Broadening the Scope of Differential Privacy Using Metrics

Konstantinos Chatzikokolakis, Miguel E. Andrés, Nicolás E. Bordenabe,
Catuscia Palamidessi

► To cite this version:

Konstantinos Chatzikokolakis, Miguel E. Andrés, Nicolás E. Bordenabe, Catuscia Palamidessi. Broadening the Scope of Differential Privacy Using Metrics. The 13th Privacy Enhancing Technologies Symposium, Jul 2013, Bloomington, Indiana, United States. pp.82-102, 10.1007/978-3-642-39077-7 . hal-00767210

HAL Id: hal-00767210

<https://hal.inria.fr/hal-00767210>

Submitted on 17 May 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Broadening the scope of Differential Privacy Using Metrics^{*}

Konstantinos Chatzikokolakis^{1,2}, Miguel E. Andrés²,
Nicolás E. Bordenabe^{3,2}, and Catuscia Palamidessi^{3,2}

¹ CNRS

² LIX, Ecole Polytechnique

³ INRIA

Abstract. Differential Privacy is one of the most prominent frameworks used to deal with disclosure prevention in statistical databases. It provides a formal privacy guarantee, ensuring that sensitive information relative to individuals cannot be easily inferred by disclosing answers to aggregate queries. If two databases are adjacent, i.e. differ only for an individual, then the query should not allow to tell them apart by more than a certain factor. This induces a bound also on the distinguishability of two generic databases, which is determined by their distance on the Hamming graph of the adjacency relation.

In this paper we explore the implications of differential privacy when the indistinguishability requirement depends on an arbitrary notion of distance. We show that we can naturally express, in this way, (protection against) privacy threats that cannot be represented with the standard notion, leading to new applications of the differential privacy framework. We give intuitive characterizations of these threats in terms of Bayesian adversaries, which generalize two interpretations of (standard) differential privacy from the literature. We revisit the well-known results stating that universally optimal mechanisms exist only for counting queries: We show that, in our extended setting, universally optimal mechanisms exist for other queries too, notably sum, average, and percentile queries. We explore various applications of the generalized definition, for statistical databases as well as for other areas, such that geolocation and smart metering.

1 Introduction

Differential privacy [1,2] is a formal definition of privacy which originated from the area of statistical databases, and it is now applied in many other domains, ranging from programming languages [3] to social networks [4] and geolocation [5].

Statistical databases are queried by analysts to obtain aggregate information about individuals. It is important to protect the privacy of the participants in the database, in the sense that it should not be possible to infer the value of an individual from the aggregate information. This can be achieved by adding random noise to the answer.

^{*} This work is partially funded by the Inria large scale initiative CAPPRIS, the EU FP7 grant no. 295261 (MEALS), and the project ANR-12-IS02-001 PACE. Nicolás E. Bordenabe was partially funded by the French Defense procurement agency (DGA) with a PhD grant.

Because of the focus on the single individual as the unit of protection, differential privacy relies in a crucial way on the notion of two databases being *adjacent*, i.e. differing only for an individual. A mechanism K is ϵ -differentially private if for any two adjacent databases x, x' , and any property Z , the probability distributions $K(x), K(x')$ differ on Z at most by e^ϵ , namely, $K(x)(Z) \leq e^\epsilon K(x')(Z)$. For two non-adjacent databases, there is no requirement other than the one induced by the transitive application of the property. Note that the set of all possible databases, together with the adjacency relation, forms a *Hamming graph*, and the graph distance $d_h(x, x')$ between x and x' is exactly the number of individuals in which x and x' differ. Then, for any databases x, x' , it is easy to see (by transitivity on a path from x to x') that $K(x)(Z) \leq e^{\epsilon d_h(x, x')} K(x')(Z)$. We can view $\epsilon d_h(x, x')$ as the distinguishability level between two generic databases x, x' : the smaller $\epsilon d_h(x, x')$ is, the more similar the probability distributions $K(x), K(x')$ are required to be.

When the sensitive information to be protected is other than the value of a single individual, it is common to consider different notions of adjacency. For example, in cases of cohesive groups with highly correlated values, we could consider adjacent two databases differing in any number of individuals of the same group. Similarly, when dealing with friendship graphs in social networks, adjacency could be defined as differing in a single edge.

We argue that in some situations the distinguishability level between x and x' should depend not only on the number of different values between x and x' , but also on the values themselves. We might require, for instance, databases in which the value of an individual is only slightly modified to be highly indistinguishable, thus protecting the *accuracy* by which an analyst can infer an individual's value.

More generally, we might want to apply differential privacy in scenarios when x, x' are not databases at all, but belong to an *arbitrary domain of secrets* $\mathcal{X}$. In such a scenario, there might be no natural notion of adjacency, but it is still reasonable to define a distinguishability level between secrets, and employ the same principle of differential privacy – i.e. the smaller the distinguishability level between x, x' is, the more similar the probability distributions $K(x), K(x')$ are required to be – to obtain a meaningful notion of privacy. For instance, when dealing with geographic locations (aka, geolocation), it might be acceptable to disclose the fact that an individual is in Paris rather than in New York. However, disclosing the *precise* location of the individual within Paris is likely to be undesired (because, for instance, the individual is currently in Moulin Rouge rather than in his office in Place d'Italie). Thus it would be useful to have a distinguishability level that depends on the geographical distance.

In this paper we assume that we have a numeric function $\epsilon(x, x')$, giving the distinguishability level between x, x' , which depends on the application at hand and the privacy guarantees we wish to express. The corresponding notion of privacy is the requirement that for an arbitrary pair x, x' we have

$$K(x)(Z) \leq e^{\epsilon(x, x')} K(x')(Z)$$

Note that standard ϵ -differential privacy is a particular case of this notion, that we obtain by setting $\epsilon(x, x') = \epsilon d_h(x, x')$.

Since ϵ models distinguishability, there are some properties that it is expected to satisfy. First, it should be the case that any element is indistinguishable from itself, i.e.

$\varepsilon(x, x) = 0$. Second, the distinguishability level of x and x' should be the same as that of x' and x , i.e. $\varepsilon(x, x') = \varepsilon(x', x)$ (symmetry). Finally, if x_1 and x_2 are hardly distinguishable from x_3 , then they should be also hardly distinguishable from each other. In other words, $\varepsilon(x_1, x_2)$ should be bounded by a function of $\varepsilon(x_1, x_3), \varepsilon(x_2, x_3)$. In this paper we assume the triangle inequality, namely $\varepsilon(x_1, x_2) \leq \varepsilon(x_1, x_3) + \varepsilon(x_3, x_2)$, which means that ε is a metric. We believe that more relaxed notions of distinguishability could potentially be useful, we leave the investigation of this possibility as future work.⁴ In the rest of this paper we use d (for “distance”) instead of ε , and we call the corresponding privacy notion “ d -privacy”.

Similarly to the standard definition, d -privacy does not explicitly talk about the adversary’s gain of knowledge. In order to better understand a privacy property, however, it is useful to provide interpretations that directly reason about the capabilities of the adversary. Two such interpretations exist for differential privacy: the first states that, regardless of side knowledge, the adversary’s gain of knowledge by observing the reported answer is the same whether or not the individual’s data were included in the database [1,6]. The second states that, an informed adversary who already knows all values except individual’s i , gains no extra knowledge from the reported answer, regardless of side knowledge about i ’s value [2].⁵

In the case of d -privacy, we provide two results that generalize the above interpretations, showing the privacy guarantees provided by a certain metric d . The first uses the concept of a *hiding* function $\phi : \mathcal{X} \rightarrow \mathcal{X}$. The idea is that ϕ can be applied to a secret x before the mechanism K , so that the latter has only access to a hidden version $\phi(x)$, instead of the real secret x . Then d -privacy implies that the adversary’s conclusions (captured by his posterior distribution) are similar (up to a factor depending on ϕ) regardless of whether ϕ is applied to the secret or not. Moreover, we show that certain classes of hiding functions are “canonical”, in the sense that if the property holds for those functions, it must hold in general.

The above characterization compares two posterior distributions and does not imply that the adversary learns no information, but that he learns the same regardless of whether the secret has been hidden or not. We then give a second characterization, comparing the adversary’s conclusions (a posterior distribution) to his initial knowledge (a prior distribution). Since some information is allowed to be revealed, we cannot expect the two to be similar. Still, if we restrict to a neighborhood N of secrets that are close to each other, we can show that d -privacy implies that an informed adversary, knowing that the secret belongs to N , can gain little more information about the exact secret, regardless of his prior knowledge within N . Similarly to the previous characterization, we also show that certain classes of neighborhoods are canonical.

We give examples of privacy problems in various contexts, and show how to define appropriate metrics. In the context of statistical databases, we consider metrics that

⁴ Several of our results do not actually depend on ε being a metric.

⁵ The knowledge increase of a non-informed adversary is not bounded by e^ϵ . Recalling the well-known example from [1], consider the side information that Terry Gross is two inches shorter than the average Lithuanian woman. Then obtaining the average height (even a noisy one) gives little additional information about Terry Gross to an informed adversary, but substantial information to a non-informed one.

depend not only on the number of different values, but also on the values themselves. First, a stronger variant of differential privacy is given in which databases differing in a single individual are hardly distinguishable, but the distinguishability level becomes even lower when the difference in the values is small. Moreover, this metric can be relaxed to obtain a privacy notion that focuses on protecting the accuracy of a value. This can be useful, for instance, in case an individual does not mind disclosing his age group, but wants to protect his exact birthday date (such precise information could in principle allow to identify the individual with little margin of error).

Departing from statistical databases, we consider smart meters, and the problem for privacy that can derive from accurate measurement of energy consumption at high frequency. Further, we consider the problem of hiding the exact position in location-based services. In all these examples, besides the proper metric notion, we construct also the canonical adversary which provides the operational interpretation.

Next, we turn our attention to the notion of utility, namely the accuracy of the reported answer, and in particular the Bayesian notion of utility [7,8], which takes into account the prior knowledge of the user. In general mechanisms may provide different degrees of utility for the same level of privacy, and obviously it is desirable to identify the optimal ones. Of particular interest are the *universally optimal* mechanisms, which provide optimal utility for all users (i.e., all priors). There are two well known results concerning universal optimality: the first [7] establishes that for counting queries the geometric and the truncated geometric mechanisms are universally optimal. The second [8] says that for any other kind of query no universally optimal mechanism exists.

We revisit these results in our framework and show that in contrast to the standard case, d -privacy allows to construct (for certain metrics) universally optimal mechanisms for many other kinds of queries. More precisely, we show that universally optimal mechanisms exist in the cases of (i) the sum, average and percentile queries for the Manhattan metric, and (ii) the average and percentile queries for the Maximum metric.

We also study the additional noise required to achieve privacy for databases queries, when we use a finer metric than the Hamming distance. Surprisingly, it turns out that in the case (i) above, the sensitivity of the queries remains the same as in the standard case. This means that, a standard ϵ -differentially private mechanism already incorporates “for free” the additional protection w.r.t. proximity of values.

Related Work. Several works in the differential privacy literature consider adjacency relations different than the standard one, effectively using a metric tailored to that application. Examples include group privacy [1] and edge privacy for graphs [9].

The generalization of differential privacy to arbitrary metrics was considered also in [10,3]. In those works, however, the purpose of extending the definition was to obtain compositional methods for proving differential privacy in programming languages, while in our work we focus on the implications of such extension for the theory of differential privacy. Namely, we aim at obtaining new meaningful definitions of privacy for various contexts through the use of different metrics (cf. the examples of the smart meters and of geolocation), and at investigating the existence of optimal mechanisms.

Another work closely related to ours is [11] in which an extended definition of differential privacy is used to capture the notion of fairness in classification. A metric d is used to model the fact that certain individuals are required to be classified similarly,

and a mechanism satisfying d -privacy is considered fair, since it produces similar results for similar individuals. We view fairness as one of the many interesting notions that can be obtained through the use of metrics in various contexts, thus it encourages our goal of studying d -privacy. With respect to the actual metrics used in this paper, the difference is that we consider metrics that depend on the individuals' values, while [11] considers metrics between individuals.

Contribution. The main contributions of this paper are summarized below:

- We study d -privacy – an extension of differential privacy to arbitrary domains endowed with a metric d – in the general case, independently from any specific metric.
- We give two operational characterizations of d -privacy that directly constraint the capabilities of the adversary.
- We show examples of applications of d -privacy to privacy scenarios both in databases and in other contexts.
- We show that several queries (including the sum, average and percentile) admit universally optimal mechanisms for certain metrics. This contrasts sharply with standard differential privacy, where such mechanisms exist only for counting queries.

Plan of the Paper. In the next section we recall some preliminary notions about mechanisms, metrics, and differential privacy. Section 3 introduces the notion of d -privacy and presents two characterization results. In Section 4 we give a sufficient and necessary condition for the privacy of an oblivious mechanism, we discuss Laplace mechanisms, and we give sufficient conditions for a mechanism to be optimal. In Sections 5 and 6 we give several examples of applications of our notions, in statistical databases with an enriched notion of privacy, and in other domains, respectively. We also show how to construct universally optimal mechanisms for some of those examples in the cases of sum, average, and percentile queries. Section 7 concludes.

All proofs can be found in the appendix.

2 Preliminaries

Mechanisms. Given two sets $\mathcal{X}$ and $\mathcal{Z}$, let $\mathcal{F}_{\mathcal{Z}}$ be a σ -algebra over $\mathcal{Z}$ and let $\mathcal{P}(\mathcal{Z})$ be the set of probability measures over $\mathcal{Z}$. A *mechanism* from $\mathcal{X}$ to $\mathcal{Z}$ is a (probabilistic) function $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$. A mechanism K can be described in terms of probability density functions (pdf's), that is by a function $D : \mathcal{X} \rightarrow \mathcal{D}(\mathcal{Z})$ (where $\mathcal{D}(\mathcal{Z})$ denotes the space of the pdf's over $\mathcal{Z}$), such that $D(x)$ is the pdf of $K(x)$.

