms for large but finite grids

The third case that we consider is that of an application within a limited region, for which users are assumed to lie on a finite grid. Still, we would like to cover a large area, and more importantly to use a fine cell size, hence the grid size should be realistically large.

Clearly, the PL and PG mechanisms, with or without remap, can still be employed in this case, by projecting (truncating) their output to the finite grid. Note that the truncation step is an instance of post-processing, since it is independent from the user’s real location x , hence it preserves geo-indistinguishability.

On the other hand, the finite nature of the domain allows us to construct alternative mechanisms, that we discuss in the following sections.

4.3.1 Optimal mechanism built on a coarse grid

Since our domain is finite, we could apply the construction of 2.4 to obtain an optimal mechanism for a certain prior π . However, if \mathcal{X} contains more than a few hundred locations, as it is common in realistic scenarios, this construction is not feasibly applicable.

A natural solution would be to construct the mechanism on a *coarser* grid \mathcal{X}' containing larger cells obtained by merging together several cells of \mathcal{X} . Once K is constructed from \mathcal{X}' , a user located at $x \in \mathcal{X}$ would draw a noisy location from $K(x')$, where x' is the cell of the coarse grid closer to x . This technique is essentially *implied* by the use of very coarse grids in the evaluation of the mechanisms in [5, 7, 8]: clearly, users cannot be assumed to always lie on the handful of points contained in these grids.

Although natural, this change of granularity has important *utility and privacy consequences*. First, from the point of view of utility, the obtained mechanism, although optimal for the coarse grid, is by no means guaranteed to be optimal for the fine one. Optimality is based on the assumption that users are located at the centers of the cells; the evaluation of Section 5 shows that when the users are placed on a fined grid, the ‘optimal’ mechanism actually performs very poorly.

Second, the change of granularity affects privacy in the geo-indistinguishability model. The optimal construction guarantees that for two locations y, y' in the coarse grid, $K(y), K(y')$ are as similar as the distance $\mathbf{d}(y, y')$ between them. If, however, the mechanism is used by users located at points x, x' of the fine grid (by projecting to the coarse one), the privacy they enjoy is as large as the *distance between the projected points*. Two points x, x' could be projected to the same coarse cell, in which case their effective distance is 0; however, two close points at a border of a coarse cell could be projected to different cells, in which case their effective distance will be much larger than $\mathbf{d}(x, x')$.

At the extreme case, the coarse grid could be so big that distances between neighbouring cells are very large, making the corresponding constraints *too weak*. The resulting mechanism will be close to an identity one: each coarse cell will report itself with probability close to 1. Using such a mechanism, however, on a finer grid is effectively similar to spacial cloaking: points within the same coarse cell will be indistinguishable, while points of different coarse cells will be completely distinguishable. This privacy guarantee is quite different that geo-indistinguishability.

Of course, this phenomenon happens for grids of all sizes. However, if the grid is sufficiently fine (eg. cell size of 100 or 200 meters), the problem is negligible.

4.3.2 Tight-constraints mechanisms

The tight-constraints mechanism was introduced by [18] in the general setting of \mathbf{d} -privacy for any finite space \mathcal{X} of secrets equipped with a privacy metric \mathbf{d} (not necessarily the Euclidean one). The main objective of this mechanism is to optimize the utility for the user by maximizing the probability of reporting the true value while preserving the required level of privacy. The tight-constraints mechanism, denoted by TC, is defined as follows.

Definition 4.1 (Tight-constraints mechanism [18]). *Given a finite set \mathcal{X} and $\epsilon > 0$, $TC : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{X})$ is a mechanism satisfying:*

$$TC(x)(z) = e^{-\epsilon \mathbf{d}(x,z)} TC(z)(z) \quad \forall x, z \in \mathcal{X}.$$

The above definition relates the probability of reporting an output z , when the real location is z itself, with the same probability from some other location x . This relation means precisely that the ϵ -geo-indistinguishability constraint for the inputs x, z and the output z is satisfied with equality. With regard to this characteristic, the mechanism is named after those ‘tight’ constraints. It is shown in [18] that TC satisfies ϵ -geo-indistinguishability; furthermore, with respect to the *binary loss* function – defined as $\mathbf{d}_Q(x, z) = 0$ iff $x = z$ and 1 otherwise – this mechanism is optimal for a set of priors called ‘regular’ priors. While this optimality is guaranteed only for the binary loss function, we experimentally show in Section 5 that TC provides also a substantial improvement in comparison to other mechanisms when the loss is measured as the Euclidean distance between the input and output locations. In par-

ticular, TC enjoys substantially better utility than the Exponential mechanism, the only other (efficient) mechanism applicable to an arbitrary privacy metric \mathbf{d} .

