

PRIVATE QUERY RELEASE VIA THE JOHNSON-LINDENSTRAUSS TRANSFORM

ALEKSANDAR NIKOLOV

Department of Computer Science, University of Toronto
e-mail address: anikolov@cs.toronto.edu

ABSTRACT. We introduce a new method for releasing answers to statistical queries with differential privacy, based on the Johnson-Lindenstrauss lemma. The key idea is to randomly project the query answers to a lower dimensional space so that the distance between any two vectors of feasible query answers is preserved up to an additive error. Then we answer the projected queries using a simple noise-adding mechanism, and lift the answers up to the original dimension. Using this method, we give, for the first time, purely differentially private mechanisms with optimal worst case sample complexity under average error for answering a workload of k queries over a universe of size N . As other applications, we give the first purely private efficient mechanisms with optimal sample complexity for computing the covariance matrix of a bounded high-dimensional distribution, and for answering 2-way marginal queries. We also show that, up to the dependence on the error, a variant of our mechanism is nearly optimal for every given query workload.

1. INTRODUCTION

One of the central problems in private data analysis is to release answers to statistical (also known as linear) queries on the data. Here, a statistical query is defined by a function $q : \mathcal{X} \rightarrow \mathbb{R}$, and the value of the query on a dataset $X = (x_1, \dots, x_n) \in \mathcal{X}^n$ is simply its average over data points, which we write as $q(X) = \frac{1}{n} \sum_{i=1}^n q(x_i)$. Counting queries, which ask what fraction of the dataset satisfies some predicate, are a special case of statistical queries, and are themselves of significant interest. Many organizations report summary statistics in the form of tables, where each cell in the table is the answer to a counting query. This is the case, for example, for many of the tables released by official statistics agencies such as the US Census Bureau: see (MMHM18) for a worked out example. As other examples, statistical queries capture CDFs of one-dimensional distributions and higher dimensional generalizations, mean estimation, and loss gradients in empirical risk minimization. Thus, when we analyze data about people whose privacy needs to be protected, the question of privately computing answers to statistical queries is of fundamental interest.

The privacy framework we adopt is differential privacy (DMNS06), which provides strong semantic guarantees for protecting the privacy of individuals represented in a dataset.

Key words and phrases: differential privacy, dimension reduction, Johnson-Lindenstrauss, query release, K-norm mechanism.

* An extended abstract of this paper was published in SODA 2023.

We say that a randomized algorithm \mathcal{M} (usually called a mechanism) that takes datasets $X \in \mathcal{X}^n$ is (ε, δ) -differentially private if for any two datasets X, X' that differ in at most one element, and for any measurable event S in the range of \mathcal{M} , we have

$$\Pr(\mathcal{M}(X) \in S) \leq e^\varepsilon \Pr(\mathcal{M}(X') \in S) + \delta.$$

The setting when $\delta = 0$ is usually called pure differential privacy, and the setting when $\delta > 0$ is called approximate differential privacy. In this paper, we focus on purely differentially private mechanisms, although our basic technique applies to approximate differential privacy, as well as to intermediate privacy notions (DR16; BS16; Mir17; DRS22). Pure differential privacy is the strictest of the standard variants of differential privacy, and has some nice properties not shared by approximate differential privacy and other relaxations: it is a single parameter definition; satisfies a simple and tight composition theorem (composing private mechanisms just requires “adding the epsilons”); and it satisfies a strong group privacy property, i.e., some measure of privacy protection is automatically offered to small groups of people and not just to individuals. These benefits, on the other hand, can sometimes come at the cost of increased error. Understanding when this is the case requires completely understanding how much error a purely differentially private algorithm needs to introduce for a given statistical task.

Coming back to statistical queries, we can see from the examples in the first paragraph that, typically, we are interested in releasing answers to many statistical queries, rather than in just a single query. Let us call a sequence of statistical queries $Q = (q_1, \dots, q_k)$ a workload. The following are some of the fundamental questions one can ask about releasing answers to a workload of statistical queries under differential privacy:

- (1) How much error is necessary, in the worst case, to answer every workload of k statistical queries over a universe \mathcal{X} of size N given a dataset of size n ?
- (2) How can we characterize the optimal trade-off between error and dataset size for a given workload Q ?
- (3) Can we achieve the optimal worst case error in question 1., or the optimal error vs dataset size tradeoff in question 2. using an efficient mechanism?

There are many variants of these questions, depending on what flavor of differential privacy we adopt (pure, or approximate, or concentrated differential privacy), and what measure of error we choose (worst case error, or average error, which we define as root mean squared error). Yet, there are few settings in which precise answers are known. A tight answer for question 1. in the setting of approximate differential privacy was given by Bun, Ullman, and Vadhan (BUV14). Question 2. is open for any variant of differential privacy and measure of error that has been studied, but there have been some partial approximate characterizations (HT10; BDKT12; NTZ16; Nik15; BBNS19). The answer to question 3. depends on how the workload is represented, but is generally wide open: for example, we do not even know an efficient mechanism that achieves the optimal error vs dataset size tradeoff for width-3 marginal queries.

In this paper we focus on the most basic of the three questions, question 1., in the case of pure differential privacy. Surprisingly, in this setting, the question is open, both for worst case error and for average error. We introduce a new algorithmic technique for query release, based on the Johnson-Lindenstrauss transform, and we show that it resolves question 1. for pure differential privacy and average error.