The composition $H \circ f$ of a deterministic function $f : \mathcal{X} \rightarrow \mathcal{Y}$ (called a *query*) and a mechanism $H : \mathcal{Y} \rightarrow \mathcal{P}(\mathcal{Z})$ is the mechanism $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ defined as $K(x) = (H \circ f)(x) = H(f(x))$. Mechanisms of this form are called *oblivious*.

Let π be a discrete probability measure on $\mathcal{X}$, called a *prior*.⁶ Starting from π and using Bayes' rule, each observation $Z \in \mathcal{Z}$ of a mechanism $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ induces a *posterior* measure $\sigma = \text{Bayes}(\pi, K, Z)$ on $\mathcal{X}$, defined as $\sigma(x) = \frac{K(x)(Z)\pi(x)}{\sum_{x' \in \mathcal{X}} K(x')(Z)\pi(x')}$.

⁶ We restrict to discrete priors for simplicity; all results could be carried to the continuous case.

Metrics. A metric on a set $\mathcal{X}$ is a function $d_{\mathcal{X}} : \mathcal{X}^2 \rightarrow [0, \infty]$ such that $d_{\mathcal{X}}(x, y) = 0$ iff $x = y$, $d_{\mathcal{X}}(x, y) = d_{\mathcal{X}}(y, x)$, and $d_{\mathcal{X}}(x, z) \leq d_{\mathcal{X}}(x, y) + d_{\mathcal{X}}(y, z)$ for all $x, y, z \in \mathcal{X}$. The *diameter* of $A \subseteq \mathcal{X}$ is defined as $d_{\mathcal{X}}(A) = \sup_{x, x' \in A} d_{\mathcal{X}}(x, x')$.

A sequence $x_1, \dots, x_n$ is called a *chain* from x_1 to x_n and denoted by $\tilde{x}$. The length $d_{\mathcal{X}}(\tilde{x})$ of a chain is defined as $d_{\mathcal{X}}(\tilde{x}) = \sum_{i=1}^{n-1} d_{\mathcal{X}}(x_i, x_{i+1})$. If $d_{\mathcal{X}}(\tilde{x}) = d_{\mathcal{X}}(x_1, x_n)$ then $\tilde{x}$ is called *tight*.

Of particular interest are metrics *induced by a graph* $(\mathcal{X}, \sim_{\mathcal{X}})$, where $\sim_{\mathcal{X}}$ is the graph's adjacency relation. In the induced metric, $d_{\mathcal{X}}(x, x')$ is the length of the shortest path from x to x' (or infinite if no path exists). Of great interest are also the Manhattan (or L_1), the Euclidean (or L_2) and the Maximum (or L_{∞}) metrics, denoted by d_1, d_2, d_{∞} respectively. The numerical distance on the reals (which coincides with all d_1, d_2, d_{∞}) will be denoted by $d_{\mathbb{R}}$ for clarity. Finally, of great interest is the metric $d_{\mathcal{P}}$ on $\mathcal{P}(\mathcal{Z})$ defined as $d_{\mathcal{P}}(\mu_1, \mu_2) = \sup_{Z \in \mathcal{F}_{\mathcal{Z}}} |\ln \frac{\mu_1(Z)}{\mu_2(Z)}|$ with the convention that $|\ln \frac{\mu_1(Z)}{\mu_2(Z)}| = 0$ if both $\mu_1(Z), \mu_2(Z)$ are zero and ∞ if only one of them is zero.

Differential Privacy. We fix a finite domain of values $\mathcal{V}$, called the *universe*. A database $x \in \mathcal{V}^n$ consists of n records from $\mathcal{V}$ - each corresponding to an individual - that is x is a tuple $\langle x[1], \dots, x[n] \rangle$, $x[i] \in \mathcal{V}$, where $x[i]$ is the value of the i -th individual in the database. We denote by $x^{[v/i]}$ the database obtained from x by substituting the value v for individual i . The case when individuals are allowed to be absent from the database can be modeled by the universe $\mathcal{V}_{\emptyset} = \mathcal{V} \cup \{\emptyset\}$ where the null value $\emptyset$ denotes absence.

A crucial notion for differential privacy is that of *adjacency*: two databases x, x' are adjacent, written $x \sim_h x'$, if they differ in exactly one element. Let d_h be the distance induced by $\sim_h$ (i.e., $d_h(x, x')$ is the number of elements in which x, x' differ). The graph $(\mathcal{V}^n, \sim_h)$ is known as *Hamming graph*, and d_h as Hamming distance.

Let $\mathcal{Z}$ be a set of query outcomes; a mechanism $K : \mathcal{V}^n \rightarrow \mathcal{P}(\mathcal{Z})$ satisfies ϵ -differential privacy if adjacent databases produce answers with probabilities that differ at most by a factor e^{ϵ} :

$$K(x)(Z) \leq e^{\epsilon} K(x')(Z) \quad \forall x \sim_h x' \in \mathcal{V}^n, Z \in \mathcal{F}_{\mathcal{Z}} \quad (1)$$

Following [3], the definition can be expressed as $d_{\mathcal{P}}(K(x), K(x')) \leq \epsilon$ for all $x \sim_h x'$. Moreover, as explained in the introduction, we can rewrite it in terms of the Hamming distance: $d_{\mathcal{P}}(K(x), K(x')) \leq \epsilon d_h(x, x')$ for all $x, x' \in \mathcal{V}^n$.

A desirable feature of this definition is that it solely depends on the mechanism itself, without explicitly talking about the adversary's side knowledge, or the information that he learns from the reported answer. However, in order to get a better understanding of a privacy definition, it is useful to give an "operational" (or "semantic") interpretation that directly restricts the abilities of the adversary. To this end, we capture the adversary's side knowledge by a prior distribution π on $\mathcal{V}^n$, and his conclusions after observing Z by the posterior distribution $\sigma = \text{Bayes}(\pi, K, Z)$.

There are two operational interpretations commonly given to differential privacy. The first can be informally stated as: "regardless of side knowledge, by observing the reported answer an adversary obtains the same information whether or not the individual's data were included in the database". This can be formalized as follows: consider a *hiding* function $\phi_{i,v} : \mathcal{V}^n \rightarrow \mathcal{V}^n$ replacing i 's value by a fixed value v , i.e.

$\phi_{i,v}(x) = x^{[v/i]}$, and let $\Phi_h = \{\phi_{i,v} \mid i \in 1..n, v \in \mathcal{V}\}$ be the set of all such functions. The mechanism $K \circ \phi_{i,v}$ behaves as K after removing i 's value; hence we require the posterior distributions induced by $K, K \circ \phi_{i,v}$ to be similar. The resulting notion (called “semantic privacy” in [6])⁷ can be shown to be implied by differential privacy.

Theorem 1 ([6]). *If a mechanism $K : \mathcal{V}^n \rightarrow \mathcal{P}(\mathcal{Z})$ satisfies ϵ -differential privacy then for all priors π on $\mathcal{V}^n$, all $\phi \in \Phi_h$, and all $Z \in \mathcal{F}_{\mathcal{Z}}$:*

$$d_{\mathcal{P}}(\sigma_1, \sigma_2) \leq 2\epsilon \quad \text{where } \sigma_1 = \mathbf{Bayes}(\pi, K, Z) \text{ and } \sigma_2 = \mathbf{Bayes}(\pi, K \circ \phi, Z)$$

Note that the above interpretation compares two *posterior* measures. This requirement does not imply that the adversary learns no information, but that he learns the same regardless of the presence of the individual's data. Both σ_1, σ_2 can be very different than the prior π , as the well-known example of Terry Gross [1] demonstrates.

A different interpretation can be obtained by comparing the posterior σ to the prior distribution π . Of course, we cannot expect those to be similar, since some information is allowed to be disclosed. Still, we can require the distributions to be similar when restricted to the value of a single individual, by assuming an informed adversary who knows all other values in the database. Let $N_i(x) = \{x^{[v/i]} \mid v \in \mathcal{V}\}$ denote the set of databases obtained from x by modifying i 's value, and let $\mathcal{N}_h = \{N_i(x) \mid x \in \mathcal{V}^n, i \in 1..n\}$. Knowing that the database belongs to a set $N \in \mathcal{N}_h$ means that we know all values except one. We denote by $\pi_{|N}$ the distribution obtained from π by restricting to N , i.e. $\pi_{|N}(x) = \pi(x|N)$. Requiring $\pi_{|N}, \sigma_{|N}$ to be similar brings us the definition of “semantic security” from [2], which is a full characterization of differential privacy.

Theorem 2 ([2]). *A mechanism $K : \mathcal{V}^n \rightarrow \mathcal{P}(\mathcal{Z})$ satisfies ϵ -differential privacy iff for all priors π on $\mathcal{V}^n$, all $N \in \mathcal{N}_h$, and all $Z \in \mathcal{F}_{\mathcal{Z}}$:*

$$d_{\mathcal{P}}(\pi_{|N}, \sigma_{|N}) \leq \epsilon \quad \text{where } \sigma = \mathbf{Bayes}(\pi, K, Z)$$

Note that if the adversary does not know $N \in \mathcal{N}_h$, then his knowledge can (and will in most cases) be increased. Note also that the above result does not imply that K allows the adversary to learn $N_i(x)$! In fact, this is clearly forbidden since it would violate the same condition for $N_j(x), j \neq i$, i.e. it would violate the other individuals' privacy.

3 Generalized Privacy

As discussed in the introduction, differential privacy can be generalized to the case of an arbitrary set of secrets $\mathcal{X}$, equipped with a metric $d_{\mathcal{X}}$.

Definition 1. *A mechanism $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$ satisfies $d_{\mathcal{X}}$ -privacy, iff $\forall x, x' \in \mathcal{X}$: $d_{\mathcal{P}}(K(x), K(x')) \leq d_{\mathcal{X}}(x, x')$, or equivalently:*

$$K(x)(Z) \leq e^{d_{\mathcal{X}}(x, x')} K(x')(Z) \quad \forall Z \in \mathcal{F}_{\mathcal{Z}}$$

⁷ The only difference between the semantic privacy of [6] and our formulation is that an “additive” metric between distributions is used instead of the “multiplicative” $d_{\mathcal{P}}$.

Intuitively, the definition requires that secrets close to each other wrt $d_{\mathcal{X}}$, meaning hardly distinguishable, should produce outcomes with similar probability. This is the same core idea as in differential privacy, which can be retrieved as $\mathcal{X} = \mathcal{V}^n, d_{\mathcal{X}} = \epsilon d_h$.

Note that Definition 1 contains no ϵ ; the distinguishability level is directly given by the metric. In practice, the desired metric can be obtained from a standard one by scaling by a proper factor ϵ (recall that a scaled metric is also a metric). For instance, in the case of standard differential privacy, the Hamming distance between adjacent databases is 1, and we want their distinguishability level to be ϵ , hence we use the scaled version ϵd_h .

Note also that an *extended* metric (allowing $d_{\mathcal{X}}(x, x') = \infty$) can be useful in cases when we allow two secrets to be completely distinguished. The understanding of Definition 1 is that the requirement is always satisfied for those secrets. Similarly, *pseudo*-metrics (allowing $d_{\mathcal{X}}(x, x') = 0$ for $x \neq x'$) could be useful when we want some secrets to be completely indistinguishable (forcing $K(x)$ and $K(x')$ to be identical). To simplify the presentation, the results of this paper assume an extended metric (but not pseudo). An approximate version of $d_{\mathcal{X}}$ -privacy can be defined, similarly to (α, δ) differential privacy [13]. We leave the study of such notion as future work.

Different metrics $d_{\mathcal{X}}, d_{\mathcal{X}}'$ on the same set $\mathcal{X}$ clearly give rise to different privacy notions. The “strength” of each notion depends on the distinguishability level assigned to each pair of secrets; $d_{\mathcal{X}}$ -privacy and $d_{\mathcal{X}}'$ -privacy are in general incomparable. However, lower distinguishability level implies stronger privacy.

Proposition 1. *If $d_{\mathcal{X}} \leq d_{\mathcal{X}}'$ (point-wise) then $d_{\mathcal{X}}$ -privacy implies $d_{\mathcal{X}}'$ -privacy.*

For example, some works consider an adjacency relation $\sim_r$ slightly different than $\sim_h$, defined as $x \sim_r x'$ iff $x' = x^{[\mathcal{O}/i]}$ (or vice versa), i.e. x' can be obtained from x by removing one individual. This relation gives rise to a metric d_r for which it holds that: $\frac{1}{2}d_r \leq d_h \leq d_r$. From Proposition 1, the two models are essentially equivalent; one can obtain ϵd_r -privacy from ϵd_h -privacy by doubling ϵ and vice versa.

Characterization 1. Similarly to standard differential privacy, $d_{\mathcal{X}}$ -privacy does not explicitly talk about the adversary’s gain of knowledge. To better understand the privacy guarantees provided by a certain metric $d_{\mathcal{X}}$, it is useful to directly reason about the capabilities of the adversary. Two such characterizations are given, generalizing the two interpretations of standard differential privacy (Theorems 1,2).

The first characterization uses the concept of a *hiding* function $\phi : \mathcal{X} \rightarrow \mathcal{X}$. The idea is that ϕ can be applied to x before the mechanism K , so that the latter has only access to a hidden version $\phi(x)$, instead of the real secret x . Let $d_{\mathcal{X}}(\phi) = \sup_{x \in \mathcal{X}} d_{\mathcal{X}}(x, \phi(x))$ be the maximum distance between a secret and its hidden version. We can show that $d_{\mathcal{X}}$ -privacy implies that the adversary’s conclusions (captured by his posterior measure) are the same (up to $2d_{\mathcal{X}}(\phi)$) regardless of whether ϕ is applied or not. Moreover, we show that certain classes of hiding functions are “canonical”, in the sense that if the property holds for those, it must hold in general. We start by defining this class.

