

Subspace Differential Privacy

Jie Gao¹, Ruobin Gong¹, Fang-Yi Yu²

¹ Rutgers University

² Harvard University

jg1555@rutgers.edu, rg915@stat.rutgers.edu, fangyiyu@seas.harvard.edu

Abstract

Many data applications have certain invariant constraints due to practical needs. Data curators who employ differential privacy need to respect such constraints on the sanitized data product as a primary utility requirement. Invariants challenge the formulation, implementation and interpretation of privacy guarantees. We propose *subspace differential privacy*, to honestly characterize the dependence of the sanitized output on confidential aspects of the data. We discuss two design frameworks that convert well-known differentially private mechanisms, such as the Gaussian and the Laplace mechanisms, to subspace differentially private ones that respect the invariants specified by the curator. For linear queries, we discuss the design of near optimal mechanisms that minimize the mean squared error. Subspace differentially private mechanisms rid the need for post-processing due to invariants, preserve transparency and statistical intelligibility of the output, and can be suitable for distributed implementation. We showcase the proposed mechanisms on the 2020 Census Disclosure Avoidance demonstration data, and a spatio-temporal dataset of mobile access point connections on a large university campus.

1 Introduction

Invariants: a challenge for data privacy Data publication that satisfies differential privacy carries the formal mathematical guarantee that an adversary cannot effectively tell the difference, in the probabilistic sense, when two databases differ in only one entry. The extent of privacy protection under differential privacy is quantified by the privacy loss budget parameters, such as in ϵ -differential privacy and (ϵ, δ) -differential privacy (Dwork et al. 2006b). While the differential privacy guarantee is rigorously formulated with the probability language, its construction does not naturally mingle with hard and truthful constraints, called *invariants* (Ashmead et al. 2019), that need to be imposed onto the sanitized data product, often as a primary utility requirement.

An important use case in which the challenge for privacy arises from invariants is the new Disclosure Avoidance System (DAS) of the 2020 Decennial Census (Abowd 2018). The new DAS tabulates noise-infused, differentially private counts into multi-way contingency tables at various geographic resolutions, from the aggregated state and county

levels to specific Census blocks. Due to the Census Bureau’s constitutional mandate and its responsibilities as the official statistics agency of the United States, all data products must be preserved in such a way that certain aspects of their values are *exactly* as enumerated. These invariants include (and are not limited to) population totals at the state level, counts of total housing units, as well as other group quarter facilities at the block level. Straightforward tabulations of the noisy measurements are most likely inconsistent with the mandated invariants.

A common method to impose invariants on a differentially private noisy query is via post-processing using distance minimization (Abowd et al. 2019; Ashmead et al. 2019). The resulting query is the solution to an optimization task, one that minimizes a pre-specified distance between the unconstrained query and the invariant-compliant space. There are two major drawbacks to this approach. First, post-processing may introduce systematic bias into the query output. Particularly troubling is that the source of such bias is poorly understood (Zhu, Hentenryck, and Fioretto 2020), in part due to the highly data-dependent nature of the post-processing procedure, and the lack of a transparent probabilistic description. The TopDown algorithm (Abowd et al. 2019), employed by the Census DAS to impose invariants on the noisy measurements, exhibits a notable bias that it tends to associate larger counts with positive errors, whereas smaller counts with negative errors, when the total count is held as invariant. Figure 1 illustrates this phenomenon using the November 2020 vintage Census demonstration data (Van Riper, Kugler, and Schroeder 2020). For all states and state-level territories with more than five counties (i.e. excluding D.C., Delaware, Hawaii and Rhode Island), a simple regression is performed between the county-level DAS errors and the log true county population sizes. Of the 48 regressions, 37 result in a negative slope estimate, out of which 11 are statistically significant at $\alpha = 0.01$ level (red circles in the left panel), indicating a systematic negative association between the DAS errors and the true counts. The bias trend is clearly visible in the right panel among the DAS errors (red squares) associated with the counties of Illinois, ordered by increasing population size. As Zhu, Hentenryck, and Fioretto (2020) discussed, the bias exhibited in the demonstration data is attributed to the non-negativity constraints imposed on the privatized counts.

The second, and more fundamental, drawback is that imposing invariants that are true aspects of the confidential data may incur additional privacy leakage, even if the invariants are also publicly released (Gong and Meng 2020). When conveyed to data users and contributors, the narrative that the invariants are imposed by “post-processing” may lend to the erroneous interpretation that no additional leakage would occur. Care needs to be taken to explain that whenever non-trivial invariants are observed, the usual (ϵ, δ) -differential privacy guarantee cannot be understood in its fullest sense.

The need to impose invariants arises in application areas involving the monitoring of stochastic events in structured spatio-temporal databases. In some cases, there are requirements to report accurate counts (service requests (City of New York 2021) or traffic volumes (Sui et al. 2016; Yang et al. 2019, 2020; Wang and Gao 2020)). In other cases, there are invariants that can be derived from external common sense knowledge –e.g., the number of vehicles entering and leaving a tunnel should be the same (when there are no other exits or parking spaces). An adversary may potentially use such information to reverse engineer the privacy protection perturbation mechanisms (Rezaei and Gao 2019; Rezaei, Gao, and Sarwate 2021). Invariants pose a new challenge to both the curators and the users of private data products, prompting its recognition as a new source of compromise to privacy that stems from external mandates.

Our contribution To meet the challenge posed by invariants, we argue that the definition of differential privacy must be recapitulated to respect the given constraints. To this end, we propose the definition of *subspace differential privacy*, which makes explicit the invariants imposed on the data product. Subspace differential privacy intends to honestly characterize the dependence of the privatized data product on truthful aspects of the confidential data, to ensure that any privacy guarantee attached to the data product is both mathematically rigorous and intuitively sensible. It enables the assessment of what kind of queries do, and do not, enjoy the protection of differential privacy.

The literature has seen attempts to generalize beyond the classic notion of differential privacy. The framework of Pufferfish privacy (Kifer and Machanavajjhala 2012, 2014; Song, Wang, and Chaudhuri 2017) specifies the potential secrets, discriminative pairs, as well as the data generation model and knowledge possessed by the potential attacker. Special cases of the Pufferfish framework include Blowfish privacy (He, Machanavajjhala, and Ding 2014) and Bayesian differential privacy (Yang, Sato, and Nakagawa 2015). Related notions of correlated differential privacy (Zhu et al. 2014) and dependent differential privacy (Liu, Chakraborty, and Mittal 2016) specifically address secrets in the form of query correlations or structural dependence of the database.

The current work differentiates itself from the existing literature in two senses. First, the theoretical focus is to provide a principled reconciliation between the hard truth constraints which the data curator must impose on the sanitized data product, and the privacy guarantee the product can enjoy. In particular, just like the classic notion of differential privacy, *subspace differential privacy* does not require the specification of a data generation model nor any knowledge

that the attacker might possess. Second, the practical emphasis is on the design of probabilistic mechanisms that impose deterministic truth constraints as they instill subspace differential privacy in the data product. This forgoes the need for additional post-processing, and preserves good statistical qualities of the output.