Existence and construction

It is important to remark that the mechanism TC may not exist in some cases, more precisely when there are no collection of conditional probability distributions (one for every input) that satisfy the tight constraints of Def 4.1. It is shown by the authors of [18] that the necessary and sufficient condition of the existence of this mechanism is related to the privacy-constraints matrix Φ with entries indexed by the elements of \mathcal{X} as follows.

$$\phi_{xz} = e^{-\epsilon \mathbf{d}(x,z)} \quad \forall x, z \in \mathcal{X}.$$

Then the mechanism TC exists if and only if there is a vector $\boldsymbol{\mu}$ indexed by \mathcal{X} such that

$$\Phi \boldsymbol{\mu} = \mathbf{1} \quad \text{and} \quad \mu_z \geq 0 \quad \forall z \in \mathcal{X}. \quad (3)$$

Every entry μ_z of the vector $\boldsymbol{\mu}$ is precisely the probability of reporting the location z when the user is at z , i.e. $\mu_z = TC(z)(z)$ for every $z \in \mathcal{X}$. Therefore using the entries of this vector and Definition 4.1, all the conditional probabilities of the mechanism TC are easily obtained by the following simple equation.

$$TC(x)(z) = \phi_{xz} \mu_z \quad \forall x, z \in \mathcal{X}. \quad (4)$$

When the domain of possible locations of the users is regarded as a discrete set of points, e.g. a grid, we show in this paper that TC provides significantly less expected loss in comparison to other well-known mechanisms, e.g. the planar Laplace mechanism. However, this advantage comes at the computational cost required to obtain $\boldsymbol{\mu}$ by solving the system of equations in (3). This system consists of $|\mathcal{X}|$ linear equations, making the cost of solving it depending on $|\mathcal{X}|$. In the following we exploit the graph symmetries between the points of \mathcal{X} to reduce this cost of constructing TC.

Reducing the construction cost using symmetries

The domain of locations \mathcal{X} together with the distances between its individual points can be regarded as a weighted graph in which the vertices are the points of \mathcal{X} , and the weight for every two vertices x, z is exactly the Euclidean distance $\mathbf{d}(x, z)$. A bijective mapping $\rho : \mathcal{X} \rightarrow \mathcal{X}$ is called an automorphism on \mathcal{X} if this mapping preserves the distance between every $x, z \in \mathcal{X}$ in the sense $\mathbf{d}(\rho(x), \rho(z)) = \mathbf{d}(x, z)$. Using the notion of automorphisms, we define symmetric points as follows.

15	14	13	12	11	11	12	13	14	15
14	10	9	8	7	7	8	9	10	14
13	9	6	5	4	4	5	6	9	13
12	8	5	3	2	2	3	5	8	12
11	7	4	2	1	1	2	4	7	11
11	7	4	2	1	1	2	4	7	11
12	8	5	3	2	2	3	5	8	12
13	9	6	5	4	4	5	6	9	13
14	10	9	8	7	7	8	9	10	14
15	14	13	12	11	11	12	13	14	15

Fig. 2. Symmetric cells in a grid of size 10×10 . The cells having the same label are symmetric to each other, i.e. in the same equivalence class. The highlighted region contains one cell of every equivalence class, and shows that the grid has 15 class.

Definition 4.2 (symmetric points). *Consider a domain of points \mathcal{X} . For any two points $x, x' \in \mathcal{X}$, we say that x is symmetric to x' , written as $x \approx x'$ if there is an automorphism $\rho : \mathcal{X} \rightarrow \mathcal{X}$ such that $\rho(x) = x'$.*

It is clear that \approx is an equivalence relation, and therefore it partitions the domain \mathcal{X} into a set of equivalence classes written as \mathcal{X}/\approx . Figure 2 illustrates a grid \mathcal{X} of 10×10 cells (locations), in which symmetric cells are annotated by the same label. In other words, the label on every cell refers to the equivalence class that the cell belongs to.

With respect to the equivalence classes of \mathcal{X} , we define a matrix Φ^\approx of which the rows and columns are indexed by elements (classes) of \mathcal{X}/\approx as follows. For every $c, c' \in \mathcal{X}/\approx$,

$$\phi_{cc'}^\approx = \sum_{z \in c'} \phi_{xz} \quad \text{where } x \text{ is any member of } c. \quad (5)$$

Then we show by the following theorem that instead of constructing the tight-constraints mechanism TC by solving the original equations (3) which involves the large matrix Φ , this mechanism can be constructed more efficiently using the smaller matrix Φ^\approx .