Definition 2. *Let Φ be a set of functions from $\mathcal{X}$ to $\mathcal{X}$, called hiding functions. A chain $\tilde{x}$ is called a maximal Φ -chain iff for every step i there exists $\phi \in \Phi$ s.t. $\phi(x_i) = x_{i+1}$, $\phi(x_{i+1}) = x_i$ and $d_{\mathcal{X}}(x_i, x_{i+1}) = d_{\mathcal{X}}(\phi)$. Then Φ is called maximally tight wrt $d_{\mathcal{X}}$ iff $\forall x, x' \in \mathcal{X}$ there exists a tight maximal Φ -chain from x to x' .*

Note that the above property requires hiding functions that *swap* the secrets x_i, x_{i+1} . This is not satisfied by the hiding functions $\phi_{i,v}$ introduced in the previous section, but will be satisfied by more general functions used later in the paper.

Theorem 3. *Let Φ be a set of hiding functions. If K satisfies $d_{\mathcal{X}}$ -privacy then for all $\phi \in \Phi$, all priors π on $\mathcal{X}$, and all $Z \in \mathcal{F}_{\mathcal{Z}}$:*

$$d_{\mathcal{P}}(\sigma_1, \sigma_2) \leq 2 d_{\mathcal{X}}(\phi) \quad \text{where } \sigma_1 = \mathbf{Bayes}(\pi, K, Z) \text{ and } \sigma_2 = \mathbf{Bayes}(\pi, K \circ \phi, Z)$$

If Φ is maximally tight then the converse also holds.

The above characterization compares two posterior distributions; hence, it does not impose that the adversary gains no information, but that this information is the same regardless of whether ϕ has been applied to the secret or not.

Characterization 2. A different approach is to compare the adversary's prior and posterior distributions, measuring how much he learned about the secret. Since we allow some information to be revealed, we cannot expect these distributions to be similar. Still, if we restrict to a neighborhood N of secrets that are close to each other, we can show that $d_{\mathcal{X}}$ -privacy implies that an informed adversary, knowing that the secret belongs to N , can gain little more information about the exact secret regardless of his side knowledge about N . Moreover, similarly to the previous characterization, we show that certain classes of neighborhoods are "canonical".

Definition 3. *Let $\mathcal{N} \subseteq 2^{\mathcal{X}}$. The elements of $\mathcal{N}$ are called neighborhoods. A chain $\tilde{x}$ is called a maximal $\mathcal{N}$ -chain iff for every step i there exists $N \in \mathcal{N}$ such that $\{x_i, x_{i+1}\} \subseteq N$ and $d_{\mathcal{X}}(x_i, x_{i+1}) = d_{\mathcal{X}}(N)$. Then $\mathcal{N}$ is called maximally tight wrt $d_{\mathcal{X}}$ iff $\forall x, x' \in \mathcal{X}$ there exists a tight maximal $\mathcal{N}$ -chain from x to x' .*

Theorem 4. *Let $\mathcal{N} \subseteq 2^{\mathcal{X}}$. If K satisfies $d_{\mathcal{X}}$ -privacy then for all $N \in \mathcal{N}$, all priors π on $\mathcal{X}$, and all $Z \in \mathcal{F}_{\mathcal{Z}}$:*

$$d_{\mathcal{P}}(\pi|_N, \sigma|_N) \leq d_{\mathcal{X}}(N) \quad \text{where } \sigma = \mathbf{Bayes}(\pi, K, Z)$$

If $\mathcal{N}$ is maximally tight then the converse also holds.

Using meaningful (and maximally tight) sets $\Phi, \mathcal{N}$, and applying the above characterizations, we can get an intuitive understanding of the privacy guarantees offered by $d_{\mathcal{X}}$ -privacy. For example, in the case of databases, it can be shown that $\mathcal{N}_h$ is maximally tight wrt the d_h metric, hence the characterization of Theorem 2 can be obtained as a special case of Theorem 4. Theorem 1 can also be obtained from Theorem 3 (even though Φ_h is not maximally tight) since it only states an implication in one direction.

4 Answering Queries

To obtain the answer to a query $f : \mathcal{X} \rightarrow \mathcal{Y}$ in a private way, we can compose it with a mechanism $H : \mathcal{Y} \rightarrow \mathcal{P}(\mathcal{Z})$, thus obtaining an oblivious mechanism $H \circ f :$

Fig. 1. Counterexample to the converse of Fact 5. The table represents the distribution H . We note that $H \circ f$ satisfies $(\ln 2)$ -privacy, and that f is 1-sensitive. However $H(y_1)(z_1) = 3/4 \not\leq 2 H(y_3)(z_1) = 2 \cdot 1/4$, hence H does not satisfy $(\ln 2)$ -privacy.

$\mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$. In this section, we first state the standard compositionality result about the privacy of $H \circ f$, relying on the notion of Δ -sensitivity (aka Lipschitz continuity), naturally extended to the case of $d_{\mathcal{X}}$ -privacy. Then, we introduce the concept of *uniform sensitivity*, and we use it to obtain the converse of the aforementioned compositionality result, which in turn allows to give optimality results later in the paper.

Definition 4. f is Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$ iff $d_{\mathcal{Y}}(f(x), f(x')) \leq \Delta d_{\mathcal{X}}(x, x')$ for all $x, x' \in \mathcal{X}$. The smallest such Δ (if exists) is called the sensitivity of f wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$.

Fact 5. Assume that f is Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$ and H satisfies $d_{\mathcal{Y}}$ -privacy. Then $H \circ f$ satisfies $\Delta d_{\mathcal{X}}$ -privacy.

Note that it is common to define a family of mechanisms $H_{\epsilon}, \epsilon > 0$, instead of a single one, where each H_{ϵ} satisfies privacy for a scaled version $\epsilon d_{\mathcal{Y}}$ of a metric of interest $d_{\mathcal{Y}}$. Given such a family and a query f , we can define a family of oblivious mechanisms $K_{\epsilon} = H_{\epsilon/\Delta} \circ f, \epsilon > 0$, each satisfying $\epsilon d_{\mathcal{X}}$ -privacy (from Fact 5).

The converse of the above result does not hold in general, see Fig. 1 for a counterexample. However, it does hold if we replace the notion of sensitivity by the stronger notion of *uniform sensitivity*.

Definition 5. Two elements $y, y' \in \mathcal{Y}$ are called Δ -expansive iff $d_{\mathcal{Y}}(y, y') = \Delta d_{\mathcal{X}}(x, x')$ for some $x \in f^{-1}(y), x' \in f^{-1}(y')$. A chain $\tilde{y}$ is Δ -expansive iff all steps y_i, y_{i+1} are Δ -expansive. Finally, f is uniformly Δ -sensitive iff it is Δ -sensitive and for all $y, y' \in \mathcal{Y}$ there exists a tight and Δ -expansive chain from y to y' .

Theorem 6. Assume that f is uniformly Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$. Then H satisfies $d_{\mathcal{Y}}$ -privacy if and only if $H \circ f$ satisfies $\Delta d_{\mathcal{X}}$ -privacy.

4.1 Laplace Mechanisms

Adding Laplace noise is the most widely used technique for achieving differential privacy. The mechanism can be naturally adapted to any metric, using a variant of the exponential mechanism [14], by providing a properly constructed scaling function. Note that in the framework of d -privacy, we can express the privacy of the mechanism itself, on its own domain, without the need to consider a query or a notion of sensitivity.

$$\begin{array}{llll}
\text{(i)} & \mathcal{Y} \subset \mathbb{R}, \mathcal{Z} = \mathbb{R} & d_{\mathcal{Y}} = \epsilon d_{\mathbb{R}} & \lambda_{\epsilon}(z) = \frac{\epsilon}{2} \\
\text{(ii)} & \mathcal{Y} \subset \mathbb{R}^2, \mathcal{Z} = \mathbb{R}^2 & d_{\mathcal{Y}} = \epsilon d_2 & \lambda_{\epsilon}(z) = \frac{\epsilon^2}{2\pi} \\
\text{(iii)} & \mathcal{Y} \subset \mathbb{R}^2, \mathcal{Z} = \mathbb{R}^2 & d_{\mathcal{Y}} = \epsilon d_1 & \lambda_{\epsilon}(z) = \frac{\epsilon^2}{4} \\
\text{(iv)} & \mathcal{Y} = \mathcal{Z} = q[0..k] & d_{\mathcal{Y}} = \epsilon d_{\mathbb{R}} & \lambda_{\epsilon}(z) = \begin{cases} \frac{e^{q\epsilon}}{e^{q\epsilon}+1} & z \in \{0, qk\} \\ \frac{e^{q\epsilon}-1}{e^{q\epsilon}+1} & 0 < z < qk \end{cases}
\end{array}$$

Fig. 2. Instantiations of the Laplace mechanism

Definition 6. Let $\mathcal{Y}, \mathcal{Z}$ be two sets, and let $d_{\mathcal{Y}}$ be a metric on $\mathcal{Y} \cup \mathcal{Z}$. Let $\lambda : \mathcal{Z} \rightarrow [0, \infty)$ be a scaling function such that $D(y)(z) = \lambda(z) e^{-d_{\mathcal{Y}}(y,z)}$ is a pdf for all $y \in \mathcal{Y}$ (i.e. $\int_{\mathcal{Z}} D(y)(z) d\nu(z) = 1$). Then the mechanism $L : \mathcal{Y} \rightarrow \mathcal{P}(\mathcal{Z})$, described by the pdf D , is called a Laplace mechanism from $(\mathcal{Y}, d_{\mathcal{Y}})$ to $\mathcal{Z}$.

Fact 7 ([14]). Any Laplace mechanism from $(\mathcal{Y}, d_{\mathcal{Y}})$ to $\mathcal{Z}$ satisfies $d_{\mathcal{Y}}$ -privacy.

Figure 4.1 provides instantiations of the general definition for various choices of $\mathcal{Y}, \mathcal{Z}$ and $d_{\mathcal{Y}}$ used in the paper, by properly adjusting $\lambda(z)$. The basic case (i) is that of the one-dimensional continuous Laplace mechanism. Similarly, we can define a two-dimensional continuous Laplace mechanism (used in Section 6.2), measuring the distance between points by either the Euclidean (ii) or the Manhattan (iii) metric. In the discrete setting, we obtain the Truncated Geometric mechanism TG_{ϵ} [7], given by (iv), using a quantized set of reals as input. We denote by $q[0..k]$ the set $\{qi \mid i \in 0..k\}$, i.e. the set of $k+1$ quantized reals with step size $q > 0$.

4.2 Mechanisms of Optimal Utility

Answering a query privately is useless if the consumer gets no information about the real answer, thus it is crucial to analyze the mechanism's utility. We consider consumers applying Bayesian inference to map the mechanism's output to a guess that minimizes their expected loss. A consumer is characterized by a prior π on the set of secrets, and a loss function l (assumed to be monotone wrt a metric of reference, which is always $d_{\mathbb{R}}$ for the needs of this paper). The utility $\mathcal{U}(H, \pi, l)$ of a mechanism H for such a consumer is given by the expected loss (under an optimal remap strategy). This is the Bayesian notion of utility [7], but our results can be extended to risk-averse consumers.

A natural question to ask, then, is whether, for a *given query* f , there exists a mechanism that universally (i.e. for all priors and loss functions) provides optimal utility. Let $\mathcal{H}_f(d_{\mathcal{X}})$ be the set of all mechanisms $H : \mathcal{Y} \rightarrow \mathcal{Z}$ (for any $\mathcal{Z}$) such that $H \circ f$ satisfies $d_{\mathcal{X}}$ -privacy. All mechanisms in $\mathcal{H}_f(d_{\mathcal{X}})$ can be used to answer f privately, hence we are interested in the one that maximizes utility.

Definition 7. A mechanism $H \in \mathcal{H}_f(d_{\mathcal{X}})$ is f - $d_{\mathcal{X}}$ -optimal iff $\mathcal{U}(H, \pi, l) \geq \mathcal{U}(H', \pi, l)$ for all $H' \in \mathcal{H}_f(d_{\mathcal{X}})$, all priors π and all loss functions l .

The existence of (universally) optimal mechanisms is far from trivial. For standard differential privacy, a well-known result from [7] states that such a mechanism does exist for *counting* queries, i.e. those of the form “how many users satisfy property P ”.

Theorem 8 ([7]). *Let $\mathcal{Y} = [0..k]$ and let $f : \mathcal{V}^n \rightarrow \mathcal{Y}$ be a counting query. Then the TG_ϵ mechanism with input $\mathcal{Y}$ is f - ϵd_h -optimal for all $\epsilon > 0$.*

On the other hand, a well-known impossibility result [8] states that counting queries are essentially the only ones for which an optimal mechanism exists. This result is based on the concept of the *induced graph* $\sim_f$ of a query $f : \mathcal{V}^n \rightarrow \mathcal{Y}$, defined as: $y \sim_f y'$ iff $\exists x \sim_h x'$ s.t. $f(x) = y, f(x') = y'$.

Theorem 9 ([8]). *Let $f : \mathcal{V}^n \rightarrow \mathcal{Y}$ be a query such that $\sim_f$ is not a path graph. Then no f - ϵd_h -optimal mechanism exists for any $\epsilon < \ln 2$.*

Thus, most interesting queries, e.g. the sum and average, have no optimal mechanisms.

However, the above negative result and the concept of the induced graph are tied to the Hamming metric d_h . This raises the question of whether this special status of counting queries holds for any metric d_x . To answer this question, we give a sufficient condition for showing the optimality of TG_ϵ for an arbitrary query f and metric d_x , based on the concept of uniform sensitivity.