A related, but different, line of work in the literature concerns the *internal* consistency of the privacy mechanism output (Barak et al. 2007; Hay et al. 2009). For example, when we query the number of students in a school and the numbers of students in each class of the school, we may expect the outputs to be non-negative, and the sum of the (privatized) numbers of students in all classes to be equal to the (privatized) number of students in the school. These internal consistency requirements, such as non-negativity and relational constraints, are *independent* of the private dataset. Therefore, they may be compatible with the classic notion of differential privacy, in which case they may be instantiated with differentially private mechanisms. However, for invariants that are nontrivial functions of the confidential data, we show in Section 2.1 that it is impossible to have differentially private mechanisms that satisfy them. It is this kind of invariants that motivate our work in this paper.

The remainder of this paper is organized as follows. Section 2 defines *subspace differential privacy* and *induced subspace differential privacy*, motivated by the pair of necessary criteria that the mechanism be simultaneously provably private and invariant-respecting. Section 3 outlines two general approaches, *projection* and *extension*, to design induced subspace differentially private mechanisms. We apply both frameworks to produce Gaussian and Laplace mechanisms for general queries, present a correlated Gaussian mechanism that is *near-optimal* (i.e. in terms of mean squared error, with a small multiplicative factor) for linear queries, and sketch the design for a k -norm mechanism that would enjoy near optimality. Section 4 discusses the statistical and implementation considerations behind the proposed mechanisms, as they enjoy transparency and statistical intelligibility that facilitate principled downstream statistical analysis. In the special case of additive spherical Gaussian mechanism, a distributional equivalence is established between the projection framework and statistical conditioning. All mechanisms can also be adapted for distributed privatization. Section 5 provides two demonstrations of the proposed induced subspace differentially private mechanisms, on the 2020 Census DAS demonstration data and spatio-temporal mobility dataset on a university campus subject to various marginal total invariants. Section 6 concludes.

2 Recapitulating Privacy under Invariants

In this work, we model private data as a database $\mathbf{x} = (x_1, \dots, x_N)^\top \in \mathcal{X}^N$ of N rows, where each row $x_i \in \mathcal{X}$ contains data about an individual i , and \mathcal{X} is finite with size d , and we set the space of all possible non-empty databases as $\mathcal{X}^* := \cup_{N \geq 1} \mathcal{X}^N$. A trusted curator holds the database $\mathbf{x} \in \mathcal{X}^*$, and provides an interface to the database through a randomized mechanism $M : \mathcal{X}^* \rightarrow \mathcal{Y}$ where $\mathcal{Y} \subseteq \mathbb{R}^n$ is the *output space* of M . We want to design good mechanisms to

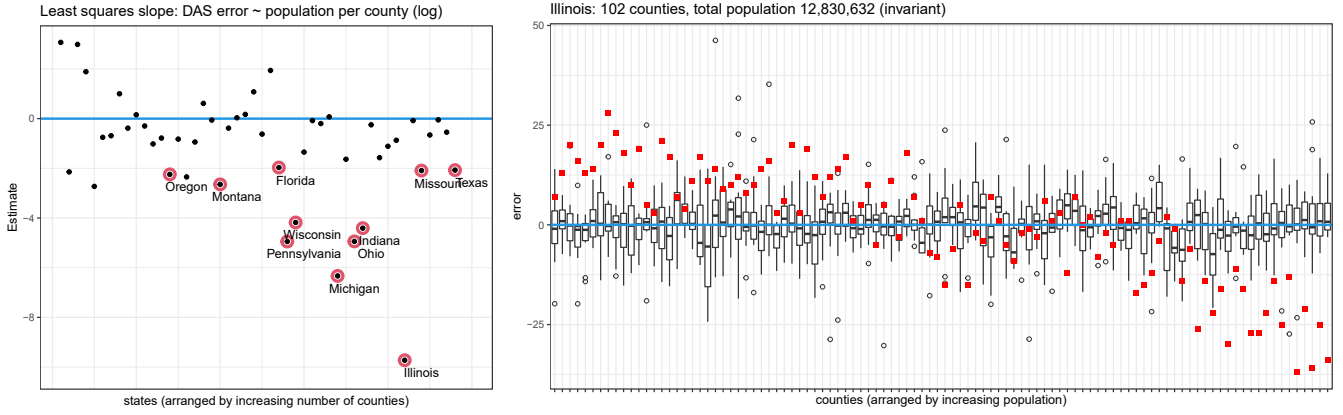


Figure 1: Left: the Census DAS associates positive errors with larger counties and negative errors with smaller counties, when the state population is held as invariant (Nov 2020 vintage demonstration files; Van Riper, Kugler, and Schroeder 2020). Eleven out of 48 simple regressions of the county-level DAS errors against log true county populations have statistically significant negative slopes ($\alpha = 0.01$), circled in red. Right: for the counties of Illinois in increasing true population sizes, DAS errors (red squares) show a clear negative trend bias. The boxplots show errors from ten runs of our proposed method (the $(\epsilon, 0)$ -induced subspace differentially private projected Laplace mechanism; Corollary 3.8). As Corollary 4.1 shows, these errors are unbiased.

answer a query $A : \mathcal{X}^* \rightarrow \mathcal{Y}$ that satisfies not only certain privacy notions, but also invariant constraints as motivated in Section 1.

We begin the discussion about mechanisms for a general query, defined by a function $A : \mathcal{X}^* \rightarrow \mathcal{Y}$, throughout the end of Section 3.3. In Section 3.4, we consider optimal mechanisms for a *linear* query A , with $a : \mathcal{X} \rightarrow \mathbb{R}^n$ so that $A(\mathbf{x}) := \sum_i a(x_i)$. Indeed, a linear query A can be represented as a linear function of the *histogram* of database \mathbf{x} , $\text{hist}(\mathbf{x}) : \mathcal{X}^* \rightarrow \mathbb{N}^d$ where $\text{hist}(\mathbf{x})_z := \sum_i \mathbf{1}[x_i = z]$ is the number of rows equal to $z \in \mathcal{X}$ in $\mathbf{x} \in \mathcal{X}^*$. With this notation, given a linear query A , we denote \mathcal{A} as a matrix where the k, z entry is $A(z)_k$ for $k \in [n]$ and $z \in \mathcal{X}$, and the linear query on a database \mathbf{x} can be written as matrix multiplication, $A(\mathbf{x}) = \mathcal{A} \cdot \text{hist}(\mathbf{x})$.