Theorem 4.3 (tight-constraints mechanism). *A tight-constraints mechanism TC exists for a domain of locations \mathcal{X} if and only if there is vector μ^\approx such that*

$$\Phi^\approx \mu^\approx = \mathbf{1} \quad \text{and} \quad \mu_c^\approx \geq 0 \quad \forall c \in \mathcal{X}/\approx. \quad (6)$$

Furthermore, TC is obtained as follows. For every $c \in \mathcal{X}/\approx$, it holds

$$\text{TC}(x)(z) = e^{-\epsilon \mathbf{d}(x,y)} \mu_c^\approx \quad \forall x \in \mathcal{X}, \forall z \in c. \quad (7)$$

Proof. In the following we will denote the equivalence class of every $x \in \mathcal{X}$ by c_x . Suppose that there is a vector μ^\approx that satisfies (6). Then it holds for every $x \in \mathcal{X}$ that $\sum_{c \in \mathcal{X}/\approx} \phi_{c_x c}^\approx \mu_c^\approx = 1$. The latter equation can be expanded using the definition (5) of $\phi_{c_x c}^\approx$ as

$$\sum_{c \in \mathcal{X}/\approx} \sum_{z \in c} \phi_{xz} \mu_c^\approx = 1.$$

From the above equation, it is clear that (3) is satisfied with the vector μ defined as $\mu_z = \mu_c^\approx$ for all $c \in \mathcal{X}/\approx$ and $z \in c$. Since also $\mu_z \geq 0$, the tight-constraints mechanism exists. In this case Eq. (4), which evaluates the probabilities of TC, yields Eq. (7) since $\phi_{xz} = e^{-\epsilon \mathbf{d}(x,y)}$ and $\mu_z = \mu_c^\approx$.

Conversely, suppose that the tight-constraints mechanism exists, i.e. there is a vector μ that satisfies the system of equations (3). We show in the following that there is vector μ^\approx that satisfies the reduced system (6). Let $\{\rho_i : i = 1, 2, \dots, n\}$ be the set of all different automorphisms on \mathcal{X} . Now consider any class $c \in \mathcal{X}/\approx$, and any member $x \in c$. Then by (3), it holds

$$\sum_{z \in \mathcal{X}} \phi_{\rho_i(x)\rho_i(z)} \mu_{\rho_i(z)} = 1, \quad 1 \leq i \leq n.$$

The above equations are exactly the instances of (3) that correspond to the locations $\rho_i(x)$ for $1 \leq i \leq n$, i.e. the members of c . Note also in the above summation, $\rho_i(z)$ considers every element of \mathcal{X} because y iterates on all these elements and ρ_i is bijective. Using the fact that $\phi_{\rho_i(x)\rho_i(z)} = \phi_{xz}$, we combine the above equations into one as follows.

$$\sum_{c' \in \mathcal{X}/\approx} \sum_{z \in c'} \phi_{xz} \left(\sum_{i=1}^n \mu_{\rho_i(z)} \right) / n = 1.$$

In the above equation, note that the quantity $(\sum_{i=1}^n \mu_{\rho_i(z)})/n$ is non-negative and depends only the class c' of y . Denoting this quantity by $\mu_{c'}^\approx$, and recalling that $\sum_{z \in c'} \phi_{xz} = \phi_{cc'}^\approx$, the above equation can be rewritten as follows.

$$\sum_{c' \in \mathcal{X}/\approx} \phi_{cc'}^\approx \mu_{c'}^\approx = 1.$$

Since c was chosen arbitrarily, the above equation holds for every member of \mathcal{X}/\approx . Since also $\mu_{c'}^\approx \geq 0$ for all c' , the system of equations in (6) is satisfied. \square

Using the above theorem, the cost of constructing the tight-constraints mechanism for a domain of locations \mathcal{X} is dramatically reduced by exploiting the symmetries between its locations. For example, in Figure 2, the grid

of 100 locations contains only 15 equivalence classes. In general for any grid of size $n \times n$, the number of classes is $n^2/8 + n/4$ if n is even and $(n+1)^2/8 + (n+1)/4$ if n is odd. Clearly this makes solving Equations (6) significantly cheaper than solving the original equations (3).