Theorem 10. *Let $\mathcal{Y} = q[0..k]$ and assume that $f : \mathcal{X} \rightarrow \mathcal{Y}$ is uniformly Δ -sensitive wrt $d_x, d_{\mathbb{R}}$. Then the TG_ϵ mechanism with input $\mathcal{Y}$ is f - Δd_x -optimal.*

In the following sections we show that this condition is indeed satisfied by several important queries, including the sum and average, for various metrics of interest.

5 Privacy in Statistical Databases

In this section, we investigate privacy notions in the context of statistical databases, other than the standard differential privacy. In contrast to the Hamming distance, which can be defined independently from the structure of the universe $\mathcal{V}$, we are interested in metrics that depend on the actual values and the distance between them. To this end, we assume that the universe is equipped with a metric d_v , measuring how far apart two values are. When the universe is numeric (i.e. $\mathcal{V} \subset \mathbb{R}$) then $d_v = d_{\mathbb{R}}$ is the natural choice. In the case of null values, we can extend a metric d_v from $\mathcal{V}$ to $\mathcal{V}_\emptyset$ by considering $\emptyset$ to be maximally distant from all other values, that is taking $d_v(\emptyset, v) = d_v(\mathcal{V}), v \in \mathcal{V}$. Note that this construction preserves the maximum distance between values, i.e. $d_v(\mathcal{V}_\emptyset) = d_v(\mathcal{V})$.

The first metric we consider, the normalized Manhattan metric, allows to strengthen differential privacy, obtaining a notion that not only protects the value of an individual, but also offers higher protection to small modifications of a value. Then we relax this metric, to obtain a weaker notion, that only protects the “accuracy” of an individual’s value, but offers higher utility.

5.1 The Normalized Manhattan Metric

Differential privacy provides indistinguishability between databases differing in a single individual, but the level of distinguishability is independent from the actual value in those databases. Consider for example a database with salary information, and two adjacent databases $x \sim_i x'$ ($\sim_i$ denoting that they differ only in the value of the i -th individual) with $x[i] = v, x'[i] = v'$. A differentially private mechanism offers distinguishability level $\varepsilon(x, x') = \epsilon$, independently from v, v' . This means that when $v = 0, v' = 1\text{M}$, the indistinguishability level between x, x' will be the same as in the case $v = 20\text{K}, v' = 20.001\text{K}$.

One might expect, however, to have better protection in the second case, since the change in the individual's data is insignificant. Being insensitive to such small changes seems a reasonable privacy requirement since many queries (e.g. sum, average, etc) are themselves insensitive to small perturbations. The equal treatment of values is particularly problematic when we are obliged to use a “weak” ϵ , due to a high sensitivity. In this case, all values are only guaranteed to be weakly protected, while we could expect that at least close values would still enjoy high protection.

The normalized Manhattan metric $\tilde{d}_1$ expresses exactly this idea. Databases differing in a single value have distance at most 1, but the distance can be substantially smaller for small modifications of values, offering higher protection in those cases. The Manhattan metric d_1 on $\mathcal{V}^n$ and its normalized version $\tilde{d}_1$ are defined as:⁸ $d_1(x, x') = \sum_{i=1}^n d_v(x[i], x'[i])$ and $\tilde{d}_1(x, x') = \frac{d_1(x, x')}{d_v(\mathcal{V})}$. Similarly to differential privacy, we use a scaled version $\epsilon\tilde{d}_1$ of the metric, to properly adjust the distinguishability level.

Concerning the operational characterizations of Section 3, the hiding functions and neighborhoods suitable for this metric are:

$$\begin{aligned} \phi_{i,w} &= x^{[w(x[i])/i]} \text{ for } w : \mathcal{V} \rightarrow \mathcal{V} & N_{i,V}(x) &= \{x^{[v/i]} \mid v \in V\} \\ \Phi_1 &= \{\phi_{i,w} \mid i \in 1..n, w : \mathcal{V} \rightarrow \mathcal{V}\} & \mathcal{N}_1 &= \{N_{i,V}(x) \mid x \in \mathcal{V}^n, i \in 1..n, V \subseteq \mathcal{V}\} \end{aligned}$$

A hiding function $\phi_{i,w}$ replaces the value of individual i by applying an arbitrary substitution of values w (instead of replacing with a fixed value as $\phi_{i,v}$ does). Moreover, for the adversary, knowing $N_{i,V}(x)$ means that he knows the values of all individuals in the database but i , and moreover he knows that the value of i lies within V . Note that $\Phi_h \subset \Phi_1$ and $\mathcal{N}_h \subset \mathcal{N}_1$. We show that $\Phi_1, \mathcal{N}_1$ are “canonical”.

Proposition 2. $\Phi_1, \mathcal{N}_1$ are maximally tight wrt both $d_1, \tilde{d}_1$.

From Theorem 3, we conclude that $\epsilon\tilde{d}_1$ -privacy is equivalent to requiring that the adversary's posterior distributions with or without hiding i 's value should be at most $2\epsilon\tilde{d}_1(\phi_{i,w})$ distant. Since $\tilde{d}_1(\phi_{i,w}) \leq 1$, hiding the individual's value in any way has small effect on the adversary's conclusions. But if i 's value is replaced by one close to it, $\tilde{d}_1(\phi_{i,w})$ can be much lower than 1, meaning that the effect on the adversary's conclusions is even smaller.

⁸ Note that in the differential privacy literature, the d_1 distance is often used on *histograms*. This metric is closely related to the standard d_h distance on $\mathcal{V}^n$ (it depends only on the record counts), and different than d_1 on $\mathcal{V}^n$ which depends on the actual values.

Then, from Theorem 4 we conclude that $\epsilon\tilde{d}_1$ -privacy is equivalent to requiring that, for an informed adversary knowing the value of all individuals but i , and moreover knowing that i 's value lies in V , his conclusions differ from his initial knowledge by at most $\epsilon \frac{d_V(V)}{d_V(\mathcal{V})}$. This difference is at most ϵ , but can be much smaller if values in V are close to each other, meaning that for an adversary who knows i 's value with high accuracy, the gain is even smaller.

Intuitively, $\epsilon\tilde{d}_1$ -privacy offers a stronger notion of privacy than ϵd_h -privacy:

Proposition 3. $\tilde{d}_1 \leq d_h$, thus $\epsilon\tilde{d}_1$ -privacy implies ϵd_h -privacy.

Since distances in $\tilde{d}_1$ can be smaller than those in d_h , the sensitivity of a query wrt $\tilde{d}_1$ is in general greater than the sensitivity wrt d_h , which means that to achieve $\epsilon\tilde{d}_1$ -privacy we need to apply more noise. However, for a general class of queries, it turns out that the two sensitivities coincide.

Definition 8. A query f belongs to the family $\mathcal{C}$ iff $d_{\mathbb{R}}(f(x), f(x')) \leq d_V(x[i], x'[i])$ for all $i \in 1..n$, $x \sim_i x' \in \mathcal{V}^n$, and moreover $\exists x \sim_i x' \in \mathcal{V}^n$ such that $d_{\mathbb{R}}(f(x), f(x')) = d_V(\mathcal{V})$.

Proposition 4. Let $f \in \mathcal{C}$. The sensitivity of f wrt both $d_h, d_{\mathbb{R}}$ and $\tilde{d}_1, d_{\mathbb{R}}$ is $d_V(\mathcal{V})$.

Intuitively, the class $\mathcal{C}$ contains queries for which the sensitivity is obtained for values that are maximally distant. For those queries, using the Truncated Geometric mechanism we can achieve a notion of privacy stronger than differential privacy *using the same amount of noise!*

Results About Some Common Queries. We now focus to some commonly used queries, namely the sum, average and p -percentile queries. Note that other commonly used queries such as the max, min and median queries are specific cases of the p -percentile query. In the following, we assume that the universe is $\mathcal{V} = q[0..k]_{\emptyset}$ with metric $d_{\mathbb{R}}$, and take $\mathcal{X} = \mathcal{V}^n \setminus \{\langle \emptyset, \dots, \emptyset \rangle\}$, that is we exclude the empty database so that the queries can be always defined.

For these queries we obtain two results: first, we show that they belong to the $\mathcal{C}$ family, which means that we can achieve $\epsilon\tilde{d}_1$ -privacy via the TG_{ϵ} mechanism, using the same amount of noise that we would need for standard differential privacy.

Proposition 5. The sum, avg, p -perc queries belong to $\mathcal{C}$.

More interestingly, we can show that the Truncated Geometric mechanism is in fact universally optimal wrt $\tilde{d}_1$ for such queries.

Theorem 11. The sum, avg and p -perc queries are all uniformly qk -sensitive wrt $\tilde{d}_1, d_{\mathbb{R}}$.

Corollary. $TG_{\epsilon/qk}$ is f - $\epsilon\tilde{d}_1$ -optimal for $f \in \{\text{sum, avg, } p\text{-perc}\}$, $\epsilon > 0$.

5.2 The Manhattan Metric

In the previous section, we used the normalized Manhattan metric $\epsilon \tilde{d}_1$, obtaining a strong privacy notion that protects an individual's value, while offering even stronger protection for small changes in an individual's value. This however, requires at least as much noise as standard differential privacy.

On the other hand, there are applications in which a complete protection of an individual's value is not required. This happens, for instance, in situations when the actual value is not sensitive, but knowing it with high accuracy might allow an adversary to identify the individual. Consider for example a database with the individuals' birthday, or the registration date and time to some social network. This information, by itself, might not be considered private, however knowing such information with minute-accuracy could easily allow to identify an individual. In such situations we might wish to protect only the accuracy of the value, thus achieving privacy with less noise and offering more accurate results.

This can be achieved by the Manhattan metric ϵd_1 (without normalization). This metric might assign a level of distinguishability higher than ϵ for adjacent databases, thus the privacy guarantees could be weaker than those of ϵ -differential privacy. However, adjacent databases with small changes in value will be highly protected, thus an adversary cannot infer an individual's value with accuracy.

Similarly to the previous section, we can obtain characterizations of ϵd_1 -privacy using the same hiding functions Φ_1 and neighborhoods $\mathcal{N}_1$. The only difference is that $\epsilon d_1(\phi_{i,w})$ and $\epsilon d_1(N_{i,V})$ can be now higher than ϵ , offering weaker protection. However, when the adversary already knows i 's value with high accuracy, meaning that values in V are close to each other, it is guaranteed that his knowledge will increase by a small factor (possibly even smaller than ϵ), ensuring that he cannot infer the value with even higher accuracy.

Note that the sensitivity of a query can be substantially lower wrt d_1 than wrt d_h . For example, the sum query is 1-sensitive wrt d_1 but qr -sensitive wrt d_h . This means that the noise we need to add could be substantially lower, offering better utility at the expense of lower privacy, but still sufficient for a given application.

Example 1. Consider a database containing the registration date on some social network, expressed as the number of days since Jan 1, 2000. We want to privately release the earliest registration date among individuals satisfying some criteria. A registration date itself is not considered sensitive, however from the result of the query it should be impossible to infer whether a particular individual belongs to that set. Since values can range between 0 and approximately 5,000, the sensitivity of the min query wrt d_h is 5,000, while wrt d_1 it is only 1. By using ϵd_h we protect (up to the intended level ϵ) an individual's registration date within the whole range of values, while by using $\frac{\epsilon}{5} d_1$ we provide the intended protection only within a radius of 5 days. More precisely: in the first case two adjacent databases will always have distinguishability level ϵ , while in the second case such level of protection is guaranteed only if the individual's registration date differs by at most 5 days in the two databases (if they differ more the distinguishability level will increase proportionally). The second case, of course, offers less privacy, but, depending on the application, confusion within 5 days can be enough

Fig. 3. Utility for various values of ϵ

to prevent an individual from being identified. On the other hand, the trade-off with utility can be much more favorable in the second case: In Figure 1 we show the utility of a Laplace mechanism for both metrics, in terms of (α, δ) -usefulness (meaning that the mechanism reports a result within distance α from the real value with probability at least $1 - \delta$).⁹ Clearly, $\frac{\epsilon}{5}d_1$ -privacy gives acceptable utility while ϵd_h -privacy renders the result almost useless.

Finally, the optimality result from the previous section also holds for d_1 .

Theorem 12. *The sum, avg and p -perc queries are all uniformly 1-sensitive wrt $d_1, d_{\mathbb{R}}$.*

Corollary. *TG_ϵ is f - ϵd_1 -optimal for $f \in \{\text{sum, avg, } p\text{-perc}\}$, $\epsilon > 0$.*

6 Privacy in Other Contexts

6.1 Smart Meters

A smart meter is a device that records the consumption of electrical energy at potentially very short time intervals, and transmits the information to the utility provider, thus offering him the capability to monitor consumption accurately and almost in real-time.

The Problem. Although smart meters can help improving energy management, they create serious privacy threats: By analyzing accurate consumption data, thanks to appliance signature libraries it is possible to identify which electric devices are being used [15]. It has even been shown that, depending on the granularity of measurement and the resolution of data, it is possible to deduce what TV channels, and which movies are being watched [16].

Several papers addressed the privacy problems of smart metering in the recent past. The solution proposed in [17] is based on the use of techniques of (standard) differential privacy in order to send sanitized sums of the readings over some period of time (e.g. an hour, a day, a month) to the service provider. Since this solution is tailored to the use of smart metering for billing purposes, the noise added is assumed to be positive.

⁹ Using Bayesian utility leads to similar results.

The Model. For the sake of generality, we assume here that the noise could be of any kind (not necessarily positive). We can regard the readings over the period $[1..n]$ as a tuple $x \in \mathcal{V}^n$, so that $x[i]$ represents the reading at the time i . Since [17] uses the standard differential privacy framework, the distinguishability metric on these tuples is assumed to be the Hamming distance, and therefore the privacy mechanism is tuned to protect the value of $x[i]$, regardless of whether the variation of this value is small or large. However, the solution proposed in [17] is general and can be adapted to a different distinguishability metric.