2.1 Privacy Guarantees and Invariants

The notion of differential privacy ensures that no individual’s data has much effect on the output of the mechanism M . That is, if we consider any two neighboring databases \mathbf{x} and \mathbf{x}' of size N that differ on one row (there exists i such that $x_i \neq x'_i$ and $x_j = x'_j$ for all $j \neq i$.) the output distribution of mechanism M on \mathbf{x} should be similar to that of M on \mathbf{x}' . Formally:

Definition 2.1 (Bounded differential privacy (Dwork et al. 2006b)). Let $(\mathcal{Y}, \mathcal{F})$ be a measurable space and $\epsilon, \delta \geq 0$. We say that a random mechanism $M : \mathcal{X}^* \rightarrow \mathcal{Y}$ is (ϵ, δ) -*differentially private* if for all neighboring databases $\mathbf{x} \sim \mathbf{x}'$ and all measurable set $S \in \mathcal{F}$, $\Pr[M(\mathbf{x}) \in S] \leq e^\epsilon \cdot \Pr[M(\mathbf{x}') \in S] + \delta$.

From the data curator’s perspective, in addition to privacy concerns, there often exists external constraints that the privatized output M must meet. These constraints can often be represented as a function of $M(\mathbf{x})$ that agrees with what’s calculated based on the confidential $A(\mathbf{x})$. In this work, we

focus on the class of invariants that can be written in the form of a system of linear equations.

Definition 2.2 (Invariants - linear equality). Given a query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, and $C \in \mathbb{R}^{n_c \times n}$ be a $n_c \times n$ matrix with rank $n_c < n$.¹ A (random) mechanism $M : \mathcal{X}^* \rightarrow \mathcal{Y} \subseteq \mathbb{R}^n$ satisfies the *linear equality invariant* C with query A , if for all \mathbf{x} , $CM(\mathbf{x}) = CA(\mathbf{x})$ with probability one over the randomness of M .

Given a linear equality invariant C , let $\mathcal{N} := \{v \in \mathbb{R}^n : Cv = 0\}$ be the *null space* of C , and $\mathcal{R} := \mathcal{N}^\perp \subseteq \mathbb{R}^n$ be the *row space* of C . Additionally, we set $\Pi_{\mathcal{N}}$ be the orthogonal projection matrix for null space \mathcal{N} , $Q_{\mathcal{N}}$ be a collection of orthonormal basis of \mathcal{N} , and $A_{\mathcal{N}} := Q_{\mathcal{N}}^\top A$ be the query function A projected into \mathcal{N} with basis $Q_{\mathcal{N}}$. We use subscript \mathcal{R} in the similar manner.

Linear equality invariants are a natural family of invariants. Here are two examples.

Example 2.3. Let $\mathbf{x} \in \mathcal{X}^N$ be a database with $|\mathcal{X}| = d$, $A = \text{hist}$ be the histogram query, and $C = \mathbf{1}^\top = (1, \dots, 1) \in \mathbb{R}^{1 \times d}$ all one vector with length d . This linear equality invariant requires the curator to accurately report the total number of individuals without error, because $CM(\mathbf{x}) = \mathbf{1}^\top \text{hist}(\mathbf{x}) = N$.

Example 2.4. Consider $\mathcal{X} = \{1, 2, 3, 4\}$, $A = \text{hist} : \mathcal{X}^* \rightarrow \mathbb{R}^4$ and $C = (1, 0, 1, 0)$. The linear equality invariant ensures the number of individual with odd feature is exact.

The invariants discussed in the Census DAS application, such as state-level population and block-level housing units and group quarter facilities, can be formulated as linear equality invariants.

¹We can always find a C' consists of a subset of independent rows of C which has same row rank as C ’s, and we can translate between these two through a linear transformation.

2.2 Subspace and Induced Subspace Differential Privacy

The mechanism we seek must meet two necessary criteria: *provably private* and *invariant-respecting*. That is, the mechanism should privatize the confidential query with mathematically provable guarantees, while at the same time the query output should conform to the invariants that a data curator chooses to impose. The two criteria culminate in the *induced subspace differential privacy* (Definition 2.6) which enjoys additional desirable properties, such as the practicalities of design and statistical intelligibility, which will be discussed from section 3 and on.

To motivate the construction, note that the classic definition of differential privacy and invariant constraints are not compatible by design. For instance, in Example 2.4 if a mechanism M respects the invariant constraint Definition 2.2, the probability ratio of event $S = \{(y_1, y_2, y_3, y_4) : y_1 + y_3 = 1\} \subset \mathbb{R}^4$ on neighboring databases $\mathbf{x} = (1, 2, 4)$ and $\mathbf{x}' = (4, 2, 4)$ is unbounded, $\frac{\Pr[M(\mathbf{x}) \in S]}{\Pr[M(\mathbf{x}') \in S]} = \infty$, because $\Pr[M(\mathbf{x}) \in S] = \Pr[CM(\mathbf{x}) = C \text{ hist}(\mathbf{x})] = 1$ but $\Pr[M(\mathbf{x}') \in S] = \Pr[CM(\mathbf{x}') \neq C \text{ hist}(\mathbf{x}')] = 0$. Thus M violates (ϵ, δ) -differential privacy for any $\epsilon > 0$ and $\delta < 1$.

Therefore, we need a new notion of differential privacy to discuss the privacy protection in the presence of mandated invariants. Below we attempt to do so by recapitulating the definition of (ϵ, δ) -differential privacy, to acknowledge the fact that if a hard linear constraint is imposed on the privacy mechanism, we can no longer offer differential privacy guarantee in the full n -dimensional space that is the image of A , but rather only within certain linear subspaces.

Definition 2.5 (Subspace differential privacy). Let \mathcal{V} be a linear subspace of \mathbb{R}^n , and $\Pi_{\mathcal{V}}$ the projection matrix onto \mathcal{V} . Given $\epsilon, \delta \geq 0$, a random mechanism $M : \mathcal{X}^* \rightarrow \mathbb{R}^n$ is \mathcal{V} -subspace (ϵ, δ) -differentially private if for all neighboring databases $\mathbf{x} \sim \mathbf{x}'$ and every Borel subset $S \subseteq \mathcal{V}$,

$$\Pr[\Pi_{\mathcal{V}}M(\mathbf{x}) \in S] \leq e^\epsilon \Pr[\Pi_{\mathcal{V}}M(\mathbf{x}') \in S] + \delta. \quad (1)$$

We are ready to formalize the notion of a provably private and invariant-respecting private mechanism, one that meets both the criteria laid out at the beginning of this subsection.

Definition 2.6 (Induced subspace differential privacy). Given $\epsilon, \delta \geq 0$, a query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$ and a linear equality invariant $C : \mathbb{R}^n \rightarrow \mathbb{R}^{n_c}$ with null space \mathcal{N} , a mechanism $M : \mathcal{X}^* \rightarrow \mathbb{R}^n$ is (ϵ, δ) -induced subspace differentially private for query A and an invariant C if 1) M is \mathcal{N} -subspace (ϵ, δ) -differentially private (Definition 2.5), and 2) M satisfies the linear equality invariant C (Definition 2.2). M may be referred to simply as *induced subspace differentially private*, whenever the context is clear about (or does not require the specification of) ϵ, δ and C .

An induced subspace differentially private mechanism delivers query outputs that meet the curator's invariant specification (C) with probability one. It is provably differentially private for all queries and their components that are orthogonal to the invariants, and is silent on privacy properties for those that are linearly dependent on the invariants.