4.4 The exponential mechanism

The exponential mechanism Exp is a generic mechanism applicable to any finite domain \mathcal{X} with an arbitrary privacy metric \mathbf{d} . It is given by:

$$\text{Exp}(x)(z) = \lambda_x e^{\frac{1}{2}\epsilon \mathbf{d}(x,z)}$$

where

$$\lambda_x = 1 / \sum_{z \in \mathcal{Z}} e^{\frac{1}{2}\epsilon \mathbf{d}(x,z)}$$

is a normalization factor that can be different for every x . Compared to the other mechanisms, the $\frac{1}{2}$ factor in the exponent compensates for the fact that the coefficient λ_x depends on the secret x . The smaller exponent leads to a greater variance of the noise, hence the utility of this mechanism is the worse among those discussed in this section, with the advantage, on the other hand, of being very simple and at the same time applicable to any metric \mathbf{d} . The exponential mechanism is used in [39] to achieve privacy wrt a constructed “elastic” metric, adapted to the semantics of each location.

5 Evaluation

In this section, we evaluate all mechanisms discussed above on two real-world datasets from the Gowalla and Brightkite social networks. For the infinite continuous and discrete case, we perform an evaluation on the *complete datasets*, showing that an efficient remapping can be performed without any restriction on the area of interest. For the finite case, we consider a large geographical region (12km \times 28km) covering most of the San Francisco peninsula as illustrated in Figure 6. This region is bounded from south and north by the latitudes 37.5395, 37.7910, and from west and east by longitudes -122.5153, -122.3789.

To ensure the practical applicability of the proposed mechanisms, we take a common machine learning approach separating between the training and testing (evaluation) data. More precisely, we split the entire location dataset into two non-overlapping parts. The first part contains the location data of approximately 80% of

the users and is seen as training data. From this part, we construct a global prior computed as the average of individual priors of all users visiting the region. This global prior is then used to optimally remap the proposed mechanisms. Note that this remap is optimized for the global prior rather than being overfitted to a specific user, but yet provides a significant improvement of the utility for most of the users.

The other part of the dataset, which contains the data of the remaining 20% of the users is seen as the testing data for evaluating the constructed mechanisms. More precisely, we construct a user-specific prior for every user having at least 20 check-ins (in total for the infinite case, or in the specified region for the bounded case), and measure the expected loss of the mechanism for the user using her own prior. Although the mechanisms are trained only on the users of the training part of the dataset, we find that they offer a low level of quality loss also for the users in the testing set.

We should emphasize that the split is performed on users (not on checkins): *no data whatsoever* for the testing users is available in the training dataset. This very conservative approach aims at simulating the case of a new user for which we have no information, other than what is generally available about the whole service. Improving the utility in this case is a strong result; clearly, if we trained on part of the user’s own data, the results could be further substantially improved.

For the continuous case, we evaluate the planar Laplace mechanism, by itself, or coupled with the continuous remapping technique of Section 3.2. The dataset Q is the whole training set consisting of 48K users and more than 5M checkins. From the list of checkins we create a kd-tree, allowing to quickly construct Q_r . In our evaluation, each remap took only a few milliseconds, despite the large dataset. The two mechanisms were evaluated in the whole training dataset of 12000 users.

The results are shown in Fig 3, where boxplots of the average loss for each user are given for various values of l . We use $\epsilon = l/r$ with $r = 0.1$ km and l ranging from $\ln(1.4)$ to $\ln(2.9)$. Note that the average loss for PL without remap is always the same, equal to $2/\epsilon$, hence the boxplots reduce to a single line. On the other hand, the average loss when remapping is employed depends on the user: since remapping is performed using a global prior, it might not always give an improvement. Still, although no user data is used for training, the results show a substantial improvement in terms of expected loss for most users. In Brightkite we get similar results, although the variance of the error is visibly higher. This