We argue that for the case of smart meters, the problem that derives from the extreme accuracy of the readings can be addressed with limited noise by adopting a metric that is sensitive also to the distance between values, and not only to the change of the value for a reading $x[i]$. The reason is the same as illustrated in previous section: if we want to protect small variations in the reading of $x[i]$, it is not a good idea to tune the sensitivity on the difference between the extremes values, because we would end up introducing a lot of noise. In fact, the experiments in [16] are performed on actual smart meters that are in the process of being deployed. These meters send readings to the service provider every 2 seconds. The solution proposed in [17] offers good privacy guarantees by completely protecting each measurement. However, such a definition is too strong if reporting values at short intervals is a requirement. With standard differential privacy, we cannot hope to fully protect each measurement without introducing too much noise. On the other hand, using a more relaxed metric, we can at least provide a meaningful privacy guarantee by protecting the accuracy of the values. Some privacy will still be lost, but the attacks described above where the individual's behaviour is completely disclosed, will be prevented.

The Manhattan distance d_1 on $\mathcal{V}^n$, however, is not suitable to model the privacy problem we have here: in fact d_1 is suitable to protect an individual $x[i]$ and its value, while here we want to protect *all the values at the same time*. This is because the adversary, i.e., the service provider, already knows *an approximation of all values*. Note the difference from the case of Section 5: there, the canonical adversary knows all exact values except $x[i]$, and for $x[i]$ he only knows an approximate value. (In the case of standard differential privacy, the canonical adversary knows all values except $x[i]$, and for $x[i]$ he does not even know an approximate value.)

The suitable distance, in this case, is the maximum distance between components, d_∞ . In fact, we should consider x, x' “indistinguishable enough” (i.e. $d(x, x') \leq \delta$, for a certain δ) if and only if for each component i , $x[i], x'[i]$ are “indistinguishable enough” (i.e. $d(x[i], x'[i]) \leq \delta$, for the same δ). It is easy to see that the only distance that satisfies this property is $d(x, x') = d_\infty(x, x') = \max_i d_v(x[i], x'[i])$.

Example 2. We illustrate the application our method to distort the digital signature of a tv program. The grey line in Fig. 4(a) represents the energy consumption of the first 5 minutes of Star Trek 11 [15]. The black line is (the approximation of) the signature produced by a smart meter that reports the true readings every 10 seconds (the samples are represented by the dots). The blue and the magenta dots in 4(b) are obtained by adding laplacian noise to the true readings, with ϵ values .1 and .5 respectively. As we can see, especially in the case of $\epsilon = .5$, the signature is not recognizable.

Fig. 4. Digital signature of a tv program (a) and its noisy reporting (b).

Concerning the characterization results, we use hiding functions substituting the value of all readings. Moreover, we use neighborhoods modelling an adversary that knows all readings with some accuracy, i.e. knows that each reading i lies within V_i .

$$\begin{aligned}\Phi_\infty &= \{\phi_{1,w_1} \circ \dots \circ \phi_{n,w_n} \mid w_i : \mathcal{V} \rightarrow \mathcal{V} \forall i \in 1..n\} \\ N_{\{V_i\}} &= \{\langle v_1, \dots, v_n \rangle \mid v_i \in V_i, i \in 1..n\} \\ \mathcal{N}_\infty &= \{N_{\{V_i\}} \mid V_i \subseteq \mathcal{V}, i \in 1..n\}\end{aligned}$$

We can show that $\Phi_\infty, \mathcal{N}_\infty$ are maximally tight.

Finally, we show that TG_ϵ is universally optimal for avg and p -perc.

Theorem 13. *The avg and p -perc queries are both uniformly 1-sensitive wrt $d_\infty, d_{\mathbb{R}}$.*

Corollary. *TG_ϵ is f - ϵd_∞ -optimal for $f \in \{\text{avg}, p\text{-perc}\}$, $\epsilon > 0$.*

6.2 Geolocation

In this subsection we briefly describe an application of our framework to privacy-aware location-based systems. We refer to [18] for more details.

Privacy notions have been already studied in previous works. Some of these works [19,20,21] propose the use of the *expectation of distance error* of the attacker as the way to quantify the privacy offered by a mechanism. Others works [22,23,24] rely on the well-known concept of k -anonymity. The notion of *relevance* is also used to measure location privacy in [25]. A strong advantage of the use of d -privacy as privacy notion is that it abstracts from the side-knowledge of the attacker.

The Problem. In several situations it is desirable to know the location of an individual or a group of individuals in order to provide a service. For instance: In census-based statistics, to determine the population density in certain areas, in transportation industry, to estimate the average number of people who need to travel between two given stations, and in smartphone applications, to obtain points of interest nearby such as restaurants.

Due to privacy concerns, an individual may refuse to disclose his exact location to the service provider. Nevertheless, he may be willing to reveal approximate location information. It is worth noting that for several location-based systems it is usually enough to obtain an approximate location to be able to provide an accurate service. Note however, that in order to guarantee a non-negligible level of privacy, the random location cannot be generated naively. Therefore, if we want to develop a method to randomize location coordinates, we have to understand what kind of privacy the user expects to have, and how much information he is willing to reveal.

The Model. In this scenario, the privacy level depends on the accuracy with which an attacker can guess an individual's location from the reported one. We will therefore aim for a distance-dependent notion of privacy, requiring points that are close in distance to each other to be *indistinguishable* from the attacker's point of view. Our method will still allow the service provider to distinguish between points that are far from each other.

We consider the problem of geolocation on the Euclidean plane, which is a good approximation of the Earth surface when the area is not "too large". In this scenario, possible locations of an individual will be modeled with a set $\mathcal{X} \subseteq \mathbb{R}^2$, and possible reported values will be represented by a set $\mathcal{Z} \subseteq \mathbb{R}^2$. The metric d_x used in this context will be the Euclidean distance d_2 .

Concerning the characterizations of Section 3, any function $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ can be used as a hiding function. Moreover, a neighborhood can be any region $N \subseteq \mathbb{R}^2$, modelling an informed adversary who knows that the user is located within N . Hence we take $\Phi_2 = \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $\mathcal{N}_2 = 2^{\mathbb{R}^2}$, both of which are maximally tight wrt d_2 .

In order to obtain a mechanism which satisfies ϵd_2 -privacy, we can use the *Laplace* mechanism L_ϵ on $\mathbb{R}^2$ mentioned in Section 4.1, that is, the one described by the pdf $D_\epsilon(x)(z) = \frac{\epsilon^2}{2\pi} e^{-\epsilon d_2(x,z)}$ $x, z \in \mathbb{R}^2$. The results in Section 4.1 ensure that such mechanism satisfies ϵd_2 -privacy.

7 Conclusion

Starting from the observation that differential privacy requires that the distinguishability of two databases depends on their Hamming distance, we have explored the consequences of extending this principle to arbitrary metrics. In this way we have obtained a rich framework suitable to model a large variety of privacy problems, and in domains other than statistical databases. Furthermore, even in statistical databases applications, whenever the privacy concern is related to disclosing small variations in the values of the individuals (rather than large ones), then our framework allows a more precise calibration of the noise necessary for achieving the intended level of privacy, and this results, in general, in a better utility than the one achievable under the constraint of standard differential privacy. We have investigated the trade-off between privacy and utility in this extended setting, and it turns out changing the metric has considerable implications on the existence of universally optimal mechanisms. In particular, for the Manhattan distance, the normalized Manhattan distance, and the max distance it is possible to define universally optimal mechanisms for several common queries like the

sum, the average, and the percentile. This contrast sharply with the case of standard differential privacy, where universally optimal mechanisms exist only for counting queries. Finally, we have shown the applicability of our framework to various privacy problems in different domains, including smart meters and geolocation.

References

1. Dwork, C.: Differential privacy. In: Proc. of ICALP. Volume 4052 of LNCS., Springer (2006) 1–12
2. Dwork, C., Mcsherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: Proc. of TCC. Volume 3876 of LNCS., Springer (2006) 265–284
3. Reed, J., Pierce, B.C.: Distance makes the types grow stronger: a calculus for differential privacy. In: Proc. of ICFP, ACM (2010) 157–168
4. Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: Proc. of S&P, IEEE (2009) 173–187
5. Machanavajjhala, A., Kifer, D., Abowd, J.M., Gehrke, J., Vilhuber, L.: Privacy: Theory meets practice on the map. In: Proc. of ICDE, IEEE (2008) 277–286
6. Ganta, S.R., Kasiviswanathan, S.P., Smith, A.: Composition attacks and auxiliary information in data privacy. In: Proc. of KDD, ACM (2008) 265–273
7. Ghosh, A., Roughgarden, T., Sundararajan, M.: Universally utility-maximizing privacy mechanisms. In: Proc. of STOC, ACM (2009) 351–360
8. Brenner, H., Nissim, K.: Impossibility of differentially private universally optimal mechanisms. In: Proc. of FOCS, IEEE (2010) 71–80
9. Nissim, K., Raskhodnikova, S., Smith, A.: Smooth sensitivity and sampling in private data analysis. In: Proc. of STOC, ACM (2007) 75–84
10. Barthe, G., Köpf, B., Olmedo, F., Béguelin, S.Z.: Probabilistic relational reasoning for differential privacy. In: Proc. of POPL, ACM (2012)
11. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.S.: Fairness through awareness. In: Proc. of ITCS, ACM (2012) 214–226
12. Chatzikokolakis, K., Andrés, Miguel, E., Bordenabe, Nicolás, E., Palamidessi, C.: Broadening the scope of Differential Privacy using metrics. Tech. rep., INRIA (2012) To appear in PETS 2013. Available at: <http://hal.inria.fr/hal-00767210>.
13. Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., Naor, M.: Our data, ourselves: Privacy via distributed noise generation. In: Proc. of EUROCRYPT. Volume 4004 of LNCS., Springer (2006) 486–503
14. McSherry, F., Talwar, K.: Mechanism design via differential privacy. In: Proc. of FOCS, IEEE (2007) 94–103
15. Lam, H., Fung, G., Lee, W.: A novel method to construct taxonomy electrical appliances based on load signatures. *IEEE Trans. on Consumer Electronics* **53**(4) (2007) 653–660
16. Greveler, U., Justus, B., Loehr, D.: Multimedia content identification through smart meter power use profiles. In: CPDP. (2012)
17. Danezis, G., Kohlweiss, M., Rial, A.: Differentially private billing with rebates. *IACR Cryptology ePrint Archive* **2011** (2011) 134
18. Andrés, M., Bordenabe, N., Chatzikokolakis, K., Palamidessi, C.: Geo-indistinguishability: Differential privacy for location-based systems. *CoRR* **abs/1212.1984** (2012)
19. Shokri, R., Theodorakopoulos, G., Boudec, J.Y.L., Hubaux, J.P.: Quantifying location privacy. In: Proc. of S&P, IEEE (2011) 247–262
20. Shokri, R., Theodorakopoulos, G., Troncoso, C., Hubaux, J.P., Boudec, J.Y.L.: Protecting location privacy: optimal strategy against localization attacks. In: Proc. of CCS, ACM (2012) 617–627

21. Hoh, B., Gruteser, M.: Protecting location privacy through path confusion. In: SecureComm, IEEE (2005) 194–205
22. Kido, H., Yanagisawa, Y., Satoh, T.: Protection of location privacy using dummies for location-based services. In: Proc. of ICDE Workshops. (2005) 1248
23. Shankar, P., Ganapathy, V., Iftode, L.: Privately querying location-based services with sybil-query. In: Proc. of UbiComp, ACM (2009) 31–40
24. Duckham, M., Kulik, L.: A formal model of obfuscation and negotiation for location privacy. In: Proc. of PERVASIVE. Volume 3468 of LNCS., Springer (2005) 152–170
25. Ardagna, C.A., Cremonini, M., Damiani, E., di Vimercati, S.D.C., Samarati, P.: Location privacy protection through obfuscation-based techniques. In: Proc. of DAS. Volume 4602 of LNCS., Springer (2007) 47–60

A Proofs of Section 3

Proposition 1. *If $d_{\mathcal{X}} \leq d_{\mathcal{X}'}'$ (point-wise) then $d_{\mathcal{X}}$ -privacy implies $d_{\mathcal{X}'}'$ -privacy.*

Proof. Trivial, since $d_{\mathcal{P}}(K(x), K(x')) \leq d_{\mathcal{X}}(x, x') \leq d_{\mathcal{X}'}'(x, x')$. $\square$

The following simple lemma states that bounding $d_{\mathcal{P}}$ is equivalent to the usual formulation of bounding the ratio between probabilities.

Lemma 1. *Let μ_1, μ_2 be probability measures on $\mathcal{Z}$. Then*

$$d_{\mathcal{P}}(\mu_1, \mu_2) \leq b \quad \text{iff} \quad \forall Z \in \mathcal{F}_{\mathcal{Z}} : e^{-b} \mu_2(Z) \leq \mu_1(Z) \leq e^b \mu_2(Z)$$

Proof. iff) We have $|\ln \frac{\mu_1(Z)}{\mu_2(Z)}| \leq d_{\mathcal{P}}(\mu_1, \mu_2) \leq b$, hence $-b \leq \ln \frac{\mu_1(Z)}{\mu_2(Z)} \leq b$, which implies $e^{-b} \leq \frac{\mu_1(Z)}{\mu_2(Z)} \leq e^b$. if) We have that $|\ln \frac{\mu_1(Z)}{\mu_2(Z)}|$ is bounded from above by b , but $d_{\mathcal{P}}(\mu_1, \mu_2)$ is the least such bound hence $d_{\mathcal{P}}(\mu_1, \mu_2) \leq b$. $\square$

The following Lemma shows the usefulness of *tight* chains.