2.3 Properties of Subspace Differential Privacy

Now we discuss several properties of subspace differential privacy. First, we show a “nestedness” property: a \mathcal{V}_1 -subspace differentially private mechanism is also \mathcal{V}_2 -subspace differentially private for all $\mathcal{V}_1 \supseteq \mathcal{V}_2$.

Proposition 2.7. *Let $\mathcal{V}_2 \subseteq \mathcal{V}_1$ be nested linear subspaces of respective dimensions $d_2 \leq d_1$. If a mechanism M is \mathcal{V}_1 -subspace (ϵ, δ) -differentially private, it is also \mathcal{V}_2 -subspace (ϵ, δ) -differentially private.*

The main idea is that because $\mathcal{V}_2 \subseteq \mathcal{V}_1$ are both linear spaces, for any measurable S , we can always find another measurable set S' such that $\Pi_{\mathcal{V}_2}^{-1}(S) = \Pi_{\mathcal{V}_1}^{-1}(S')$. Thus \mathcal{V}_1 -subspace differential privacy implies \mathcal{V}_2 -subspace differential privacy. We include the proof to the appendix.

Proposition 2.7 implies that a differentially private mechanism is subspace differential private, as shown in Corollary 2.8. Thus, we can call an (ϵ, δ) -differentially private mechanism the \mathbb{R}^n -subspace (ϵ, δ) -differentially private mechanism.

Corollary 2.8. *If $M : \mathcal{X}^* \rightarrow \mathbb{R}^n$ is a (ϵ, δ) -differentially private mechanism, it is \mathcal{V} -subspace (ϵ, δ) -differentially private for any linear subspace $\mathcal{V} \subseteq \mathbb{R}^n$.*

Induced subspace differential privacy inherits the composition property from differential privacy in the following sense (with proof deferred to the full version).

Proposition 2.9 (Composition). *Given $\epsilon_1 \geq 0$ and $\epsilon_2 \geq 0$, if M_1 is induced subspace $(\epsilon_1, 0)$ -differentially private for query A_1 and linear invariant C_1 and M_2 is induced subspace $(\epsilon_2, 0)$ -differentially private for query A_2 and linear invariant C_2 , the composed mechanism (M_1, M_2) is induced subspace $(\epsilon_1 + \epsilon_2, 0)$ -differentially private for query $A_{1,2} := (A_1, A_2)$ and linear invariant $C_{1,2} := (C_1, C_2)$.*

The above is a preliminary answer to the composition property of subspace differential privacy. However, when C_1 and C_2 are different, the composition of invariant constraints may reveal additional information about the underlying confidential data. In general, the composition of logically independent and informative invariants is not unlike a database linkage attack. For instance, $A_1 = A_2$ is a two-way contingent table that reports the counts of individuals with ages and zip code, the invariant C_1 ensures the accurate marginal counts of individuals within each age bracket, and C_2 ensures the accurate count of individuals within each zip code. The composition of these two invariant constraints may allow an adversary to infer each individual's information.

The subspace differential privacy is naturally immune to any post-processing mapping which only acts on the subspace \mathcal{N} . However, it is not readily clear how to define post-processing under mandated invariant constraints. Specifically, if a mechanism satisfies a nontrivial invariant constraint C with row space \mathcal{R} and null space \mathcal{N} , we can apply a post-processing mapping that outputs the invariant component in \mathcal{R} – here the output is revealed precisely. Unfortunately, such will be true for any invariant-respecting privatized product, regardless of what notion of privacy is attached to it.

Finally, we note that if we consider a mixture of two \mathcal{V} -subspace differentially privacy mechanisms the resulting mechanism is also \mathcal{V} -subspace differentially private. This property is known as the privacy axiom of choice (Kifer and Lin 2010).

3 Mechanism Design

This section introduces two frameworks for designing induced subspace differentially private mechanisms with linear equality invariant C . The data curator would invoke the *projection* framework if seeking to impose invariants onto an *existing* differentially private mechanism, and the *extension* framework if augmenting a smaller private query in a manner compatible with the invariants. Both frameworks are applied to revise existing differentially private mechanisms in Section 3.3, notably the Gaussian and the Laplace mechanisms, for general queries. For linear queries, Section 3.4 presents a near optimal correlated Gaussian mechanism, and sketches the design of a near optimal k -norm mechanism.

3.1 The Projection Framework

Suppose the data curator already employs a differentially private mechanism M to answer a general query A , and would like to impose a linear equality invariant C on the query output. The projection framework, outlined in Theorem 3.1, can transform the existing mechanism M to induced subspace differentially private for A and C , with little overhead on the curator’s part.

Theorem 3.1 (Projection framework). *Given $\epsilon, \delta \geq 0$, a general query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, and a linear equality invariant C with null space \mathcal{N} , if $M : \mathcal{X}^* \rightarrow \mathbb{R}^n$ is (ϵ, δ) -differentially private, then $\mathcal{M}(\mathbf{x}) := A(\mathbf{x}) + \Pi_{\mathcal{N}}(M(\mathbf{x}) - A(\mathbf{x}))$, for all $\mathbf{x} \in \mathcal{X}^*$ is (ϵ, δ) -induced subspace differentially private for query A and invariant C .*

Informally, we conduct projection on a differentially private mechanism in order to 1) remove the noise in the row space of C to respect the invariant constraint C ; and 2) preserve noise in the null space \mathcal{N} and satisfy \mathcal{N} -subspace differential privacy.

Proof of Theorem 3.1. Because for all $\mathbf{x} \in \mathcal{X}^*$, $C\mathcal{M}(\mathbf{x}) = CA(\mathbf{x}) + C\Pi_{\mathcal{N}}(M(\mathbf{x}) - A(\mathbf{x})) = CA(\mathbf{x})$, \mathcal{M} satisfies the invariant C . Since M is \mathcal{N} -subspace (ϵ, δ) -differentially private by Corollary 2.8, and $\Pi_{\mathcal{N}}\mathcal{M}(\mathbf{x})$ equals $\Pi_{\mathcal{N}}M(\mathbf{x})$ in distribution, \mathcal{M} is also \mathcal{N} -subspace differentially private. \square

The projection framework in Theorem 3.1 is particularly useful for revising *additive* mechanisms, having the form

$$\mathcal{M}(\mathbf{x}) = A(\mathbf{x}) + \mathbf{e}, \quad (2)$$

where \mathbf{e} is a noise component independent of $A(\mathbf{x})$. Examples of additive mechanisms include the classic Laplace and Gaussian mechanisms (Dwork et al. 2006b), t - (Nissim, Raskhodnikova, and Smith 2007), double Geometric (Fioretto and Van Hentenryck 2019), and k -norm mechanisms (Hardt and Talwar 2010; Bhaskara et al. 2012). In contrast, the Exponential mechanism (McSherry and Talwar

2007) is in general not additive because the sampling process depends on the utility function non-additively.

When the existing differentially private mechanism M is additive, the projection construction of an induced subspace differentially private mechanism based on M can be simplified, by first sampling the noise \mathbf{e} , and outputting the query value $A(\mathbf{x})$ with the projected noise added to it.