In order to discuss our results and prior work more precisely, we now introduce some notation. Let us use $Q(X)$ to denote the vector of query answers $(q_1(X), \dots, q_k(X))$. Let us

say the queries in the workload are bounded if, for each $i \in [k]$ and each $x \in \mathcal{X}$, $|q_i(x)| \leq 1$. We define, respectively, the worst case (or ℓ_∞), and the average (or root mean squared, or ℓ_2) error of a mechanism \mathcal{M} on a workload Q and datasets of size n as

$$\begin{aligned} \text{err}^{\ell_\infty}(\mathcal{M}, Q, n) &= \sup_{X \in \mathcal{X}^n} \mathbb{E} \|\mathcal{M}(X) - Q(X)\|_\infty; \\ \text{err}^{\ell_2}(\mathcal{M}, Q, n) &= \sup_{X \in \mathcal{X}^n} \mathbb{E} \left[\frac{1}{\sqrt{k}} \|\mathcal{M}(X) - Q(X)\|_2 \right]. \end{aligned}$$

Above, we assume that on input X , the mechanism \mathcal{M} outputs a random k -dimensional vector $\mathcal{M}(X)$, and the expectation is taken with respect to the mechanism's randomness. The sample complexity of the mechanism on the workload Q is the smallest dataset size for which the mechanism can achieve error at most α , denoted $\text{sc}^{\ell_p}(\mathcal{M}, Q, \alpha) = \inf\{n : \text{err}^{\ell_p}(\mathcal{M}, Q, n) \leq \alpha\}$, where $p \in \{2, \infty\}$. Note that $\text{err}^{\ell_2}(\mathcal{M}, Q, n) \leq \text{err}^{\ell_\infty}(\mathcal{M}, Q, n)$, and, therefore, $\text{sc}^{\ell_2}(\mathcal{M}, Q, \alpha) \leq \text{sc}^{\ell_\infty}(\mathcal{M}, Q, \alpha)$. Finally, we define the optimal sample complexity of a workload Q under ε -differential privacy by

$$\text{sc}_\varepsilon^{\ell_p}(Q, \alpha) = \inf\{\text{sc}^{\ell_p}(\mathcal{M}, Q, \alpha) : \mathcal{M} \text{ is } \varepsilon\text{-differentially private}\}.$$

Now, question 1. can be formulated precisely as the problem of giving tight bounds on the worst case value of $\text{sc}_\varepsilon^{\ell_p}(Q, \alpha)$ over workloads Q of k bounded queries on a universe of size N . The Laplace noise mechanism (DMNS06), a variant of the K-norm mechanism (SU17), and the SmallDB mechanism (BLR13) show that¹

$$\text{sc}_\varepsilon^{\ell_2}(Q, \alpha) \lesssim \min \left\{ \frac{k}{\varepsilon\alpha}, \frac{\log(N)}{\varepsilon\alpha^3} \right\}; \quad \text{sc}_\varepsilon^{\ell_\infty}(Q, \alpha) \lesssim \min \left\{ \frac{k}{\varepsilon\alpha}, \frac{\log(k)\log(N)}{\varepsilon\alpha^3} \right\}.$$

Stronger bounds are known using methods tailored to average error. The K-norm mechanism,² and a combination of the K-norm mechanism and the projection mechanism (NTZ16) give the bounds

$$\text{sc}_\varepsilon^{\ell_2}(Q, \alpha) \lesssim \min \left\{ \frac{\log^{3/2}(k)\sqrt{k\log(N)}}{\varepsilon\alpha}, \frac{\log^2(k)\log^{3/2}(N)}{\varepsilon\alpha^2} \right\}.$$

The first term in the minimum on the right hand side can be boosted to an upper bound for $\text{sc}_\varepsilon^{\ell_\infty}(Q, \alpha)$ using non-private boosting for queries (see Section 6.1 in (NTZ16)). We do not know how to boost the second term, since the error per query is private information, and private boosting incurs too big of a loss in the sample complexity in the pure differential privacy setting. The second term is the only known sample complexity bound in the pure differential privacy setting which has quadratic dependence on $\frac{1}{\alpha}$, and polylogarithmic dependence on the number of queries and universe size. Unfortunately, the powers of the logarithmic terms are not known to be tight. In particular, the best known lower bounds on sample complexity come from packing arguments (HT10; Har11; De12), and show that

¹Here and in the rest of the paper we use the notation $A \lesssim B$ to signify that there exists an absolute constant $C > 0$ such that $A \leq CB$. We use $A \gtrsim B$ to denote $B \lesssim A$.

²The first bound is not stated in (HT10), but follows directly from their results and known estimates of the volume of a polytope contained in a ball.

there exists a workload Q of k bounded queries on a universe of size N such that

$$\begin{aligned} \text{sc}_\varepsilon^{\ell_2}(Q, \alpha) &\gtrsim \min \left\{ \frac{k}{\varepsilon\alpha}, \frac{\sqrt{k \log N}}{\varepsilon\alpha}, \frac{\log(N)}{\varepsilon\alpha^2} \right\}; \\ \text{sc}_\varepsilon^{\ell_\infty}(Q, \alpha) &\gtrsim \min \left\{ \frac{k}{\varepsilon\alpha}, \frac{\sqrt{k \log N}}{\varepsilon\alpha}, \frac{\log(k) \log(N)}{\varepsilon\alpha^2} \right\}. \end{aligned}$$

We note that these appear to be the best lower bounds that can be proved with packing arguments. Thus, closing the gap between upper and lower bounds would require either new algorithmic ideas, or developing a lower bound method for pure differential privacy that goes beyond packing.