Lemma 2. *Let $x_1, \dots, x_n$ be a tight chain. If K satisfies $d_{\mathcal{X}}$ -privacy on all adjacent elements of the chain, then it also satisfies it for x_1, x_n . That is*

$$d_{\mathcal{P}}(K(x_i), K(x_{i+1})) \leq d_{\mathcal{X}}(x_i, x_{i+1}) \quad \forall 1 \leq i < n$$

implies $d_{\mathcal{P}}(K(x_1), K(x_n)) \leq d_{\mathcal{X}}(x_1, x_n)$.

Proof. Using the fact that $d_{\mathcal{P}}$ is itself a metric, we have

$$\begin{aligned} d_{\mathcal{P}}(K(x_1), K(x_n)) &\leq \sum_{i=1}^{n-1} d_{\mathcal{P}}(K(x_i), K(x_{i+1})) && \text{triangle ineq. for } d_{\mathcal{P}} \\ &\leq \sum_{i=1}^{n-1} d_{\mathcal{X}}(x_i, x_{i+1}) && \text{hypothesis} \\ &= d_{\mathcal{X}}(x_1, x_n) && \text{tightness} \end{aligned}$$

$\square$

Theorem 3. Let Φ be a set of hiding functions. If K satisfies $d_{\mathcal{X}}$ -privacy then for all $\phi \in \Phi$, all priors π on $\mathcal{X}$, and all $Z \in \mathcal{F}_{\mathcal{Z}}$:

$$d_{\mathcal{P}}(\sigma_1, \sigma_2) \leq 2d_{\mathcal{X}}(\phi) \quad \text{where } \sigma_1 = \mathbf{Bayes}(\pi, K, Z) \text{ and } \sigma_2 = \mathbf{Bayes}(\pi, K \circ \phi, Z)$$

If Φ is maximally tight then the converse also holds.

Proof. Assume that K satisfies $d_{\mathcal{X}}$ -privacy and let π be a prior, $\phi \in \Phi$ and $Z \in \mathcal{F}_{\mathcal{Z}}$. We need to show that

$$\forall x \in \mathcal{X} : e^{-2d_{\mathcal{X}}(\phi)}\sigma_1(x) \leq \sigma_2(x) \leq e^{2d_{\mathcal{X}}(\phi)}\sigma_1(x)$$

(then conclude by applying Lemma 1). Let $x \in \mathcal{X}$, we have:

$$\begin{aligned} \sigma_2(x) &= \frac{(K \circ \phi)(x)(Z)\pi(x)}{\sum_{x' \in \mathcal{X}} (K \circ \phi)(x')(Z)\pi(x')} && \text{def. of Bayes} \\ &= \frac{K(\phi(x))(Z)\pi(x)}{\sum_{x' \in \mathcal{X}} K(\phi(x'))(Z)\pi(x')} \\ &\leq \frac{e^{d_{\mathcal{X}}(x, \phi(x))}K(x)(Z)\pi(x)}{\sum_{x' \in \mathcal{X}} e^{-d_{\mathcal{X}}(x', \phi(x'))}K(x')(Z)\pi(x')} && d_{\mathcal{X}}\text{-privacy} \\ &\leq \frac{e^{d_{\mathcal{X}}(\phi)}K(x)\pi(x)(Z)}{e^{-d_{\mathcal{X}}(\phi)}\sum_{x' \in \mathcal{X}} K(x')(Z)\pi(x')} && d_{\mathcal{X}}(x, \phi(x)) \leq d_{\mathcal{X}}(\phi) \\ &\leq e^{2d_{\mathcal{X}}(\phi)}\sigma_1(x) && \text{def. of Bayes} \end{aligned}$$

and symmetrically for $\sigma_2(x) \geq e^{-2d_{\mathcal{X}}(\phi)}\sigma_1(x)$.

For the opposite direction, assume that Φ is maximally tight (Def 2), that $d_{\mathcal{P}}(\sigma_1, \sigma_2) \leq 2d_{\mathcal{X}}(\phi)$ holds for all π, ϕ, Z , but $d_{\mathcal{X}}$ -privacy is violated for some $x, x' \in \mathcal{X}$. From Def 2, there exist a tight maximal Φ -chain $\tilde{x}$ from x to x' . Then from Lemma 2, we get that $d_{\mathcal{X}}$ -privacy is also violated for some adjacent x_i, x_{i+1} in the chain, that is:

$$K(x_i)(Z) > e^{d_{\mathcal{X}}(x_i, x_{i+1})}K(x_{i+1})(Z) \quad \text{for some } Z \quad (2)$$

We fix Z to the one above. Since $\tilde{x}$ is a maximal Φ -chain, there exists $\phi \in \Phi$ such that $\phi(x_i) = x_{i+1}, \phi(x_{i+1}) = x_i$ and $d_{\mathcal{X}}(x_i, x_{i+1}) = d_{\mathcal{X}}(\phi)$. Fixing this ϕ , we define a function $f : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ as follows:

$$f(\pi) = \frac{\sum_{x' \in \mathcal{X}} K(x')(Z)\pi(x')}{\sum_{x' \in \mathcal{X}} K(\phi(x'))(Z)\pi(x')}$$

Let $\delta(x)$ denote the Dirac measure assigning probability 1 to x , from (2) we have that

$$\begin{aligned} f(\delta(x_i)) &= \frac{K(x_i)(Z)}{K(x_{i+1})(Z)} > e^{-d_{\mathcal{X}}(x_i, x_{i+1})} \\ f(\delta(x_{i+1})) &= \frac{K(x_{i+1})(Z)}{K(x_i)(Z)} < e^{-d_{\mathcal{X}}(x_i, x_{i+1})} \end{aligned}$$

From the continuity of f on the line between $\delta(x_i)$ and $\delta(x_{i+1})$, there exists a prior $\pi = t\delta(x_i) + (1-t)\delta(x_{i+1})$, $t \in (0, 1)$, such that $f(\pi) = e^{-d_{\mathcal{X}}(x_i, x_{i+1})}$. Note that since π is distinct from $\delta(x_i), \delta(x_{i+1})$, it holds that $\pi(x_i) > 0, \pi(x_{i+1}) > 0$. By applying the hypothesis for this π , we get

$$\begin{aligned}
d_{\mathcal{P}}(\sigma_1, \sigma_2) &\leq 2d_{\mathcal{X}}(\phi) && \Rightarrow \\
\sigma_1(x_i) &\leq e^{2d_{\mathcal{X}}(\phi)} \sigma_2(x_i) && (\text{Lemma 1}) \Rightarrow \\
\frac{K(x_i)(Z)\pi(x_i)}{\sum_{x' \in \mathcal{X}} K(x')(Z)\pi(x')} &\leq e^{2d_{\mathcal{X}}(\phi)} \frac{K(\phi(x_i))(Z)\pi(x_i)}{\sum_{x' \in \mathcal{X}} K(\phi(x'))(Z)\pi(x')} && (\text{Def. of } \sigma_1, \sigma_2) \Rightarrow \\
K(x_i)(Z) &\leq e^{2d_{\mathcal{X}}(\phi)} f(\pi) K(\phi(x_i))(Z) && (\pi(x_i) > 0) \Rightarrow \\
K(x_i)(Z) &\leq e^{d_{\mathcal{X}}(x_i, x_{i+1})} K(x_{i+1})(Z)
\end{aligned}$$

which contradicts (2). $\square$

Theorem 4. Let $\mathcal{N} \subseteq 2^{\mathcal{X}}$. If K satisfies $d_{\mathcal{X}}$ -privacy then for all $N \in \mathcal{N}$, all priors π on $\mathcal{X}$, and all $Z \in \mathcal{F}_{\mathcal{Z}}$:

$$d_{\mathcal{P}}(\pi_{|N}, \sigma_{|N}) \leq d_{\mathcal{X}}(N) \quad \text{where} \quad \sigma = \mathbf{Bayes}(\pi, K, Z)$$

If $\mathcal{N}$ is maximally tight then the converse also holds.

Proof. Assume that K satisfies $d_{\mathcal{X}}$ -privacy. We fix some $N \in \mathcal{N}, \pi \in \mathcal{P}(\mathcal{X}), Z \in \mathcal{F}_{\mathcal{Z}}$ and let $\sigma = \mathbf{Bayes}(\pi, K, Z)$. Note that $\pi_{|N}, \sigma_{|N}$ are distributions on N . From Lemma 1 we need to show that

$$e^{-d_{\mathcal{X}}(N)} \pi_{|N}(x) \leq \sigma_{|N}(x) \leq e^{d_{\mathcal{X}}(N)} \pi_{|N}(x) \quad \forall x \in N$$

Fixing some $x \in N$, we have:

$$\begin{aligned}
\sigma_{|N}(x) &= \sigma(x|N) && \text{def. of } \sigma_{|N} \\
&= \frac{\sigma(x)}{\sum_{x' \in N} \sigma(x')} \\
&= \frac{\pi(x)K(x)(Z)}{\sum_{x' \in N} \pi(x')K(x')(Z)} && \text{def. of Bayes} \\
&\leq \frac{\pi(x)K(x)(Z)}{\sum_{x' \in N} \pi(x')e^{-d_{\mathcal{X}}(x, x')}K(x')(Z)} && d_{\mathcal{X}}\text{-privacy} \\
&\leq e^{d_{\mathcal{X}}(N)} \frac{\pi(x)}{\sum_{x' \in N} \pi(x')} && d_{\mathcal{X}}(x, x') \leq d_{\mathcal{X}}(N) \\
&= e^{d_{\mathcal{X}}(N)} \pi_{|N}(x)
\end{aligned}$$

and symmetrically for $\sigma_{|N}(x) \geq e^{-d_{\mathcal{X}}(N)} \pi_{|N}(x)$.

For the opposite direction, assume that $\mathcal{N}$ is maximally tight (Def 3) but $d_{\mathcal{X}}$ -privacy is violated for some $x, x' \in \mathcal{X}$. From Def 3, there exist a tight $\mathcal{N}$ -chain $\tilde{x}$ from x to x' .

Then from Lemma 2, we get that $d_{\mathcal{X}}$ -privacy is also violated for some adjacent x_i, x_{i+1} in the chain, that is:

$$K(x_i)(Z) > e^{d_{\mathcal{X}}(x_i, x_{i+1})} K(x_{i+1})(Z) \quad \text{for some } Z \quad (3)$$

Since $\tilde{x}$ is an $\mathcal{N}$ -chain, there exist $N \in \mathcal{N}$ such that $\{x_i, x_{i+1}\} \subseteq N$ and $d_{\mathcal{X}}(x_i, x_{i+1}) = d_{\mathcal{X}}(N)$. We define a prior distribution $\pi_t(x)$ as

$$\pi_t(x) = \begin{cases} t & x = x_i \\ 1 - t & x = x_{i+1} \\ 0 & \text{otherwise} \end{cases}$$

Using that prior for $t > 0$, we fix some $Z \in \mathcal{F}_Z$ and let $\sigma_t = \text{Bayes}(\pi_t, K, Z)$. We have

$$\begin{aligned} \sigma_{t|N}(x_i) &\leq e^{d_{\mathcal{X}}(N)} \pi_{t|N}(x_i) && (\text{hypoth., Lemma 1}) \Rightarrow \\ \sigma_t(x_i) &\leq e^{d_{\mathcal{X}}(N)} \pi_t(x_i) && (\pi_t(N) = \sigma_t(N) = 1) \Rightarrow \\ \frac{\pi_t(x_i) K(x_i)(Z)}{\sum_{x' \in \mathcal{X}} \pi_t(x') K(x')(Z)} &\leq e^{d_{\mathcal{X}}(N)} \pi_t(x_i) && (\text{def. of Bayes}) \Rightarrow \\ \frac{t K(x_i)(Z)}{t K(x_i)(Z) + (1-t) K(x_{i+1})(Z)} &\leq e^{d_{\mathcal{X}}(N)} t && (\text{def. of } \pi_t) \Rightarrow \\ \frac{K(x_i)(Z)}{t K(x_i)(Z) + (1-t) K(x_{i+1})(Z)} &\leq e^{d_{\mathcal{X}}(N)} && (t > 0) \Rightarrow \\ \frac{K(x_i)(Z)}{t K(x_i)(Z) + (1-t) K(x_{i+1})(Z)} &\leq e^{d_{\mathcal{X}}(x_i, x_{i+1})} && (d_{\mathcal{X}}(x_i, x_{i+1}) = d_{\mathcal{X}}(N)) \end{aligned}$$

The above inequality holds for all $t > 0$. Finally, taking the $\lim_{t \rightarrow 0}$ on both sides we get

$$K(x_i)(Z) \leq e^{d_{\mathcal{X}}(x_i, x_{i+1})} K(x_{i+1})(Z)$$

which is a contradiction of (3). $\square$

B Proofs of Section 4

Fact 5. Assume that f is Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$ and H satisfies $d_{\mathcal{Y}}$ -privacy. Then $H \circ f$ satisfies $\Delta d_{\mathcal{X}}$ -privacy.

Proof. Assume that H satisfies $d_{\mathcal{Y}}$ -privacy and let $x, x' \in \mathcal{X}$. We have:

$$\begin{aligned} d_{\mathcal{P}}((H \circ f)(x), (H \circ f)(x')) &= d_{\mathcal{P}}(H(f(x)), H(f(x'))) \\ &\leq d_{\mathcal{Y}}(f(x), f(x')) && d_{\mathcal{Y}}\text{-privacy} \\ &\leq \Delta d_{\mathcal{X}}(x, x') && \Delta\text{-sensitivity} \end{aligned}$$

$\square$

Theorem 6. Assume that f is uniformly Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$. Then H satisfies $d_{\mathcal{Y}}$ -privacy if and only if $H \circ f$ satisfies $\Delta d_{\mathcal{X}}$ -privacy.