Corollary 3.2. *If the mechanism M in Theorem 3.1 is furthermore additive, i.e. of the form (2), the modified mechanism, $\mathcal{M}(\mathbf{x}) := A(\mathbf{x}) + \Pi_{\mathcal{N}}\mathbf{e}$, is (ϵ, δ) -induced subspace differentially private for query A and invariant C .*

Corollary 3.2 will be useful for the derivation of the projection mechanisms, as well as the various statistical properties of additive mechanisms (including the crucial *unbiasedness* property), to be discussed in the ensuing sections.

3.2 The Extension Framework

The projection framework in Section 3.1 transforms an existing differentially private mechanism to one that is subspace differentially private and respects the invariant C . The extension framework, on the other hand, enables the design of an induced subspace differentially private mechanism without a full mechanism in place yet. Theorem 3.3 discusses how to extend a differential private mechanism with image contained in \mathcal{N} to an induced subspace differentially private mechanism. Moreover, the converse also holds— any induced subspace differentially private mechanism can be written as a differential private mechanism with a translation. Thus, the extension framework provides the optimal trade-off between privacy and accuracy, as Corollary 3.10 will show. We defer the proof to supplementary material.

Theorem 3.3 (Extension framework). *Given $\epsilon, \delta \geq 0$, a general query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, and a linear equality invariant C with null space \mathcal{N} and row space \mathcal{R} , \mathcal{M} is (ϵ, δ) -induced subspace differentially private for query A and invariant C , if and only if $\mathcal{M}(\mathbf{x}) := \hat{M}(\mathbf{x}) + \Pi_{\mathcal{R}}A(\mathbf{x})$ where \hat{M} is (ϵ, δ) -differentially private and its image is contained in $\mathcal{N} \subseteq \mathbb{R}^n$.*

3.3 Induced Subspace Differentially Private Mechanisms for General Queries

We now describe the use of the above two frameworks to construct induced subspace differentially private mechanisms. We first introduce induced subspace differentially private mechanisms with Gaussian noises, then the pure (i.e. $\delta = 0$) versions with Laplace noises. All mechanisms discussed in this section are additive mechanisms, having a general functional form as (2).

Let the ℓ_p sensitivity of the query function $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$ be $\Delta_p(A) = \sup_{\mathbf{x} \sim \mathbf{x}'} \|A(\mathbf{x}) - A(\mathbf{x}')\|_p$, which measures how much a single individual’s data can change the output of the query A . We measure the performance of a mechanism M in terms of the expected squared error. Given any query function $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, the worst-case expected squared error of a mechanism M for A is defined as $\text{err}_M(A) := \sup_{\mathbf{x} \in \mathcal{X}^*} \mathbb{E} [\|M(\mathbf{x}) - A(\mathbf{x})\|_2^2]$.

Gaussian mechanisms Recall the standard Gaussian mechanism, which adds a spherical noise to the output that depends on the ℓ_2 sensitivity of A . By an abuse of notation,

the superscript n over a probability distribution denotes the n -dimensional product distribution with the said marginal.

Lemma 3.4 (Gaussian mechanism (Dwork et al. 2006a)). *For all $\epsilon, \delta > 0$, and general query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, let $c_{\epsilon, \delta} := \epsilon^{-1}(1 + \sqrt{1 + \ln(1/\delta)})$. Then an additive mechanism for A with noise $\mathbf{e}_G(A, \epsilon, \delta) \stackrel{d}{=} N(0; \Delta_2(A)c_{\epsilon, \delta})^n$ where $N(0; \sigma)$ is the unbiased Gaussian distribution with variance σ^2 is (ϵ, δ) -differentially private.*

Given $\epsilon, \delta > 0$, a general query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, and linear equality invariant $C \in \mathbb{R}^{n_c \times n}$, by projection (Theorem 3.1) and extension (Theorem 3.3), we can derive two induced subspace differentially private Gaussian mechanisms. The proofs of Corollaries 3.5 and 3.6 are both deferred to the full version.

Corollary 3.5 (Projected Gaussian mechanism). *An additive mechanism \mathcal{M}_{PG} for A with noise $\Pi_{\mathcal{N}}\mathbf{e}_G(A, \epsilon, \delta)$ where \mathbf{e}_G is defined in Lemma 3.4 is (ϵ, δ) -induced subspace differentially private for query A and invariant C . Moreover, $\text{err}_{\mathcal{M}_{PG}}(A) = (n - n_c)c_{\epsilon, \delta}^2 \Delta_2(A)^2$.*

To apply the extension framework of Theorem 3.3, one complication is how to design a differentially private mechanism with image in \mathcal{N} . To handle this, we first project the query A to the null space \mathcal{N} and have a new query $A_{\mathcal{N}} = Q_{\mathcal{N}}^{\top}A : \mathcal{X}^* \rightarrow \mathbb{R}^{n-n_c}$. Then, compute the sensitivity of $A_{\mathcal{N}}$, $\Delta_2(A_{\mathcal{N}})$, and sample $\mathbf{e}_{\mathcal{N}}$ which consists of $n - n_c$ iid Gaussian noise with variance $(c_{\epsilon, \delta} \Delta_2(A_{\mathcal{N}}))^2$. Finally, convert the noise to the original space \mathbb{R}^n and add the true query $A(\mathbf{x})$. We define the mechanism formally below.

Algorithm 1: Gaussian induced subspace differentially private mechanism through extension

Input: a database \mathbf{x} , a query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, linear equality invariant $C \in \mathbb{R}^{n_c \times n}$ with rank n_c , $\epsilon \in (0, 1)$, and $\delta \in (0, 1)$.

- 1: Compute $Q_{\mathcal{N}} \in \mathbb{R}^{n \times n-n_c}$ an collection of an orthonormal basis of \mathcal{N} , and $A_{\mathcal{N}} := Q_{\mathcal{N}}^{\top}A$.
 - 2: Let $c_{\epsilon, \delta} = \epsilon^{-1}(1 + \sqrt{1 + \ln(1/\delta)})$, and sample $\mathbf{e}_{\mathcal{N}} \stackrel{d}{=} N(0; c_{\epsilon, \delta} \Delta_2(A_{\mathcal{N}}))^{n-n_c}$.
 - 3: **return** $A(\mathbf{x}) + Q_{\mathcal{N}}\mathbf{e}_{\mathcal{N}}$.
-

Corollary 3.6 (Extended Gaussian mechanism). *An additive mechanism \mathcal{M}_{EG} for A with noise $\mathbf{e}_{EG}(A, \epsilon, \delta) = Q_{\mathcal{N}}\mathbf{e}_{\mathcal{N}}$ where $\mathbf{e}_{\mathcal{N}} \stackrel{d}{=} N(0; c_{\epsilon, \delta} \Delta_2(Q_{\mathcal{N}}^{\top}A))^{n-n_c}$ is (ϵ, δ) -induced subspace differentially private for query A and invariant C . Moreover, $\text{err}_{\mathcal{M}_{EG}}(A) = (n - n_c)c_{\epsilon, \delta}^2 \Delta_2(Q_{\mathcal{N}}^{\top}A)^2$.*

Since $Q_{\mathcal{N}}$ consists of orthonormal columns, the ℓ_2 sensitivity of $Q_{\mathcal{N}}^{\top}A$ is less than or equal to the sensitivity of the original query A , and the error in Corollary 3.6 is no more than the error in Corollary 3.5.