1.1. Results. In this paper, we introduce a simple new algorithmic idea for the query release problem. While our main technique applies to any flavor of differential privacy, we focus on its application in the setting of pure differential privacy, where it gives a complete answer to question 1. above for average error. In particular, we prove the following theorem.

Theorem 1.1. *For any workload Q of k bounded queries on a universe of size N , we have*

$$\text{sc}_\varepsilon^{\ell_2}(Q, \alpha) \lesssim \min \left\{ \frac{k}{\varepsilon\alpha}, \frac{\sqrt{k \log N}}{\varepsilon\alpha}, \frac{\log N}{\varepsilon\alpha^2} \right\}. \quad (1.1)$$

Moreover, this sample complexity is achieved by a mechanism running in time polynomial in $k, N, n, \frac{1}{\varepsilon}$.

The sample complexity bounds in (1.1) are tight for random queries (see Section 7.1), and resolve Open Problem 2. from (NU21). The first term on the right hand side of (1.1) is known, and can be shown using either the Laplace noise or the K-norm mechanism. To get the second term, we optimize the privacy budget of the K-norm mechanism in order to remove unnecessary logarithmic terms. Our main contribution is in establishing the last term by giving a new private query release mechanism, which we call the JL mechanism. The key new idea is to use a variant of the Johnson-Lindenstrauss lemma to first randomly project the query answers to a lower dimensional space, then add noise in the lower dimensional space, and finally lift up the answers to the original dimension. To explain why this gives the sample complexity in (1.1), let us define the sensitivity polytope K_Q of the workload Q , given by $K_Q = \text{conv}\{Q(x) : x \in \mathcal{X}\}$, where $Q(x) = (q_1(x), \dots, q_k(x))$. K_Q contains all possible values of $Q(X)$ for all possible dataset sizes, and conv is the convex hull operator. A theorem of Liaw, Mehrabian, Plan, and Vershynin (LMPV17) guarantees that multiplication by a suitable random matrix $T \in \mathbb{R}^{\ell \times k}$ with $\ell \lesssim \frac{\log N}{\alpha^2}$ preserves all distances between points in K_Q up to $\pm\alpha\sqrt{k}$. Therefore, to achieve error $O(\alpha)$, it suffices to answer the projected queries $TQ(X)$ with an answer vector $\tilde{Y} \in TK_Q$ with error $O(\alpha)$, and output a vector $\hat{Y} \in K_Q$ such that $T\hat{Y} = \tilde{Y}$. This works because, by the result of Liaw et al., the distance between \hat{Y} and $Q(X)$ is approximately the same as the distance between \tilde{Y} and $TQ(X)$ with high probability. The benefit of using the random projection (i.e., multiplication by T) is that now we only have to answer the ℓ -dimensional workload TQ , for which we can use the K-norm mechanism, giving us error α with sample complexity $\frac{\sqrt{\ell \log N}}{\varepsilon\alpha} \lesssim \frac{\log N}{\varepsilon\alpha^2}$, as was our goal. The only caveat is that the K-norm mechanism does not necessarily produce answers in TK_Q , and multiplication by T is only promised to preserve distances in K_Q . This

problem has a simple fix: a least squares projection of the noisy answers back onto TK_Q does not increase the error.

We note that the Johnson-Lindenstrauss lemma has been used before in differential privacy in other contexts, e.g. (BBDS12; KKMM13; Sta21). To the best of our knowledge, the application to query release we present here is new. Our algorithm is similar to a single iteration of the gradient descent algorithm in (Kas21), except for the choice of noise. The analysis required for query release vs. optimization is different, however.

A more careful analysis of this JL mechanism allows giving the sample complexity bounds in terms of a natural geometric property of the queries. Intuitively, we expect that the sample complexity required to compute Q with differential privacy scales with the “size” of K_Q . In many situations (NTZ16; BBNS19), the appropriate measure of size appears to be the Gaussian mean width, defined for a set $S \subseteq \mathbb{R}^k$ by

$$w(S) = \mathbb{E} \sup_{y \in S} \langle y, G \rangle,$$

where G is a standard k -dimensional Gaussian vector. The novel bounds in Theorem 1.1 are implied and refined by the next theorem. The theorem also refines the computational complexity guarantee in Theorem 1.1, by showing that our mechanism runs in time polynomial in $k, n, \frac{1}{\epsilon}$, and in the running time of a separation oracle for K_Q . Such an oracle can be implemented in time polynomial in k and N using linear programming.

Theorem 1.2. *For any workload of k queries Q for which the diameter of K_Q is bounded, and any $0 \leq \alpha \leq \frac{\text{diam}(K_Q)}{\sqrt{k}}$, there exists an ϵ -differentially private mechanism \mathcal{M} with sample complexity*

$$\text{sc}^{\ell_2}(\mathcal{M}, Q, \alpha) \lesssim \min \left\{ \frac{w(K_Q)}{\alpha \epsilon}, \frac{w(K_Q)^2}{k \epsilon \alpha^2} \right\}.$$

Moreover, given a separation oracle for K_Q , this sample complexity is achieved by a mechanism running in time polynomial in $k, n, \frac{1}{\epsilon}$, and in the running time of the separation oracle.