Proof. The only if part is Fact 5. For the if part, fix some $y, y' \in \mathcal{Y}$ and let $y_1, \dots, y_n$ be the tight Δ -expansive chain from y to y' guaranteed to exist by the definition of uniform Δ -sensitivity. Then, for all $1 \leq i < n$, since y_i, y_{i+1} are Δ -expansive, there exist

$$x \in f^{-1}(y_i), x' \in f^{-1}(y_{i+1}) \quad \text{such that} \quad d_{\mathcal{Y}}(f(x), f(x')) = \Delta d_{\mathcal{X}}(x, x')$$

Hence

$$\begin{aligned} d_{\mathcal{P}}(H(y_i), H(y_{i+1})) &= d_{\mathcal{P}}(H(f(x)), H(f(x'))) \\ &\leq \Delta d_{\mathcal{X}}(x, x') && \Delta d_{\mathcal{X}}\text{-privacy of } H \circ f \\ &= d_{\mathcal{Y}}(y_i, y_{i+1}) \end{aligned}$$

So H satisfies $d_{\mathcal{Y}}$ -privacy for all adjacent elements in the chain, hence from Lemma 2 it also satisfies it for y, y' . $\square$

The following result is standard, we provide the proof for completeness.

Fact 7 ([14]). Any Laplace mechanism from $(\mathcal{Y}, d_{\mathcal{Y}})$ to $\mathcal{Z}$ satisfies $d_{\mathcal{Y}}$ -privacy.

Proof. For the pdf describing the mechanism we have:

$$\begin{aligned} D(y)(z) &= \lambda(z) e^{-d_{\mathcal{Y}}(y, z)} \\ &\leq \lambda(z) e^{-(d_{\mathcal{Y}}(y', z) - d_{\mathcal{Y}}(y, y'))} && \text{triangle ineq.} \\ &= e^{d_{\mathcal{Y}}(y, y')} \lambda(z) e^{-d_{\mathcal{Y}}(y', z)} \\ &= e^{d_{\mathcal{Y}}(y, y')} D(y')(z) \end{aligned}$$

for all $y, y' \in \mathcal{Y}, z \in \mathcal{Z}$. The above inequality can be directly extended from the pdf to the measures, thus we conclude that the mechanism satisfies $d_{\mathcal{Y}}$ -privacy. $\square$

B.1 Proofs of Section 4.2

To prove our sufficient condition for optimality (Theorem 10), we introduce the concept of a mechanism being optimal not wrt a specific query, but wrt *the metric of its input domain*. Let $\mathcal{H}(d_{\mathcal{Y}})$ be the set of all mechanisms $H : \mathcal{Y} \rightarrow \mathcal{Z}$ (for any $\mathcal{Z}$) satisfying $d_{\mathcal{Y}}$ -privacy.

Definition 9. A mechanism $H \in \mathcal{H}(d_{\mathcal{Y}})$ is $d_{\mathcal{Y}}$ -optimal iff $\mathcal{U}(H, \pi, l) \geq \mathcal{U}(H', \pi, l)$ for all $H' \in \mathcal{H}(d_{\mathcal{Y}})$, all priors π and all loss functions l .

Notice the difference between $d_{\mathcal{Y}}$ -optimal and f - $d_{\mathcal{X}}$ -optimal; the latter refers to a specific query. The two notions can be related in the case of uniformly sensitive queries by the following result:

Proposition 6. Assume that f is uniformly Δ -sensitive wrt d_x, d_y . Then H is f - Δd_x -optimal iff it is d_y -optimal.

Proof. From uniform Δ -sensitivity and Theorem 6, we get that $\mathcal{H}(d_y) = \mathcal{H}_f(\Delta d_x)$. Then the result follows directly from the definition of optimality (Definitions 7,9). $\square$

The importance of the induced graph $\sim_f$ in optimality results comes from the fact that f is always uniformly sensitive wrt the metric induced by $\sim_f$.

Proposition 7. Let f be a query with induced graph $\sim_f$, and let d_f be the metric induced by $\sim_f$. Then f is uniformly 1-sensitive wrt d_h, d_f .

Proof. Let $x, x' \in \mathcal{X}$ and let $n = d_h(x, x')$. We first need to show that f is 1-sensitive wrt d_h, d_f , that is $d_f(f(x), f(x')) \leq n$. Since d_h is induced by $\sim_h$, there exists a $\sim_h$ -path $x_1 \dots, x_n$ such that $x = x_1, x' = x_n$. By definition of $\sim_f$ we have that $f(x_i) \sim_f f(x_{i+1})$, thus $f(x_1), \dots, f(x_n)$ is a $\sim_f$ -path of length n from $f(x)$ to $f(x')$. Since $d_f(f(x), f(x'))$ is the length of the shortest such path, we have that $d_f(f(x), f(x')) \leq n$.

For the “uniformly” part, let $y, y' \in \mathcal{Y}$ and $n = d_f(y, y')$. We need to show that there exists a tight and 1-expansive chain from y to y' .

Since d_f is induced by $\sim_f$, there exist a $\sim_f$ -path $\tilde{y} = y_1, \dots, y_n$ such that $y = y_1, y' = y_n$. This implies that $d_f(y_i, y_{i+1}) = 1$ and thus $d_f(\tilde{y}) = n = d_f(y, y')$ so the chain is tight.

Moreover, from the definition of $\sim_f$ we have that there exist $x \sim_h x'$ such that $f(x) = y_i, f(x') = y_{i+1}$, so $d_h(x, x') = 1$ which means that y_i, y_{i+1} are 1-expansive, and this happens for all $1 \leq i < n$ so the chain is 1-expansive. $\square$

We can now show the optimality of TG_ϵ with input $q[0..k]$ wrt the $\epsilon d_{\mathbb{R}}$ metric, independently from any query.

Proposition 8. Let $\mathcal{Y}_q = q[0..k]$. The TG_ϵ mechanism with input $\mathcal{Y}_q$ is $\epsilon d_{\mathbb{R}}$ -optimal for all $\epsilon > 0$.

Proof. Fix $\mathcal{Y} = 0..k$ and $\mathcal{Y}_q = q[0..k]$ for some $k \in \mathbb{N}, q > 0$, and let $TG_\epsilon(\mathcal{Y}), TG_\epsilon(\mathcal{Y}_q)$ denote the Truncated Geometric mechanisms with input $\mathcal{Y}, \mathcal{Y}_q$ respectively.

From Theorem 8 we know that $TG_\epsilon(\mathcal{Y})$ is f - ϵd_h -optimal when f is a counting query. For counting queries, d_f (the metric that corresponds to their induced graph) and $d_{\mathbb{R}}$ coincide, thus from Prop 7 we get that f is uniformly 1-sensitive wrt $d_h, d_{\mathbb{R}}$. Then from Prop 6 we have that $TG_\epsilon(\mathcal{Y})$ is $\epsilon d_{\mathbb{R}}$ -optimal. This mechanism has pdf

$$D_\epsilon(y)(z) = \lambda_\epsilon(z) e^{-\epsilon d_{\mathbb{R}}(y, z)} \quad \lambda_\epsilon(z) = \begin{cases} \frac{\epsilon^\epsilon}{\epsilon^\epsilon + 1} & z \in \{0, k\} \\ \frac{\epsilon^\epsilon - 1}{\epsilon^\epsilon + 1} & 0 < z < k \end{cases}$$

We now show that $TG_\epsilon(\mathcal{Y}_q)$ is also $\epsilon d_{\mathbb{R}}$ -optimal. The metric spaces $(\mathcal{Y}, \epsilon d_{\mathbb{R}})$ and $(\mathcal{Y}_q, \epsilon \frac{1}{q} d_{\mathbb{R}})$ are isometric. So we can obtain a mechanism D'_ϵ with input $\mathcal{Y}_q$ by replacing i with qi and $\epsilon d_{\mathbb{R}}$ with $\epsilon \frac{1}{q} d_{\mathbb{R}}$. The pdf of this mechanism is:

$$D'_\epsilon(y)(z) = \lambda_\epsilon(z) e^{-\epsilon \frac{1}{q} d_{\mathbb{R}}(y, z)} \quad \lambda_\epsilon(z) = \begin{cases} \frac{\epsilon^\epsilon}{\epsilon^\epsilon + 1} & z \in \{0, qk\} \\ \frac{\epsilon^\epsilon - 1}{\epsilon^\epsilon + 1} & 0 < z < k \end{cases}$$

Due to the isometry, D_ϵ satisfies $\epsilon d_{\mathbb{R}}$ -privacy iff D'_ϵ satisfies $\epsilon \frac{1}{q} d_{\mathbb{R}}$ -privacy, thus (from the optimality of D_ϵ) it follows that D'_ϵ is $\epsilon \frac{1}{q} d_{\mathbb{R}}$ -optimal for all $\epsilon > 0$.

Finally we define:

$$D''_\epsilon = D'_{q\epsilon}$$

From the optimality of D'_ϵ we get that D''_ϵ is $(q\epsilon) \frac{1}{q} d_{\mathbb{R}}$ -optimal, i.e. it is $\epsilon d_{\mathbb{R}}$ -optimal.

This concludes the proof, since D''_ϵ is exactly the pdf of $TG_\epsilon(\mathcal{Y}_q)$. $\square$

The results above bring us directly to our sufficient condition.

Theorem 10. *Let $\mathcal{Y} = q[0..k]$ and assume that $f : \mathcal{X} \rightarrow \mathcal{Y}$ is uniformly Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathbb{R}}$. Then the TG_ϵ mechanism with input $\mathcal{Y}$ is $f\text{-}\Delta d_{\mathcal{X}}$ -optimal.*

Proof. Direct corollary of Prop 8 and Prop 6. $\square$

C Proofs of Section 5

Proposition 2. $\Phi_1, \mathcal{N}_1$ are maximally tight wrt both $d_1, \tilde{d}_1$.

Proof. We first consider d_1 . Let $x, x' \in \mathcal{V}^n$, we show that there exist a tight chain from x to x' that is both a maximal Φ_1 -chain and a maximal $\mathcal{N}_1$ -chain. We recursively create a chain $x_1, \dots, x_{n+1}$ from x to x' by modifying one element at a time:

$$\begin{aligned} x_1 &= x \\ x_{i+1} &= x_i[x'[i]/i] \quad i \in 1..n \end{aligned}$$

It is easy to see that $d_1(x, x') = \sum_{i=1}^n d_1(x_i, x_{i+1})$ so the chain is tight wrt d_1 .

Fix any $i \in 1..n$ and let $w : \mathcal{V} \rightarrow \mathcal{V}$ defined as:

$$w(v) = \begin{cases} x'[i] & \text{if } v = x[i] \\ x[i] & \text{if } v = x'[i] \\ v & \text{otherwise} \end{cases}$$

For the hiding function $\phi_{i,w} \in \Phi_1$ it holds that

$$\phi_{i,w}(x_i) = x_{i+1} \quad \phi_{i,w}(x_{i+1}) = x_i \quad d_1(x_i, x_{i+1}) = d_1(\phi_{i,w}(x_i), \phi_{i,w}(x_{i+1}))$$

hence the chain is a maximal Φ_1 -chain.

Moreover, let $V = \{x[i], x'[i]\}$. For the neighborhood $N_{i,V}(x_i) \in \mathcal{N}_1$ it holds that

$$\{x_i, x_{i+1}\} \subseteq N_{i,V}(x_i) \quad d_1(x_i, x_{i+1}) = d_1(N_{i,V}(x_i))$$

so the chain is a maximal $\mathcal{N}_1$ -chain.

The case of $\tilde{d}_1$ is similar, since it is a scaled version of d_1 . $\square$

Proposition 3. $\tilde{d}_1 \leq d_h$, thus $\epsilon \tilde{d}_1$ -privacy implies ϵd_h -privacy.

Proof. Fix $x, x' \in \mathcal{V}^n$ and let $I = \{i \in 1..n \mid x[i] \neq x'[i]\}$. Then

$$\tilde{d}_1(x, x') = \frac{\sum_{i \in I} d_{\mathcal{V}}(x[i], x'[i])}{d_{\mathcal{V}}(\mathcal{V})} \leq \frac{\sum_{i \in I} d_{\mathcal{V}}(\mathcal{V})}{d_{\mathcal{V}}(\mathcal{V})} = |I| = d_h(x, x')$$

□

We continue by introducing a relaxed version of the concept of maximal $\mathcal{N}$ -tightness (Def 3), called $\mathcal{N}$ -tightness, by simply dropping the requirement $d_{\mathcal{X}}(x_i, x_{i+1}) = d_{\mathcal{X}}(N)$ from Def 3.

Definition 10. Let $\mathcal{N} \subseteq 2^{\mathcal{X}}$. A chain $\tilde{x}$ is called an $\mathcal{N}$ -chain iff for every step i there exists $N \in \mathcal{N}$ such that $\{x_i, x_{i+1}\} \subseteq N$. Then $\mathcal{N}$ is called tight wrt $d_{\mathcal{X}}$ iff $\forall x, x' \in \mathcal{X}$ there exists a tight $\mathcal{N}$ -chain from x to x' .

We can now show the (maximal wrt d_h , simple wrt $d_1, \tilde{d}_1$) tightness of $\mathcal{N}_h$, which will be useful later on.

Proposition 9. $\mathcal{N}_h$ is maximally tight wrt d_h and tight wrt both $d_1, \tilde{d}_1$.