Fig. 3. Results for the continuous case




Fig. 4. Results for the discrete infinite case

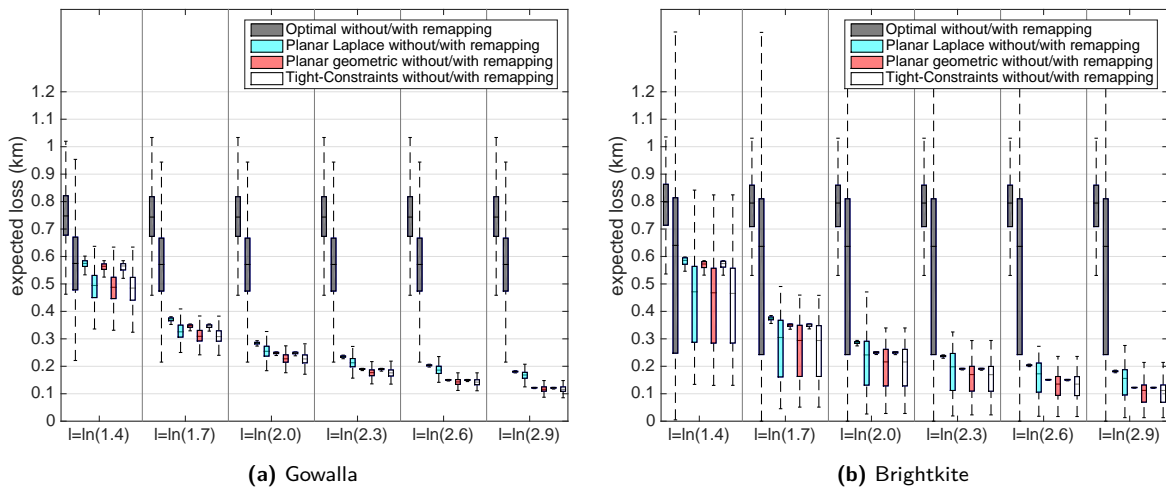


Fig. 5. The expected loss of privacy mechanisms using Gowalla (a) and Brightkite (b) datasets on the geographical region of Figure 6. The optimal mechanism is constructed on a coarse grid (6×14 cells) covering the region, and its remapped version is constructed on the large fine grained grid (60×140 cells). The other mechanisms are constructed directly on the fine grained grid. The parameter l describes the level of privacy within 0.1km. Higher values correspond to more relaxed privacy.



Fig. 6. A rectangular region of size 12 km \times 28 km covering most of the San Francisco peninsula

can be explained by the lower quality of data in the Brightkite dataset: many users have checkins in a single location, repeated thousands of times.

For the infinite discrete case, we perform the same evaluation, mapping all points in the Gowalla dataset to a infinite grid of cell size 0.1 km (hence, despite the discretization, locations are given with a relatively high precision). The planar Laplace and planar Geometric mechanisms are evaluated, with and without remapping. For PG, we use the continuous remapping of Section 3.2, since we are interested in the Euclidean loss metric and this technique is substantially more efficient than a direct Bayesian remap (given that we need to repeat it thousands of times). The results are shown in Figure 4; similarly to the continuous case, remapping is shown to be effective, despite the use of a generic prior. Moreover, PG slightly outperforms PL in most cases, which is expected since it is tailored to the discrete grid.

Finally, for the evaluation on a large but finite domain, we use the San Francisco region shown in Figure 6 for which the boundaries were specified earlier. In the case of the Gowalla dataset, our training set has 5216 users having 103052 check-ins in the considered region. On the other hand, the testing part consists of 273 users, where each one has at least 20 check-ins in the region. To allow acceptable precision of this discretization, we set the side length of every cell to be only 200 meters, hence splitting our region in a 60×140 grid of 8400 cells. Note that, although still finite, this size is two orders of magnitude larger than the one typically considered for constructing optimal mechanisms [5, 7], typically performed for domains of 50 to 100 cells.

For this grid, we construct the truncated planar Laplace mechanism, the truncated planar geometric mechanism, and the tight-constraints mechanism on the 8400 cells covering the required region. We also use a global prior from the training data to construct remapped versions of these mechanisms

Finally, we also construct the optimal ϵ -geo-indistinguishable mechanism described in [7]. But, as discussed in Section 4.3.1, since its construction is computationally expensive, we construct this mechanism on a *coarser* grid of 6×14 cells with cell spacing 2.0 km. Note that this mechanism is optimal only for the centers of the coarse grid, while being non-optimal for the fine grid consisting of 8400 cells. Based on this limitation, we finally investigate remapping the outputs of this mechanism to the cells of the required fine grid, using the global prior constructed from the training data.

For every mechanism we evaluate the expected loss for the individual users in the testing data, using their own priors. We demonstrate the results of this evaluation through the boxplot in Figure 5. In this figure we show the results for various levels of privacy. Similarly to the infinite case, we set $\epsilon = l/r$, where $r = 0.1$ km and l ranging between $\ln(1.4)$ and $\ln(2.9)$.

One clear implication from Figure 5 is that the optimal mechanism constructed on a coarse grid (6×14 cells) of the region is no longer optimal for the more realistic grid of 60×140 cells. Although this situation is improved by constructing a remapped version on the fine-grained grid (using a global prior), the remapped optimal mechanism still incurs higher quality loss, i.e. worse utility, relative to the other mechanisms, namely the Laplace, geometric, and the tight-constraints ones.

Another important observation is that the optimal mechanism and its remapped version maintain the same expected loss as the privacy level l grows. This phenomenon, discussed in Section 4.3.1, is due to the fact that the optimal mechanism satisfies ϵ -geo-indistinguishability only for the centers of the ‘large’ cells of the coarse grid. This makes the privacy constraints between these centers very weak since they are separated by large distances (2km), and therefore turns the mechanism into a deterministic ‘cloaking’, almost always revealing the center of the big cell containing the user’s location, regardless of the privacy level. This cloaking offers a privacy guarantee quite different than the intended geo-indistinguishability on the finer grid.