To see that Theorem 1.2 is a refinement of Theorem 1.1, we note that when Q is a workload of k bounded queries over a universe of size N , a standard concentration of measure argument shows that $w(K_Q) \lesssim \sqrt{k \log N}$. Moreover, since for bounded queries $K_Q \subseteq [-1, +1]^k$, we also have $w(K_Q) \leq w([-1, +1]^k) \lesssim k$. These two bounds and Theorem 1.2 give Theorem 1.1. For some workloads, Theorem 1.2 can give tighter bounds than Theorem 1.1. Moreover, it can apply in situations when the universe \mathcal{X} is not finite. We will give one such example soon. To prove the theorem, we first show, using Urysohn’s inequality tying Gaussian width and volume, that the error of the K-norm mechanism can be controlled in terms of the Gaussian width of K_Q . This gives the first term in the sample complexity bound. We then apply the JL mechanism with this noise bound, together with the fact that multiplication by the random matrix T does not asymptotically increase Gaussian width in expectation, which we prove.

Theorems 1.1 and 1.2 give improved, and in fact tight sample complexity bounds for natural and important workloads, e.g., for workloads of constant width marginal queries, and workloads deriving from covariance estimation. Let $\mathcal{X} = \{0, 1\}^d$ for a positive integer d , and, for $S \subseteq [d]$, define the query $q_S(x) = \prod_{i \in S} x_i$, i.e., the conjunction of the bits indexed by S . The workload of width- ω marginals over d -dimensional boolean data, for a positive integer ω , is defined as $Q_{\omega, d} = \{q_S : S \subseteq [d], |S| = \omega\}$. Marginals capture many public releases

of statistics in the form of tables, and privately approximating marginal queries is one of the better studied problems in differential privacy (BCD⁺07; HRS12; CKKL12; CTUW14; TUV12; DNT15). Using Theorem 1.1, we immediately get new sample complexity bounds for computing marginals with pure differential privacy. In particular, (1.1) implies that, for any constant integer ω ,

$$\text{sc}_\varepsilon^{\ell_2}(Q_{\omega,d}, \alpha) \lesssim \min \left\{ \frac{d^\omega}{\varepsilon\alpha}, \frac{d^{(\omega+1)/2}}{\varepsilon\alpha}, \frac{d}{\varepsilon\alpha^2} \right\}. \quad (1.2)$$

Note the the first term in the minimum is never smaller than the second term. The third term is new to this paper. We also prove a lower bound showing that (1.2) is tight.

As another application, we also study the problem of estimating the covariance of a bounded distribution in the Frobenius norm. The next theorem crucially uses the refined Theorem 1.2, since covariance matrices are naturally modeled using statistical queries on an infinite domain.

Theorem 1.3. *There exists an ε -differentially private mechanism that, given n samples from a distribution μ on the unit Euclidean ball B_2^d with covariance matrix Σ , and outputs a matrix $\widehat{\Sigma}$ such that, for any $\alpha \in (0, 1)$,*

$$\mathbb{E} \|\widehat{\Sigma} - \Sigma\|_F \leq \alpha$$

as long as $n \geq n_0$ where

$$n_0 \lesssim \left(\min \left\{ \frac{d^{1.5}}{\alpha\varepsilon}, \frac{d}{\alpha^2\varepsilon} \right\} + \frac{1}{\alpha^2} \right).$$

Moreover, the mechanism runs in time polynomial in $d, n, \frac{1}{\varepsilon}$.

Theorem 1.3 improves on the recent work of Dong, Liang, and Yi (DLY22), who give an algorithm with sample complexity on the order of $\frac{d^{1.5}}{\alpha^2\varepsilon}$. Moreover, it shows that the optimal sample complexity bound (1.2) can be achieved in time polynomial in d when $w = 2$, since 2-way marginals can be modeled as a special case of covariance estimation. Such efficient algorithms are not known for $w \geq 3$.

So far, we argued that the JL mechanism is optimal for worst-case query workloads. We can in fact show that, in the constant error regime (e.g., when we want to achieve error $\alpha = \frac{1}{100}$), the sample complexity of a variant of the JL mechanism is also tight with respect to the optimal sample complexity for the given workload. This gives a partial answer to question 2. from the beginning of the introduction for pure differential privacy. The next theorem, capturing this optimality guarantee, follows from combining Theorem 1.2 with the methods of Blasiok, Bun, Nikolov, and Steinke (BBNS19).

Theorem 1.4. *For any workload Q of k queries over a universe of size N , and any $0 \leq \alpha \leq \frac{\text{diam}(K_Q)}{\sqrt{k}}$, there exists an ε -differentially private mechanism \mathcal{M} with running time polynomial in $k, N, n, \frac{1}{\varepsilon}$, and with sample complexity*

$$\text{sc}^{\ell_2}(\mathcal{M}, Q, \alpha) \lesssim \frac{\text{diam}(K_Q)}{\alpha\sqrt{k}} \log \left(\frac{\text{diam}(K_Q)}{\alpha\sqrt{k}} \right)^2 \text{sc}_\varepsilon^{\ell_2}(Q, \alpha/8).$$

In particular, when k is a workload of bounded queries,

$$\text{sc}^{\ell_2}(\mathcal{M}, Q, \alpha) \lesssim \frac{1}{\alpha} \log \left(\frac{1}{\alpha} \right)^2 \text{sc}_\varepsilon^{\ell_2}(Q, \alpha/8).$$

Let us remark that the bounded queries result follows immediately from the more general result, since, by definition, bounded queries take values in $[-1, +1]$, so $K_Q \subseteq [-1, +1]^k$, and, therefore, $\text{diam}(K_Q) \leq 2\sqrt{k}$. In the case of bounded queries, Theorem 1.4 implies that, as long as α is a constant, no other mechanism can achieve much better sample complexity and more than a constant factor better error at the same time. It is important here to emphasize that this holds for an arbitrarily small constant α . A similar result was proved previously for pure differential privacy by Roth for the SmallDB mechanism (Rot11) (see also (BS16; Vad17)). Our algorithm runs in time polynomial in N , whereas Roth's algorithm needs to enumerate over a set of size N^{1/α^2} . A similar result is also known for concentrated differential privacy (BBNS19).