Laplace mechanisms Similarly, the standard Laplace mechanism adds independent product Laplace noise to the output that depends on the ℓ_1 sensitivity of A . In what follows, $\text{Lap}(b)$ denotes the univariate Laplace distribution with scale $b > 0$, with density function $\frac{1}{2b} \exp\left(-\frac{|x|}{b}\right)$.

Lemma 3.7 (Laplace mechanism (Dwork et al. 2006b)). *Given $\epsilon > 0$, and a query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, an additive mechanism for A with noise $\mathbf{e}_L(A, \epsilon) \stackrel{d}{=} \text{Lap}(\Delta_1(A)/\epsilon)^n$ is $(\epsilon, 0)$ -differentially private.*

Given $\epsilon > 0$, a query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, and a linear equality invariant $C \in \mathbb{R}^{n_c \times n}$, we have the following two induced subspace differentially private Laplace mechanisms.

Corollary 3.8 (Projected Laplace mechanism). *The additive mechanism for A with noise $\Pi_{\mathcal{N}}\mathbf{e}_L$ where \mathbf{e}_L is defined in Lemma 3.7 is $(\epsilon, 0)$ -induced subspace differentially private for query A and invariant C .*

Corollary 3.9 (Extended Laplace mechanism). *An additive mechanism \mathcal{M}_{LE} for A with noise $\mathbf{e}_{EL}(A, \epsilon) = A(\mathbf{x}) + Q_{\mathcal{N}}w_{\mathcal{N}}$ where $w_{\mathcal{N}} \stackrel{d}{=} \text{Lap}(\Delta_1(Q_{\mathcal{N}}^{\top}A))/\epsilon)^{n-n_c}$ is $(\epsilon, 0)$ -induced subspace differentially private for query A and invariant C .*

We have thus far discussed four mechanisms, respectively derived using the projection and extension frameworks, and employing Gaussian and Laplace errors. In practice, a data curator would choose either projected mechanisms if seeking to impose invariants on an existing differentially private mechanism, and either extension mechanisms if augmenting a smaller private query while staying compatible with the invariants. The curator would prefer the Laplace mechanisms over the Gaussian ones if a pure (i.e. $\delta = 0$) subspace differential privacy guarantee is sought, although at the expense of heavier-tailed noises which may be undesirable for utility purposes. In what follows, we discuss mechanism options for the curator, if utility considerations are the most salient.

3.4 Optimal Induced Subspace Differentially Private Mechanisms for Linear Queries

As a consequence of Theorem 3.3, Corollary 3.10 translates optimal accuracy enjoyed by a differentially private mechanism to optimal accuracy by an induced subspace differentially private mechanism. Let $\text{opt}_{\epsilon, \delta}(A)$ be the optimal error achievable by any (ϵ, δ) -differentially private mechanism, and $\text{opt}_{\epsilon, \delta}^C(A)$ be the optimal error by any (ϵ, δ) -induced subspace differentially private mechanism for query A and invariant C .

Corollary 3.10. *For all $\epsilon, \delta \geq 0$, general query $A : \mathcal{X}^* \rightarrow \mathbb{R}^n$, and linear equality invariant C , $\text{opt}_{\epsilon, \delta}^C(A) = \text{opt}_{\epsilon, \delta}(\Pi_{\mathcal{N}}A)$.*

We defer the proof to the full version. Informally, for any differentially private mechanism we use the extension framework in Theorem 3.3 to construct an induced subspace differentially private mechanism. Because our proof is constructive, we can translate existing near optimal differentially private mechanisms to induced subspace differentially private ones.

We demonstrate this translation with a near-optimal (i.e. mean squared error with a small multiplicative factor) correlated Gaussian mechanism for linear queries from Nikolov, Talwar, and Zhang (2013). Specifically, first design a differential private mechanism \hat{M} for $A_{\mathcal{N}} := Q_{\mathcal{N}}^{\top}A$, and extend

it to a subspace differentially private mechanism by Theorem 3.3. Because the mean squared error is invariant under rotation, $\text{err}_{\hat{M}}(\Pi_{\mathcal{N}}A) = \text{err}_{\hat{M}}(A_{\mathcal{N}})$. Therefore, if \hat{M} is the near optimal correlated Gaussian noise mechanism for $A_{\mathcal{N}}$, the resulting induced subspace differentially private mechanism is also near optimal by Corollary 3.10. Details of this mechanism is spelled out in the full version which we, together with the proof for Theorem 3.11 below, defer to the full version.

Theorem 3.11. *Given $\epsilon, \delta > 0$, a linear query $A : \mathbb{R}^d \rightarrow \mathbb{R}^n$, and a linear equality invariant $C \in \mathbb{R}^{n_c \times n}$, the correlated Gaussian mechanism is an efficient (ϵ, δ) -induced subspace differentially private mechanism such that for all small enough ϵ and all δ small enough with respect to ϵ satisfies*

$$\text{err}_{\mathcal{M}}(A) = O(\log^2(n - n_c) \log 1/\delta) \text{opt}_{\epsilon, \delta}^C(A).$$

We may use the same idea to convert an k -norm mechanism (Hardt and Talwar 2010; Bhaskara et al. 2012) to an $(\epsilon, 0)$ -induced subspace differentially private one. Bhaskara et al. (2012) proposed an $(\epsilon, 0)$ -differentially private k -norm mechanism whose approximation ratio of mean squared error is $O((\log n)^2)$ for any linear query with image in \mathbb{R}^n . We can run the k -norm mechanism on query $A_{\mathcal{N}}$ whose mean squared error is $O((\log(n - n_c))^2) \text{opt}_{\epsilon, 0}(A_{\mathcal{N}})$. Then, by Corollary 3.10, the output can be converted to an $(\epsilon, 0)$ -induced subspace differentially private mechanism, with an approximation ratio $O((\log n - n_c)^2)$.

4 Statistical and Practical Considerations

Unbiasedness of projection algorithms The projection algorithms proposed in this paper, be they Laplace or Gaussian, are provably *unbiased* due to their additive construction. In fact, we have the following result.