In contrast to the uniform behavior of the optimal mechanism and its remapped version, we observe that the expected losses of other mechanisms are significantly influenced by the imposed privacy level and provide bet-




Fig. 7. The expected loss (using Gowalla and Brightkite datasets) of the planar Laplace, planar geometric, and tight-constraints mechanisms constructed on the large grid (60×140 cells) covering San Francisco.

ter trade-off between the privacy requirements and the utility. For more clarity we show in Figure 7 the expected losses of the (truncated) Laplace, geometric and the tight-constraints mechanisms constructed directly on the fine grid of 8400 cells, and also their remapped versions constructed using the global prior.

The results show that the tight-constraints and the geometric mechanism clearly outperform the others. It is important here to remark that each one of these two mechanisms has an advantage over the other one. The geometric mechanism enjoys the feature that it exists for all levels of privacy due to the symmetry of the infinite grid on which it is constructed. The tight-constraints mechanism, on the other hand, might not exist for very strong levels of privacy; however, it has the advantage of being applicable to any privacy metric \mathbf{d} , not necessarily the Euclidean one.

Finally, we compare the tight-constraints and the exponential mechanism, the only two (efficient) ones that are applicable to an arbitrary privacy metric \mathbf{d} . Apart from the Euclidean one, we evaluate the two mechanisms using the well known maximum³ (or Chebyshev) metric, defined as $\mathbf{d}_\infty(x, x') = \epsilon \max_i (|x_i - x'_i|)$. This metric corresponds to a privacy property requiring the same level of privacy within a square area, instead of a circle. Note that in both cases $\epsilon = l/r$ scales the corresponding metric, deciding the radius (or the length of the square’s sides) in which a privacy level l is required.

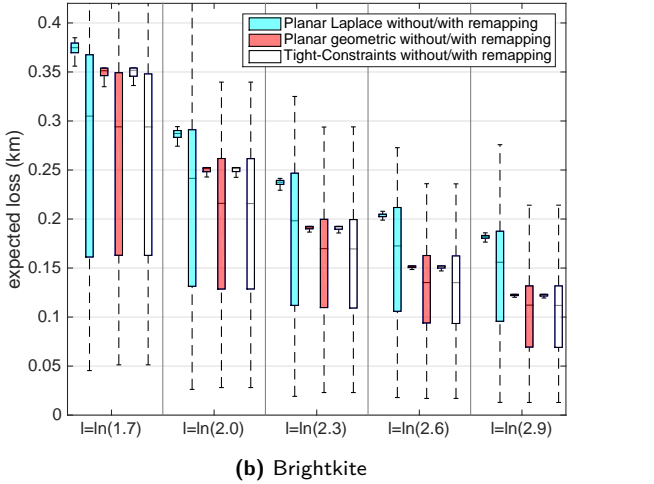


Figure 8 illustrates the utility of the two mechanisms for the Euclidean metric, showing that the tight-constraints mechanism significantly outperforms the exponential mechanism. For the maximum metric the results for the stronger levels of privacy $\ln(1.4)$ and $\ln(1.7)$, while it exists for other levels in which cases it outperforms the exponential mechanism. Hence we can conclude that whenever the tight-constraints mechanism exists, it outperforms the exponential one, while the latter is simpler to compute and always applicable.

6 Conclusion

In this work, we studied mechanisms for location privacy with emphasis on being *practical*, for realistic domain sizes and without unreasonable assumptions about the prior information available about the user. We discussed such solutions under various constraints for the domain of locations: the complete unconstrained continuous plain, the discrete plain or large but finite domains, using a Bayesian remap as a key ingredient. An extended evaluation was provided on two real-world datasets, showing considerable improvements wrt the standard planar Laplace mechanism.

³ Note that the maximum metric is a natural choice for location traces [6, 40], although we here use it for single locations.