The main idea in the proof of Theorem 1.4 is to run the JL mechanism on an approximation $K_\alpha \subseteq K_Q$ of K_Q . We construct K_α by taking the convex hull of a net of the extreme points of K_Q . This guarantees that any point in K_Q is within distance $\ll \alpha\sqrt{k}$ from some point in K_α . At the same time, in some cases, the Gaussian width of K_α can be a lot smaller than that of K_Q , giving us better error via Theorem 1.2.

2. PRELIMINARIES

In this section we record some notation and preliminaries that we need in the rest of the paper.

In terms of general notation, we use $\|x\|_p = (\sum_{i=1}^k |x_i|^p)^{1/p}$ to denote the ℓ_p norm of a vector $x \in \mathbb{R}^k$. We use $B_p^k = \{x \in \mathbb{R}^k : \|x\|_p \leq 1\}$ to denote the unit ℓ_p ball in \mathbb{R}^k . We use $\|A\|_F = \sqrt{\text{tr}(AA^T)}$ for the Frobenius norm of a matrix A , which is just its ℓ_2 norm when treated as a vector.

For a set $S \subseteq \mathbb{R}^k$ and a real number t , we use $tS = \{tx : x \in S\}$ to denote scaling. The Minkowski sum of two sets $S, T \subseteq \mathbb{R}^k$ is denoted $S + T = \{x + y : x \in S, y \in T\}$, and the special case when one of the sets is a singleton is denoted $x + S$ rather than the more cumbersome $\{x\} + S$.

We use $A \otimes B$ to denote the Kronecker product of two matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{k \times \ell}$, which is a $mk \times n\ell$ matrix whose rows and columns are indexed, respectively, by pairs of rows and columns of A and B . The entry $(i, p), (j, q)$ of $A \otimes B$ is equal to $a_{i,j}b_{p,q}$. We also use $A^{\otimes w}$ for the w -fold Kronecker product of A with itself. An important property of Kronecker products is that $(A \otimes B)(x \otimes y) = (Ax) \otimes (By)$ for matrices A, B and vectors x, y of compatible dimensions.

2.1. Statistical Queries and Differential Privacy. Let us recall the main notation on statistical queries from the Introduction. We consider datasets $X \in \mathcal{X}^n$, and query workloads $Q = \{q_1, \dots, q_k\}$, where each q_i is defined by a function $q_i : \mathcal{X} \rightarrow \mathbb{R}$, and, overloading notation, the query itself is defined by

$$q_i(X) = \frac{1}{n} \sum_{x \in X} q_i(x).$$

We say the queries in the workload are bounded if, for each $i \in [k]$ and each $x \in \mathcal{X}$, $|q_i(x)| \leq 1$. Again overloading notation, we write $Q(X)$ for the vector $(q_1(X), \dots, q_k(X))$, and, for any

$x \in \mathcal{X}$, we write $Q(x)$ for $(q_1(x), \dots, q_k(x))$. Then, approximating $Q(X)$ is equivalent to approximating the mean of k -dimensional vectors:

$$Q(X) = \frac{1}{n} \sum_{x \in X} Q(x).$$

We define the sensitivity polytope $K_Q = \text{conv}\{Q(x) : x \in \mathcal{X}\}$, containing all possible values of $Q(X)$ for all possible dataset sizes. In general, we may allow \mathcal{X} to be infinite, as long as the set $\{Q(x) : x \in \mathcal{X}\}$ is bounded; in case \mathcal{X} is infinite, K_Q is defined as the closure of $\text{conv}\{Q(x) : x \in \mathcal{X}\}$, and is more appropriately called the sensitivity body.

We define the ℓ_p error of a mechanism \mathcal{M} on a workload Q and datasets of size n as

$$\text{err}^{\ell_p}(\mathcal{M}, Q, n) = \max_{X \in \mathcal{X}^n} \mathbb{E} \left[\frac{1}{\sqrt{k}} \|\mathcal{M}(X) - Q(X)\|_p \right].$$

Above, we assume that on input X , the mechanism \mathcal{M} outputs a random k -dimensional vector $\mathcal{M}(X)$. The sample complexity of the mechanism on queries Q is the smallest dataset size for which the mechanism can achieve error at most α , denoted

$$\text{sc}^{\ell_p}(\mathcal{M}, Q, \alpha) = \inf\{n : \text{err}(\mathcal{M}, Q, n) \leq \alpha\}.$$

The optimal sample complexity of a workload Q under ε -differential privacy is given by

$$\text{sc}_\varepsilon^{\ell_p}(Q, \alpha) = \inf\{\text{sc}^{\ell_p}(\mathcal{M}, Q, \alpha) : \mathcal{M} \text{ is } \varepsilon\text{-differentially private}\}.$$

In the rest of the paper we use $\text{err}(\mathcal{M}, Q, n)$, $\text{sc}(\mathcal{M}, Q, \alpha)$, and $\text{sc}_\varepsilon(Q, \alpha)$ without superscripts to denote, respectively, the ℓ_2 error and the corresponding sample complexity functions.

In the rest of the paper we call an ε -differentially private mechanism for answering a workload Q of k queries on datasets of size n efficient if it runs in polynomial time in k, n , and $\frac{1}{\varepsilon}$ when given an evaluation oracle for each query in Q , and a separation oracle for the sensitivity polytope K_Q .