Proof. Let $x, x' \in \mathcal{V}^n$. We need to show that there exists a tight (also maximal in the case of d_h) $\mathcal{N}_h$ -chain from x to x' . We recursively create a chain $x_1, \dots, x_{n+1}$ from x to x' by modifying one element at a time:

$$\begin{aligned} x_1 &= x \\ x_{i+1} &= x_i[x'[i]/i] \quad i \in 1..n \end{aligned}$$

It is easy to see that $d(x, x') = \sum_{i=1}^n d(x_i, x_{i+1})$, for all $d \in \{d_h, d_1, \tilde{d}_1\}$, so the chain is tight wrt all three metrics. Moreover, we have that $\{x_i, x_{i+1}\} \subseteq N_i(x_i)$ so the chain is an $\mathcal{N}_h$ -chain wrt all metrics.

For d_h , it also holds that $d_h(x_i, x_{i+1}) = d_h(N_i(x_i)) = 1$, so the chain is a maximal $\mathcal{N}_h$ -chain wrt d_h . □

We continue with a lemma that facilitates proofs of Δ -sensitivity by reducing the pairs of secrets x, x' that one needs to check.

Lemma 3. Let $\mathcal{N} \subseteq 2^{\mathcal{X}}$ be tight (Def 10) and assume:

$$d_{\mathcal{Y}}(f(x), f(x')) \leq \Delta d_{\mathcal{X}}(x, x') \quad \forall N \in \mathcal{N}, x, x' \in N$$

Then f is Δ -sensitive wrt $d_{\mathcal{X}}, d_{\mathcal{Y}}$.

Proof. Fix $x, x' \in \mathcal{X}$, we need to show that $d_{\mathcal{Y}}(f(x), f(x')) \leq \Delta d_{\mathcal{X}}(x, x')$. Since $\mathcal{N}$ is tight there exist a tight $\mathcal{N}$ -chain $x = x_1, \dots, x_n = x'$, such that each step x_i, x_{i+1}

belongs to some set $N \in \mathcal{N}$. We have:

$$\begin{aligned}
& d_{\mathcal{V}}(f(x), f(x')) \\
& \leq \sum_{i=1}^{n-1} d_{\mathcal{V}}(f(x_i), f(x_{i+1})) && \text{triangle ineq.} \\
& \leq \Delta \sum_{i=1}^{n-1} d_{\mathcal{X}}(x_i, x_{i+1}) && \text{hypoth., } x_i, x_{i+1} \in N \\
& = \Delta d_{\mathcal{X}}(x, x') && \text{tightness of chain}
\end{aligned}$$

□

Proposition 4. *Let $f \in \mathcal{C}$. The sensitivity of f wrt both $d_h, d_{\mathbb{R}}$ and $\tilde{d}_1, d_{\mathbb{R}}$ is $d_{\mathcal{V}}(\mathcal{V})$.*

Proof. Let $f \in \mathcal{C}$. We first show that f is $d_{\mathcal{V}}(\mathcal{V})$ -sensitive wrt both $d_h, d_{\mathbb{R}}$ and $\tilde{d}_1, d_{\mathbb{R}}$. From Prop 9 together with Theorem 3, we only need to show the sensitivity for databases x, x' from some set $N_i(x) \in \mathcal{N}_h$, i.e. $x \sim_i x'$.

For d_h , we have $d_h(x, x') = 1$, thus

$$d_{\mathcal{V}}(f(x), f(x')) \leq d_{\mathcal{V}}(x[i], x'[i]) \leq d_{\mathcal{V}}(\mathcal{V})d_h(x, x')$$

For $\tilde{d}_1$ we have:

$$d_{\mathcal{V}}(f(x), f(x')) \leq d_{\mathcal{V}}(x[i], x'[i]) = d_{\mathcal{V}}(\mathcal{V})\tilde{d}_1(x, x')$$

Then, for any $\Delta < d_{\mathcal{V}}(\mathcal{V})$, f is not Δ -sensitive for neither metric, since from Def 8 there exists $x \sim_i x'$ such that

$$d_{\mathcal{V}}(f(x), f(x')) = d_{\mathcal{V}}(\mathcal{V}) > \Delta d_h(x, x')$$

and similarly for $\tilde{d}_1$. □

Proposition 5. *The sum, avg, p-perc queries belong to $\mathcal{C}$.*

Proof. The universe is assumed to be $\mathcal{V} = q[0..k]_{\emptyset}$ for some $k \in \mathbb{N}, q > 0$. Let $x \sim_i x' \in \mathcal{V}^n$. We first show that $d_{\mathbb{R}}(f(x), f(x')) \leq d_{\mathcal{V}}(x[i], x'[i])$.

For sum, it is easy to see that

$$|\text{sum}(x), \text{sum}(x')| = \begin{cases} d_{\mathcal{V}}(x[i], x'[i]) & x[i] \neq \emptyset, x'[i] \neq \emptyset \\ x'[i] & x[i] = \emptyset \\ x[i] & x'[i] = \emptyset \end{cases}$$

Note that $x'[i] \leq d_{\mathcal{V}}(x[i], x'[i]) = qr$ in the case $x[i] = \emptyset$ (and similarly for $x'[i] = \emptyset$).¹⁰

¹⁰ It is crucial here that $\mathcal{V}$ contains 0, so that $v \leq d_{\mathcal{V}}(\mathcal{V})$ for all non-null v . If $0 \notin \mathcal{V}$, we can achieve a similar result for sum by adapting the way $d_{\mathcal{V}}$ treats $\emptyset$.

Consider now $f \in \{\text{avg}, p\text{-perc}\}$. From Theorem 13 we know that both queries are 1-sensitive wrt $d_\infty, d_{\mathbb{R}}$. And since $d_\infty(x, x') = d_{\mathcal{V}}(x[i], x'[i])$ we have:

$$d_{\mathbb{R}}(f(x), f(x')) \leq d_\infty(x, x') = d_{\mathcal{V}}(x[i], x'[i])$$

Finally, we need to show that there exist $x \sim_i x' \in \mathcal{V}^n$ such that $d_{\mathbb{R}}(f(x), f(x')) = d_{\mathcal{V}}(\mathcal{V}) = qk$. We construct $x = \langle 0, \emptyset, \dots, \emptyset \rangle$, $x' = \langle qk, \emptyset, \dots, \emptyset \rangle$. These databases satisfy $d_{\mathbb{R}}(f(x), f(x')) = qk$ for all queries. $\square$

Theorem 11. *The sum, avg and p-perc queries are all uniformly qk -sensitive wrt $\tilde{d}_1, d_{\mathbb{R}}$.*

Proof. First we have to show that the queries are qk -sensitive wrt $\tilde{d}_1, d_{\mathbb{R}}$. This comes from Prop 4 and 5, since all queries belong to the family $\mathcal{C}$.

We now show the uniform sensitivity of sum. Let $y, y' \in q[0..nk]$ and assume that $y \geq y'$. It is easy to see that we can construct databases x, x' such that $\text{sum}(x) = y$, $\text{sum}(x') = y'$ and $x[i] \geq x'[i]$ for all $i \in 1..n$. For x, x' we have

$$\begin{aligned} d_1(x, x') &= \sum_i |x[i] - x'[i]| \\ &= |\sum_i x[i]| - |\sum_i x'[i]| && x[i] \geq x'[i] \\ &= d_{\mathbb{R}}(\text{sum}(x), \text{sum}(x')) \end{aligned}$$

Thus $d_{\mathbb{R}}(y, y') = qk \tilde{d}_1(x, x')$, which means that the chain y, y' is qk -expansive wrt $\tilde{d}_1$.

Finally, for $f \in \{\text{avg}, p\text{-perc}\}$ let $y, y' \in q[0..k]$. We construct two databases x, x' with a single present individual as follows:

$$x = \langle y, \emptyset, \dots, \emptyset \rangle \quad x' = \langle y', \emptyset, \dots, \emptyset \rangle$$

It is easy to see that $f(x) = y$, $f(x') = y'$ and $d_1(x, x') = d_{\mathbb{R}}(y, y')$. Thus $d_{\mathbb{R}}(y, y') = qk \tilde{d}_1(x, x')$ which means that the chain y, y' is qk -expansive wrt $\tilde{d}_1$. $\square$

Theorem 12. *The sum, avg and p-perc queries are all uniformly 1-sensitive wrt $d_1, d_{\mathbb{R}}$.*

Proof. Direct consequence of Theorem 11, since $d_1 = qk \tilde{d}_1$. $\square$

D Proofs of Section 6

Proposition 10. $\Phi_\infty, \mathcal{N}_\infty$ are maximally tight wrt d_∞ .

Proof. Let $x, x' \in \mathcal{V}^n$, and consider the trivial 1-step chain x, x' . This chain is trivially tight, we need to show that it is both a maximal Φ_∞ -chain and a maximal $\mathcal{N}_\infty$ -chain.

For each $i \in 1..n$ we define a function $w_i : \mathcal{V} \rightarrow \mathcal{V}$ as:

$$w_i(v) = \begin{cases} x'[i] & \text{if } v = x[i] \\ x[i] & \text{if } v = x'[i] \\ v & \text{otherwise} \end{cases}$$

For the hiding function $\phi = \phi_{1,w_1} \circ \dots \circ \phi_{n,w_n}$ we have that $\phi \in \Phi_\infty$ and moreover:

$$\phi(x) = x' \quad \phi(x') = x \quad d_\infty(x, x') = d_\infty(\phi)$$

hence the chain is a maximal Φ_∞ -chain.

Moreover, let $V_i = \{x[i], x'[i]\}, i \in 1..n$. For the neighborhood $N_{\{V_i\}} \in \mathcal{N}_\infty$ it holds that

$$\{x, x'\} \subseteq N_{\{V_i\}} \quad d_\infty(x, x') = d_\infty(N_{\{V_i\}})$$

so the chain is a maximal $\mathcal{N}_\infty$ -chain. $\square$

Theorem 13. *The avg and p-perc queries are both uniformly 1-sensitive wrt $d_\infty, d_\mathbb{R}$.*

Proof. The universe is assumed to be $\mathcal{V} = q[0..k]_\emptyset$ for some $k \in \mathbb{N}, q > 0$. The p -percentile query ($0 \leq p < 100$) is defined as $p\text{-perc}(x) = \text{sort}(x)[l]$ for $l = \lfloor \frac{p}{100}m + 1 \rfloor$, where m is the number of non-null values in x and sort returns a sorted version of x (after removing the null values). We also define $I_\emptyset(x) = \{i \in 1..n \mid x[i] = \emptyset\}$.

We first show that both queries are 1-sensitive wrt $d_\infty, d_\mathbb{R}$. Let $x, x' \in \mathcal{V}^n$. If $I_\emptyset(x) \neq I_\emptyset(x')$ then x, x' are maximally distant, i.e. $d_\infty(x, x') = d_\infty(\mathcal{V}^n) = qk$. Then for both queries it trivially holds that $d_\mathbb{R}(f(x), f(x')) \leq qk = d_\infty(x, x')$ since their range is $q[0..k]$.

It remains to show 1-sensitivity for the case $I_\emptyset(x) = I_\emptyset(x') = I$. For the average query we have

$$\begin{aligned} & d_\mathbb{R}(\text{avg}(x), \text{avg}(x')) \\ &= \frac{1}{|I|} |(\sum_{i \in I} x[i]) - (\sum_{i \in I} x'[i])| \\ &= \frac{1}{|I|} |\sum_{i \in I} (x[i] - x'[i])| \\ &\leq \frac{1}{|I|} \sum_{i \in I} |x[i] - x'[i]| && \text{subadditivity of } |\cdot| \\ &\leq \frac{1}{|I|} \sum_{i \in I} d_\infty(x, x') && |x[i] - x'[i]| \leq d_\infty(x, x') \\ &= d_\infty(x, x') \end{aligned}$$

For the p -perc query it holds that $p\text{-perc}(x) = \text{sort}(x)[l]$ and $p\text{-perc}(x') = \text{sort}(x')[l]$ for the same l (since $I_\emptyset(x) = I_\emptyset(x')$). Let $h, h' \in 1..n$ such that $x[h] = \text{sort}(x)[l]$ and $x'[h'] = \text{sort}(x')[l]$.

Assume that $x[h] \leq x[h']$ (the case $x[h] \geq x[h']$ is symmetric). By the definition of sort , there are at least l elements $j \in 1..n$ such that $x[j] \leq x[h]$ (including h itself). Moreover, there are at most $l - 1$ elements $j \in 1..n$ such that $x'[j] < x'[h']$. Hence, there exists at least one $j \in 1..n$ such that

$$x[j] \leq x[h] \quad \text{and} \quad x'[j] \geq x'[h']$$

It also holds that $|x[i] - x'[i]| \leq d_\infty(x, x')$, i.e.

$$x[i] - d_\infty(x, x') \leq x'[i] \leq x[i] + d_\infty(x, x') \quad \forall i \in 1..n \quad (4)$$

From $x[h] \leq x[h']$ and (4) we get

$$x[h] - d_\infty(x, x') \leq x'[h']$$

Moreover, it holds that

$$x'[h'] \leq x'[j] \leq x[j] + d_\infty(x, x') \leq x[h] + d_\infty(x, x')$$

thus

$$d_\mathbb{R}(p\text{-perc}(x), p\text{-perc}(x')) = |x[h] - x'[h']| \leq d_\infty(x, x')$$

For the “uniformly” part, let $y, y' \in q[0..k]$; we construct $x = \langle y, \emptyset, \dots, \emptyset \rangle$, $x' = \langle y', \emptyset, \dots, \emptyset \rangle$, for which it holds that $f(x) = y, f(x') = y'$ (for both queries) and $d_\infty(x, x') = d_\mathbb{R}(y, y')$.

Note that for $p\text{-perc}$ we can construct x, x' without $\emptyset$ values with the same property. However, for the avg query this is not possible; its uniform optimality depends on the fact that $\emptyset$ values are allowed. If $\emptyset \notin \mathcal{V}$ then avg is essentially equivalent to sum , which is not uniformly optimal wrt $d_\infty, d_\mathbb{R}$. $\square$