References

- [1] K. Orland, "Stalker Victims Should Check For GPS." The Associated Press, 2003. <http://www.cbsnews.com/news/stalker-victims-should-check-for-gps/>.
- [2] J. Brownlee, "This Creepy App Isn't Just Stalking Women Without Their Knowledge, It's A Wake-Up Call About Facebook Privacy (Update)," 2012. <http://www.cultofmac.com/157641/>.
- [3] J. Simerman, "FasTrak to courthouse." East Bay Times, 2007. <http://www.eastbaytimes.com/2007/06/05/fastrak-to-courthouse/>.
- [4] D. Ashbrook and T. Starner, "Using gps to learn significant locations and predict movement across multiple users," *Personal and Ubiquitous Computing*, vol. 7, no. 5, pp. 275–286, 2003.
- [5] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. L. Boudec, "Protecting location privacy: optimal strategy against localization attacks," in *Proc. of CCS*, pp. 617–627, ACM, 2012.
- [6] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: differential privacy for location-based systems," in *Proc. of CCS*, pp. 901–914, ACM, 2013.
- [7] N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Optimal geo-indistinguishable mechanisms for location privacy," in *Proc. of CCS*, 2014.
- [8] R. Shokri, "Privacy games: Optimal user-centric data obfuscation," *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 2, pp. 299–315, 2015.
- [9] C. Dwork, "Differential privacy," in *Proc. of ICALP*, vol. 4052 of LNCS, pp. 1–12, Springer, 2006.
- [10] "Location guard." <https://github.com/chatziko/location-guard>.
- [11] K. Fawaz and K. G. Shin, "Location privacy protection for smartphone users," in *Proc. of CCS*, pp. 239–250, ACM Press, 2014.
- [12] K. Fawaz, H. Feng, and K. G. Shin, "Anatomization and protection of mobile apps' location privacy threats," in *Proc. of USENIX Security 2015*, pp. 753–768, USENIX Association, 2015.
- [13] C. Ma and C. W. Chen, "Nearby friend discovery with geo-indistinguishability to stalkers," *Procedia Computer Science*, vol. 34, pp. 352 – 359, 2014.
- [14] "Qgis processing provider plugin." https://github.com/SpatialVision/differential_privacy.
- [15] L. Pournajaf, L. Xiong, V. Sunderam, and X. Xu, "Stac: Spatial task assignment for crowd sensing with cloaked participant locations," in *Proceedings of the 23rd SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems, GIS '15*, pp. 90:1–90:4, ACM, 2015.
- [16] Y. Xiao and L. Xiong, "Protecting locations with differential privacy under temporal correlations," in *Proc. of CCS*, pp. 1298–1309, ACM, 2015.
- [17] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," in *Proc. of STOC*, pp. 351–360, ACM, 2009.
- [18] E. ElSalamouny, K. Chatzikokolakis, and C. Palamidessi, "Generalized differential privacy: Regions of priors that admit robust optimal mechanisms," in *Horizons of the Mind*, vol. 8464 of LNCS, pp. 292–318, Springer Int. Publishing, 2014.
- [19] M. Gruteser and D. Grunwald, "Anonymous usage of location-based services through spatial and temporal cloaking," in *Proc. of MobiSys*, USENIX, 2003.
- [20] P. Samarati and L. Sweeney, "Generalizing data to provide anonymity when disclosing information (abstract)," in *Proc. of PODS*, pp. 188–188, ACM Press, 1998.
- [21] L. Sweeney, "k-anonymity: A model for protecting privacy," *Int. Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 5, pp. 557–570, 2002.
- [22] L. Sweeney, "Achieving k-anonymity privacy protection using generalization and suppression," *Int. Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 5, pp. 571–588, 2002.
- [23] P. Samarati, "Protecting respondents' identities in micro-data release," *IEEE Trans. Knowl. Data Eng.*, vol. 13, no. 6, pp. 1010–1027, 2001.
- [24] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian, "l-diversity: Privacy beyond k-anonymity," *ACM Trans. on Knowledge Discovery from Data (TKDD)*, vol. 1, no. 1, p. 3, 2007.
- [25] N. Li, T. Li, and S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity," in *ICDE*, vol. 7, pp. 106–115, 2007.
- [26] A. Solanas, F. Seb e, and J. Domingo-Ferrer, "Micro-aggregation-based heuristics for p-sensitive k-anonymity: one step beyond," in *Proc. of PAIS 2008*, ACM Int. Conf. Proceeding Series, pp. 61–69, ACM, 2008.
- [27] A. R. Beresford and F. Stajano, "Location privacy in pervasive computing," *IEEE Pervasive Computing*, vol. 2, no. 1, pp. 46–55, 2003.
- [28] A. Machanavajjhala, D. Kifer, J. M. Abowd, J. Gehrke, and L. Vilhuber, "Privacy: Theory meets practice on the map," in *Proc. of ICDE*, pp. 277–286, IEEE, 2008.
- [29] S.-S. Ho and S. Ruan, "Differential privacy for location pattern mining," in *Proc. of SPRINGL*, pp. 17–24, ACM, 2011.
- [30] R. Dewri, "Local differential perturbations: Location privacy under approximate knowledge attackers," *IEEE Trans. on Mobile Computing*, vol. 99, no. PrePrints, p. 1, 2012.
- [31] F. Durr, P. Skvortsov, and K. Rothermel, "Position sharing for location privacy in non-trusted systems," in *Proc. of PerCom 2011*, pp. 189–196, IEEE, 2011.
- [32] E. ElSalamouny and S. Gamba, "Differential privacy models for location-based services," *Trans. on Data Privacy*, vol. 9, no. 1, pp. 15–48, 2016.
- [33] C. A. Ardagna, M. Cremonini, E. Damiani, S. D. C. di Vimercati, and P. Samarati, "Location privacy protection through obfuscation-based techniques," in *Proc. of DAS*, vol. 4602 of LNCS, pp. 47–60, Springer, 2007.
- [34] B. Bamba, L. Liu, P. Pesti, and T. Wang, "Supporting anonymous location queries in mobile environments with privacygrid," in *Proc. of WWW*, pp. 237–246, ACM, 2008.
- [35] M. Duckham and L. Kulik, "A formal model of obfuscation and negotiation for location privacy," in *Proc. of PERVASIVE*, vol. 3468 of LNCS, pp. 152–170, Springer, 2005.
- [36] M. Xue, P. Kalnis, and H. Pung, "Location diversity: Enhanced privacy protection in location based services," in *Proc. of LoCA*, vol. 5561 of LNCS, pp. 70–87, Springer,