We also mention two important properties of differential privacy: composition and invariance under post-processing, both captured by the following lemma. See the monograph of Dwork and Roth (DR14) for a proof.

Lemma 2.1. *If $\mathcal{M}_1(X)$ is an ε_1 -differentially private mechanism, and $\mathcal{M}_2(X, y)$ is ε_2 -differentially private with respect to X for each y in the range of \mathcal{M}_1 , then the composition $\mathcal{M}(X) = \mathcal{M}_2(X, \mathcal{M}_1(X))$ is $(\varepsilon_1 + \varepsilon_2)$ -differentially private. In particular, for any ε -differentially private mechanism \mathcal{M} , and any, potentially randomized, function f on the range of \mathcal{M} , $f(\mathcal{M}(X))$ is ε -differentially private.*

2.2. Packing Lower Bounds. The following packing lower bound for statistical queries is likely folklore, and was used in the work of Hardt and Talwar (HT10) (see also Theorem 2.1 in (De12)).

Lemma 2.2. *Suppose that there exists a workload Q of k queries, and datasets $X_1, \dots, X_M \in \mathcal{X}^n$ such that for any $i \neq j$, $i, j \in [M]$, $\frac{1}{\sqrt{k}} \|Q(X_i) - Q(X_j)\|_2 > \alpha$, and X_i and X_j differ in at most Δ elements. Then, as long as $\Delta \leq \frac{\log(M/2)}{\varepsilon}$, we have $\text{sc}_\varepsilon(Q, \alpha/2) \geq n$.*

The next lemma is a direct consequence of Lemma 2.2, and is often, but not always how the packing bound is used. In it, we use the notion of a separation number: for a compact set $K \subseteq \mathbb{R}^k$ and a centrally symmetric convex set L , the separation number $P(K, L)$ is the maximum size of a finite set $S \subseteq K$ such that for all distinct $y, y' \in S$ we have $y \notin y' + L$.

Lemma 2.3. *Let Q be a workload of k queries, and let $S_Q = \{Q(x) : x \in \mathcal{X}\}$ be the set of vectors of query answers on each universe element. Then for any $\alpha \in (0, t]$, $sc_\varepsilon(Q, \alpha/2) \gtrsim \frac{t \log(P(S_Q, t\sqrt{k}B_2^k)/2)}{\alpha\varepsilon}$.*

Proof. Let x_1, \dots, x_M achieve $P(S_Q, t\sqrt{k}B_2^k)$. This means that, for any two distinct x_i, x_j , $\frac{1}{\sqrt{k}}\|Q(x_i) - Q(x_j)\|_2 > t$. Let $n = \frac{t \log(M/2)}{\alpha\varepsilon}$. We pick an arbitrary x_0 , and define the dataset X_i to contain $\frac{\alpha}{t}n$ copies of x_i and $(1 - \frac{\alpha}{t})n$ copies of x_0 . Then $\Delta = \frac{\log(M/2)}{\varepsilon}$, and, for any $i \neq j$, $\frac{1}{\sqrt{k}}\|Q(X_i) - Q(X_j)\|_2 > \alpha$. We can now apply Lemma 2.2 to finish the proof. \square

2.3. A Johnson-Lindenstrauss Lemma with Additive Error. Our algorithms rely on a variant of the Johnson-Lindenstrauss (JL) lemma with additive error, due to Liaw, Mehrabian, Plan, and Vershynin (LMPV17). Their result shows that the pairwise distances between elements in any bounded (but potentially infinite) set are preserved under suitable random projection onto a lower dimensional subspace, up to an additive approximation. Before stating this result precisely, we recall the definitions of subgaussian random variables and random vectors, and of the associated ψ_2 Orlicz norm.

The ψ_2 norm of, respectively, a real-valued random variable A , and an ℓ -dimensional random vector V , are defined by

$$\|A\|_{\psi_2} = \inf\{t : \mathbb{E}\psi_2(A/z) \leq 1\}; \quad \|V\|_{\psi_2} = \sup_{\theta \in \mathbb{R}^\ell: \|\theta\|_2=1} \|\langle V, \theta \rangle\|_{\psi_2},$$

where $\psi_2(a) = e^{a^2} - 1$. The random variable, respectively random vector, is *subgaussian* if its ψ_2 norm is finite. Background on subgaussian random variables, including other equivalent definitions, their properties, and examples, can be found in Vershynin's book (Ver18). Here we mention two examples:

- A Gaussian random vector $G \sim N(0, \sigma^2 I)$ is subgaussian with $\|G\|_{\psi_2} \lesssim \sigma$;
- A scaled symmetric Bernoulli random vector B , where B_i is chosen independently and uniformly in $\{-\sigma, +\sigma\}$, is subgaussian with $\|B\|_{\psi_2} \lesssim \sigma$.

Next we define the class of random projections for which the additive JL lemma is known to hold.

Definition 2.4. A random matrix $T \in \mathbb{R}^{\ell \times k}$ is a *JL-matrix with parameter C* if it has independent rows T_1, \dots, T_ℓ , each satisfying

$$\mathbb{E}T_i T_i^\top = \frac{1}{\ell} I; \quad \|T_i\|_{\psi_2} \leq \frac{C}{\sqrt{\ell}}.$$

Examples of JL-matrices with parameter $C \lesssim 1$ include matrices with independent entries, each entry picked either from $N(0, \frac{1}{\ell})$, or uniformly from $\{-\frac{1}{\sqrt{\ell}}, \frac{1}{\sqrt{\ell}}\}$.