Corollary 4.1. *Any mechanism of the form $\mathcal{M}(\mathbf{x}) := A(\mathbf{x}) + \Pi_{\mathcal{N}}\mathbf{e}$ as defined in Corollary 3.2, where \mathbf{e} is random noise with $\mathbb{E}(\mathbf{e}) = 0$, is unbiased in the sense that*

$$\mathbb{E}[\mathcal{M}(\mathbf{x}) \mid A(\mathbf{x})] = A(\mathbf{x}).$$

Corollary 4.1 stands because the conditional expectation of its noise component, $\mathbb{E}[\Pi_{\mathcal{N}}\mathbf{e} \mid A(\mathbf{x})]$, is zero, due to the independence of \mathbf{e} from $A(\mathbf{x})$, and the nature of the projection operation. All projection mechanisms proposed in this paper are of this type, hence are unbiased. As for the proposed extension algorithms, their purpose is to augment existing DP mechanisms in a way that satisfy mandated invariants, thus their unbiasedness hinge on the unbiasedness of the initial mechanism which they extend. For applications in which the data curator has the freedom to design the privacy mechanism from scratch, projection mechanisms are the recommended way to proceed. Indeed, both numerical demonstrations in Section 5 applied to the county-level 2020 Census demonstration data and the spatio-temporal university campus data utilize the projection mechanisms, guaranteeing the unbiasedness of the sanitized data products under their respective invariant constraints.

Transparency and statistical intelligibility Subspace differentially private additive mechanisms carry a special advantage when examined through the lens of downstream statistical analysis of the output query. All Gaussian and Laplace mechanisms examined in this paper (corollaries 3.5, 3.6, 3.8 and 3.9 and theorem 3.11), be they obtained via projection or extension, linearly combine the confidential query with a noise term that is publicly specified. Just like standard differential privacy, mechanisms of subspace differential privacy described in this paper are *transparent* (Abowd and Schmutte 2016), a prized property that brought revolutionary change to the literature of statistical disclosure limitation by ridding obscure procedures. Moreover, the employed noise terms have probability distributions that are fully characterized and independent of the confidential query. This grants the mechanisms *statistical intelligibility* (Gong and Meng 2020), making the output query eligible for both analytical and simulation-based (such as bootstrap) statistical analysis and uncertainty quantification.

In the special case that the original unconstrained differentially private mechanism is spherical Gaussian, defined in Lemma 3.4, the induced subspace differentially private mechanism resulting from projection produces a random query that is distributionally equivalent to that obtained via the standard probabilistic conditioning of the unconstrained mechanism, where the conditioning event is precisely the invariants that the curator seeks to impose.

Theorem 4.2. *If the additive mechanism M in Corollary 3.2 is spherical Gaussian as defined in Lemma 3.4, the corresponding modified mechanism \mathcal{M} has a probability distribution equivalent to the distribution of M conditional on the invariant being true. That is,*

$$\mathcal{M}(\mathbf{x}) \stackrel{d}{=} M(\mathbf{x}) \mid CM(\mathbf{x}) = CA(\mathbf{x}).$$

The proof of Theorem 4.2 is deferred to the full version. The equivalence with conditionalization is particularly valuable for Bayesian inference based on the privatized query, as the analyst may coherently utilize all available information. We note here however, that Theorem 4.2 results from unique properties of the spherical Gaussian distribution. In general, the projection operation aligns closer with marginalization, and cannot produce the equivalent distribution as conditionalization. Nevertheless, the happy statistical consequence of Theorem 4.2 may still be widely impactful, thanks to the ubiquity of the spherical Gaussian mechanism.

Implementation: distributed privatization In local differential privacy, we consider the identity mapping as the query function A , and the private mechanism directly infuses entry-wise noise into the entire confidential dataset before releasing it to the public. The confidential dataset, \mathbf{x} , is often gathered by a number of local agents – nodes, sensors, survey workers – each responsible for one (or more) entries of \mathbf{x} . Distributed privatization can be valuable in local differential privacy, as it ensures individual data contributors’ information is protected the moment it leaves the local agent.

For all additive subspace differentially private mechanisms proposed in this work, distributed privatization may be achieved, if the local agents are capable of simulating

the same noise component. The synchronized simulation can be implemented – hardware permitting – by sharing a common seed across the different local sensors or workers. An instance of distributed privatization is in the full version, which works for arbitrary linear equality invariant C .

5 Numerical Examples

5.1 2020 Census Demonstration Data

We consider the publication of induced subspace differentially private county-level Census population counts, subject to the invariant of state population size, using the November 2020 vintage privacy-protected demonstration files curated by IPUMS NHGIS (Van Riper, Kugler, and Schroeder 2020). These data files link together the original tables from the 2010 Census Summary Files (CSF), here treated as the confidential values, and the trial runs of the Census Bureau’s 2020 Disclosure Avoidance System (DAS) applied to the CSF. All these datasets are publicly available at the cited source. For our demonstration, the privacy loss budget is set to accord exactly to the Census Bureau’s specification, with $\epsilon = 0.192 = 4$ (total) $\times 0.16$ (county level) $\times 0.3$ (population query).

Right panel of Figure 1 showcases the county-level errors from ten runs of the projected Laplace $(\epsilon, 0)$ -induced subspace differentially private mechanism of Corollary 3.8, applied to the counties of Illinois arranged in increasing true population sizes. Compared with the DAS errors (red squares) which show a clear negative bias trend, the proposed mechanism is provably unbiased, due to its additive errors being projected from unbiased and unconstrained random variables. On the other hand, these errors span a similar scale compared to the DAS errors. In the full version, we further show the application of the projected Laplace $(\epsilon, 0)$ -induced subspace differentially private mechanism to an additional ten states, for which the TopDown algorithm incurred decidedly negatively biased errors. Details of how these states were identified are given in Section 1.

5.2 Spatio-temporal Dataset

We consider the publication of time series derived from WiFi log data on connections of mobile devices with nearby access points from a large university campus (Tsinghua University) (Sui et al. 2016) consisting of 3243 fully anonymized individuals and the top 20 most popularly visited buildings in one day.² The raw data recorded whether an individual appears in each of the building in each of the hours on one day. The data were tabulated into hourly time series for 14 clusters of individuals obtained through simple K -means, to represent hypothetical group memberships with distinct travel patterns.

The invariants we consider are of two types, motivated by needs for building management, energy-control, and group activity scheduling: 1) the total number of person-hours spent at each building every hour from all groups, and 2)

²Data was collected under the standard consent for Wifi access on university campus. Interested reader may contact the authors of (Sui et al. 2016) to inquire access to the dataset.

the total number of person-hours spent at each building by every group over 24 hours. The query under consideration is 14 (groups) $\times 24$ (hours) $\times 20$ (location) = 6720 dimensional, subject to a $(24 + 14 - 1) \times 20 = 740$ -dimensional linear constraint.

We apply the projection Gaussian mechanism in Corollary 3.5 with the scale of the elementwise Gaussian noise set to 1. The mechanism is again provably unbiased, but due to the numerous linear constraints imposed, the errors exhibit a slight loss of scale. The median standard deviation of the elementwise additive errors over 50 repetitions is 0.88, with 5% and 95% quantiles at (0.86, 0.91) respectively. Detailed results of the simulation can be found in the full version.

6 Conclusion and Future Work

In this paper, we proposed the concept of subspace differential privacy to make explicit the mandated invariants imposed on private data products, and discussed the projection and extension designs of induced differentially private mechanisms. The invariants we consider are in the form of linear equalities, including sums and contingency table margins as often encountered in applications including the U.S. Decennial Census and spatio-temporal datasets.