2009.

- [37] B. Gedik and L. Liu, "Location privacy in mobile systems: A personalized anonymization model," in *Proc. of ICDCS*, pp. 620–629, IEEE, 2005.
- [38] K. Chatzikokolakis and C. Palamidessi, "Making random choices invisible to the scheduler," *Inf. and Comp.*, vol. 208, no. 6, pp. 694–715, 2010.
- [39] K. Chatzikokolakis, C. Palamidessi, and M. Stronati, "Constructing elastic distinguishability metrics for location privacy," *PoPETS*, vol. 2015, no. 2, pp. 156–170, 2015.
- [40] K. Chatzikokolakis, C. Palamidessi, and M. Stronati, "A predictive differentially-private mechanism for mobility traces," in *Proc. of PETS*, vol. 8555 of *LNCS*, pp. 21–41, Springer, 2014.

A Additional Figures

We include here figures removed from the paper due to space constraints.




Fig. 8. The expected loss of the exponential and tight-constraints mechanisms that satisfy geo-indistinguishability on the centers of the large grid (60×140 cells) covering San Francisco.

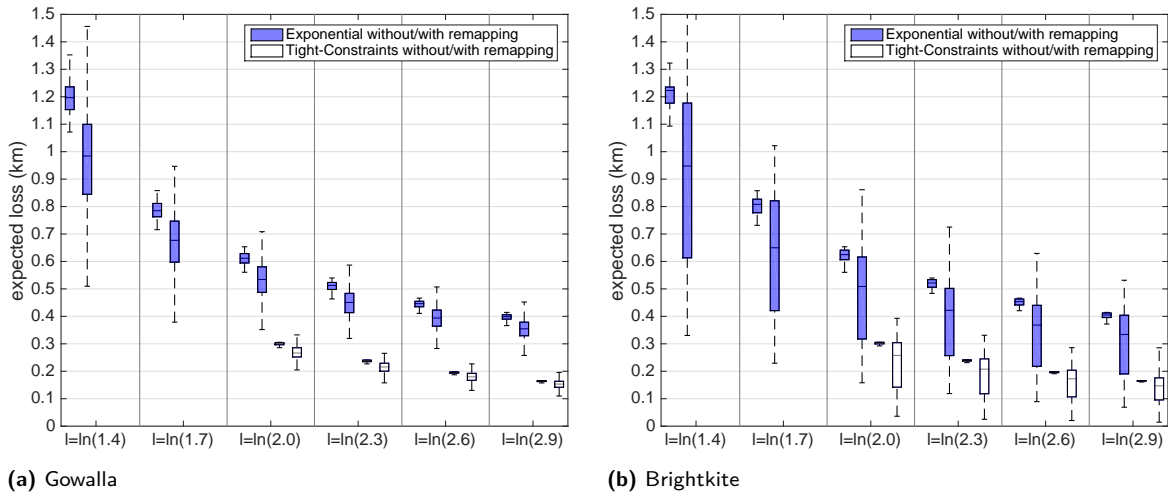


Fig. 9. The expected loss of the exponential and tight-constraints mechanisms for the maximum metric d_∞