We can now state the additive JL lemma that we use.

Theorem 2.5 ((LMPV17)). *Let $T \in \mathbb{R}^{\ell \times k}$ be a JL matrix with parameter C . Then, for any bounded set $S \subseteq \mathbb{R}^k$, we have that*

$$\mathbb{E} \sup_{y, y' \in S} \left| \|Ty - Ty'\|_2 - \|y - y'\|_2 \right| \lesssim \frac{C^2 w(S)}{\sqrt{\ell}}.$$

In the remainder of this section, we derive a useful consequence of Theorem 2.5. We prove that, for any bounded set S , the expected Gaussian width of TS is bounded by the Gaussian width of S . To this end, we need a lemma about Gaussian processes, stating that contracting distances does not increase the Gaussian mean width. For a short proof of a slightly weaker lemma (which would suffice for our application), see Corollary 3.14 in Ledoux and Talagrand's book (LT11).

Lemma 2.6 (Theorem 3.15 (LT11)). *Let $S \subseteq \mathbb{R}^k$ be a bounded set, and let $f : \mathbb{R}^k \rightarrow \mathbb{R}^\ell$ be a function such that*

$$\forall y, z \in S : \|f(y) - f(z)\|_2 \leq C\|y - z\|_2.$$

Then $w(f(S)) \leq Cw(S)$.

We use Lemma 2.6 to prove a similar result when we allow an additive distortion in pairwise distances.

Lemma 2.7. *Let $S \subseteq \mathbb{R}^k$ be a bounded set, and let $f : \mathbb{R}^k \rightarrow \mathbb{R}^\ell$ be a function such that*

$$\forall y, y' \in S : \|f(y) - f(y')\|_2 \leq \|y - y'\|_2 + \alpha$$

for some $\alpha \geq 0$. Then $w(f(S)) \lesssim w(S) + \alpha\sqrt{\ell}$.

Proof. Let $T \subseteq S$ be an inclusion maximal set such that $f(T)$ is a 2α -separated subset of $f(S)$. I.e., for any two distinct points $y, y' \in T$, we have $\|f(y) - f(y')\|_2 > 2\alpha$, and, moreover, we cannot add any point of S to T while still satisfying this property. This means that for any $z \in f(S)$ there is some $z' \in f(T)$ such that $\|z - z'\|_2 \leq 2\alpha$; let us fix one such $z' \in f(T)$ for each $z \in f(S)$ and denote it $\pi(z)$. Taking G to be a standard Gaussian random vector in \mathbb{R}^ℓ , we have

$$\begin{aligned} w(f(S)) &= \mathbb{E} \sup_{z \in f(S)} \langle z, G \rangle \\ &\leq \mathbb{E} \sup_{z \in f(S)} \langle \pi(z), G \rangle + \mathbb{E} \sup_{z \in f(S)} \langle z - \pi(z), G \rangle \\ &\leq \mathbb{E} \sup_{z' \in f(T)} \langle z', G \rangle + \sup_{z \in f(S)} \|z - \pi(z)\|_2 \mathbb{E} \|G\|_2 \end{aligned} \tag{2.1}$$

$$\leq w(f(T)) + 2\alpha\sqrt{\ell}. \tag{2.2}$$

Above, inequality (2.1) follows by the Cauchy-Schwarz inequality, and (2.2) follows from Jensen's inequality, which implies $\mathbb{E} \|G\|_2 \leq \sqrt{\mathbb{E} \|G\|_2^2} \leq \sqrt{\ell}$.

At the same time, we can use Lemma 2.6 with T and $f(T)$, because $f(T)$ is well-separated. Indeed, for any two distinct $y, y' \in T$, by assumption, and because $f(T)$ is 2α -separated, we have

$$\|y - y'\|_2 \geq \|f(y) - f(y')\|_2 - \alpha > \frac{1}{2} \|f(y) - f(y')\|_2.$$

Therefore, by Lemma 2.6,

$$w(f(T)) \leq 2w(T) \leq 2w(S),$$

with the final inequality following since $T \subseteq S$. Combining this bound with (2.2) proves the lemma. \square

Lemma 2.8. *Let $T \in \mathbb{R}^{\ell \times k}$ be a JL-matrix with parameter C . Then, for any bounded set $S \subseteq \mathbb{R}^k$, we have*

$$\mathbb{E}w(TS) \lesssim C^2w(S).$$

Proof. Lemma 2.7 implies that pointwise for each $T \in \mathbb{R}^{\ell \times k}$ we have

$$w(TK) \lesssim w(K) + \left(\sup_{y, y' \in K} \left| \|Ty - Ty'\|_2 - \|y - y'\|_2 \right| \right) \sqrt{\ell}.$$

Taking expectations of both sides over T and using Theorem 2.5, we get

$$\mathbb{E}w(TK) \lesssim w(K) + \frac{C^2w(K)}{\sqrt{\ell}}\sqrt{\ell} = (1 + C^2)w(K).$$

This finishes the proof, since $C \gtrsim 1$ for Definition 2.4 to hold (as follows, e.g., from item (ii) of Proposition 2.5.2 in (Ver18)). \square

3. THE JL MECHANISM

We consider algorithms that follow a general template, presented in Algorithm 1. Our main lemma for the analysis of this algorithm is the following.