An important type of invariants not addressed in this paper are inequalities, such as nonnegativity and relational constraints (e.g. the population size must be larger or equal to the number of households in a geographic area). However, we note that an important premise to the unbiasedness achieved by subspace differentially private mechanisms, as discussed in Section 4, is that the mechanism admits equality invariants only. If inequality invariants must be imposed, unbiased privacy mechanisms can be inherently difficult to design. As Zhu, Hentenryck, and Fioretto (2020) discussed, the bias induced by projection-type post-processing of noisy measurements is attributable to the non-negativity constraints imposed on them. This raises the question of the appropriateness of inequality invariants on the sanitized output, if unbiasedness is simultaneously required.

Also not considered are invariants for binary and categorical attributes, taking values in a discrete space. These invariants differ from real-valued linear equality invariants, because in general they cannot be realized by an additive mechanism with a noise term independent of the confidential data value. While the notion of subspace differential privacy can be extended to these cases, the design of accompanying privacy mechanisms that also enjoy good statistical and implementation properties remains a subject of future research.

Acknowledgments

The authors gratefully acknowledge research support from the National Science Foundation (OAC-1939459, CCF-2118953, CCF-1934924, DMS-1916002, and ISS-2007887).

References

Abowd, J.; Ashmead, R.; Simson, G.; Kifer, D.; Leclerc, P.; Machanavajjhala, A.; and Sexton, W. 2019. Census Top-Down: Differentially Private Data, Incremental Schemas,

and Consistency with Public Knowledge. Technical report, US Census Bureau.

Abowd, J. M. 2018. The US Census Bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2867–2867. ACM.

Abowd, J. M.; and Schmutte, I. M. 2016. Economic analysis and statistical disclosure limitation. *Brookings Papers on Economic Activity*, 2015(1): 221–293.

Ashmead, R.; Kifer, D.; Leclerc, P.; Machanavajjhala, A.; and Sexton, W. 2019. Effective Privacy After Adjusting for Invariants with Applications to the 2020 Census. Technical report, US Census Bureau.

Barak, B.; Chaudhuri, K.; Dwork, C.; Kale, S.; McSherry, F.; and Talwar, K. 2007. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 273–282.

Bhaskara, A.; Dadush, D.; Krishnaswamy, R.; and Talwar, K. 2012. Unconditional differentially private mechanisms for linear queries. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 1269–1284.

City of New York. 2021. 311 Service Requests from 2010 to Present: NYC Open Data.

Dwork, C.; Kenthapadi, K.; McSherry, F.; Mironov, I.; and Naor, M. 2006a. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 486–503. Springer.

Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006b. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, 265–284. Springer.

Fioretto, F.; and Van Hentenryck, P. 2019. Differential privacy of hierarchical census data: An optimization approach. In *International Conference on Principles and Practice of Constraint Programming*, 639–655. Springer.

Gong, R.; and Meng, X.-L. 2020. Congenial differential privacy under mandated disclosure. In *Proceedings of the 2020 ACM-IMS on Foundations of Data Science Conference*, 59–70.

Hardt, M.; and Talwar, K. 2010. On the geometry of differential privacy. In *Proceedings of the forty-second ACM symposium on Theory of computing*, 705–714.

Hay, M.; Rastogi, V.; Miklau, G.; and Suci, D. 2009. Boosting the accuracy of differentially-private histograms through consistency. *arXiv preprint arXiv:0904.0942*.

He, X.; Machanavajjhala, A.; and Ding, B. 2014. Blowfish Privacy: Tuning Privacy-utility Trade-offs Using Policies. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, 1447–1458. New York, NY, USA: ACM.

Kifer, D.; and Lin, B.-R. 2010. Towards an axiomatization of statistical privacy and utility. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 147–158.

Kifer, D.; and Machanavajjhala, A. 2012. A rigorous and customizable framework for privacy. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, 77–88.

Kifer, D.; and Machanavajjhala, A. 2014. Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems (TODS)*, 39(1): 1–36.

Liu, C.; Chakraborty, S.; and Mittal, P. 2016. Dependence Makes You Vulnerable: Differential Privacy Under Dependent Tuples. In *NDSS*, volume 16, 21–24.

McSherry, F.; and Talwar, K. 2007. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, 94–103. IEEE.

Nikolov, A.; Talwar, K.; and Zhang, L. 2013. The geometry of differential privacy: the sparse and approximate cases. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 351–360.

Nissim, K.; Raskhodnikova, S.; and Smith, A. 2007. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, 75–84. ACM.

Rezaei, A.; and Gao, J. 2019. On Privacy of Socially Contagious Attributes. In *Proceedings of the 19th IEEE International Conference on Data Mining (ICDM'19)*, 1294–1299.

Rezaei, A.; Gao, J.; and Sarwate, A. D. 2021. Influencers and the Giant Component: the Fundamental Hardness in Privacy Protection for Socially Contagious Attributes. In *Proceedings of the SIAM International Conference on Data Mining (SDM'2021)*, 217–225.

Song, S.; Wang, Y.; and Chaudhuri, K. 2017. Pufferfish privacy mechanisms for correlated data. In *Proceedings of the 2017 ACM International Conference on Management of Data*, 1291–1306.

Sui, K.; Zhou, M.; Liu, D.; Ma, M.; Pei, D.; Zhao, Y.; Li, Z.; and Moscibroda, T. 2016. Characterizing and improving wifi latency in large-scale operational networks. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, 347–360.

Van Riper, D.; Kugler, T.; and Schroeder, J. 2020. IPUMS NHGIS Privacy-Protected 2010 Census Demonstration Data, version 20201116. Minneapolis, MN: IPUMS.

Wang, H.; and Gao, J. 2020. Distributed Human Trajectory Sensing and Partial Similarity Queries. In *Proceedings of the 19th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN 2020*, 253–264. ACM/IEEE.

Yang, B.; Sato, I.; and Nakagawa, H. 2015. Bayesian Differential Privacy on Correlated Data. In Sellis, T. K.; Davidson, S. B.; and Ives, Z. G., eds., *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, 747–762. ACM.

Yang, D.; Qu, B.; Yang, J.; and Cudre-Mauroux, P. 2019. Revisiting user mobility and social relationships in lbsns: a hypergraph embedding approach. In *The World Wide Web Conference*, 2147–2157.

Yang, D.; Qu, B.; Yang, J.; and Cudre-Mauroux, P. 2020. LBSN2Vec++: Heterogeneous Hypergraph Embedding for Location-Based Social Networks. *IEEE Transactions on Knowledge and Data Engineering*.

Zhu, K.; Hentenryck, P. V.; and Fioretto, F. 2020. Bias and Variance of Post-processing in Differential Privacy. *CoRR*, abs/2010.04327.

Zhu, T.; Xiong, P.; Li, G.; and Zhou, W. 2014. Correlated differential privacy: Hiding information in non-IID data set. *IEEE Transactions on Information Forensics and Security*, 10(2): 229–242.