Algorithm 1: The JL-Release mechanism template

Public input: A workload Q of k queries over the universe \mathcal{X} ;
Private input: A dataset $X \in \mathcal{X}^n$;
 Pick a JL-matrix $T \in \mathbb{R}^{\ell \times k}$ with parameter $C \lesssim 1$;
 $\tilde{Y} \leftarrow TQ(X) + \frac{1}{n}Z$;
 /* $Z \in \mathbb{R}^k$ is additive noise that ensures \tilde{Y} is ε -differentially private. */
 Output $\hat{Y} \in \arg \min \left\{ \|\tilde{Y} - Ty\|_2 : y \in K_Q \right\}$;

Lemma 3.1. *Let \mathcal{M}_Z be the mechanism described in Algorithm 1, where T is a JL-matrix with parameter C . Then there exists an absolute constant $C' > 0$ such that*

$$\text{err}(\mathcal{M}_Z, Q, n) \leq \frac{C'C^2w(K_Q)}{\sqrt{k\ell}} + \frac{1}{\sqrt{kn}}\mathbb{E}\|Z\|_2. \quad (3.1)$$

Moreover, if the mechanism that outputs $\tilde{Y} = TQ(X) + \frac{1}{n}Z$ on input X satisfies ε -differential privacy (resp. (ε, δ) -differential privacy, ρ -zCDP, f -DP) with respect to X , then so does \mathcal{M}_Z . \mathcal{M}_Z is also efficient assuming that Z can be sampled efficiently.

Proof. We first claim that

$$\|T\hat{Y} - TQ(X)\|_2 \leq \|\tilde{Y} - TQ(X)\|_2 = \frac{1}{n}\|Z\|_2. \quad (3.2)$$

This is a standard fact about ℓ_2 projection onto a convex set. Indeed, by first-order optimality conditions, the set $\{y \in \mathbb{R}^k : \langle T\hat{Y} - \tilde{Y}, Ty - T\hat{Y} \rangle \geq 0\}$ contains all of K_Q . Since $Q(X) \in K_Q$, this means that

$$\langle \tilde{Y} - T\hat{Y}, T\hat{Y} - TQ(X) \rangle = \langle T\hat{Y} - \tilde{Y}, TQ(X) - T\hat{Y} \rangle \geq 0.$$

We have

$$\begin{aligned} \|\tilde{Y} - TQ(X)\|_2^2 &= \|\tilde{Y} - T\hat{Y} + T\hat{Y} - TQ(X)\|_2^2 \\ &= \|\tilde{Y} - T\hat{Y}\|_2^2 + 2\langle \tilde{Y} - T\hat{Y}, T\hat{Y} - TQ(X) \rangle + \|T\hat{Y} - TQ(X)\|_2^2 \\ &\geq \|T\hat{Y} - TQ(X)\|_2^2. \end{aligned}$$

This proves (3.2). From (3.2) and Theorem 2.5, we have that, for an absolute constant $C' > 0$,

$$\begin{aligned} \mathbb{E}\|\hat{Y} - Q(X)\|_2 &\leq \mathbb{E}\left|\|T\hat{Y} - TQ(X)\|_2 - \|\hat{Y} - Q(X)\|_2\right| + \mathbb{E}\|T\hat{Y} - TQ(X)\|_2 \\ &\leq \mathbb{E} \sup_{y, y' \in K_Q} \left| \|Ty - Ty'\|_2 - \|y - y'\|_2 \right| + \frac{1}{n} \mathbb{E}\|Z\|_2 \end{aligned} \quad (3.3)$$

$$\leq \frac{C' C^2 w(K_Q)}{\sqrt{\ell}} + \frac{1}{n} \mathbb{E}\|Z\|_2. \quad (3.4)$$

Above, in inequality (3.3) we use (3.2) and the fact that \hat{Y} and $Q(X)$ are both elements of K_Q , and in inequality (3.4) we use Theorem 2.5. The desired bound on error now follows by dividing through by \sqrt{k} .

The privacy guarantee follows directly from the post-processing property of differential privacy (Lemma 2.1). Efficiency follows since $\|\tilde{Y} - Ty\|_2$, being a convex function, can be minimized over $y \in K_Q$ in time polynomial in k given a separation oracle for K_Q (see, e.g., the Frank-Wolfe implementation of this projection step in (DNT15)). \square

Note that, if T is a JL-matrix with parameter $C \lesssim 1$, then, in order to make the first term in the error bound (3.1) at most $\frac{\alpha}{2}$, it is enough to choose $\ell \gg \frac{w(K_Q)^2}{k\alpha^2}$.

Finally, let us remark when T is deterministically chosen to be the identity matrix, Algorithm 1 reduces to the Projection Mechanism from (NTZ16), i.e., we just add noise in order to preserve differential privacy, and perform a least squares projection of the noisy query answers onto K_Q . In the case of (ϵ, δ) -differential privacy and bounded queries on a universe of size N , the Projection Mechanism, instantiated with Gaussian noise, achieves sample complexity (with respect to ℓ_2 error) on the order of $\frac{\sqrt{\log N \log 1/\delta}}{\alpha^2 \epsilon}$, which is also optimal for worst-case queries (NTZ16; BUV14). The same guarantee also follows directly from Lemma 3.1 when Z is chosen to be Gaussian noise. Nevertheless, it is interesting that, in that case, the dimension reduction step appears to be unnecessary. We do not know if it is necessary for pure differential privacy, i.e., if the Projection Mechanism can recover optimal sample complexity with noise achieving pure differential privacy. Nikolov, Talwar, and Zhang analyzed the Projection Mechanism with K-norm noise (NTZ16), but their bound is suboptimal. The main hurdle to getting an optimal bound using their approach is that one-dimensional marginals of K-norm noise are not subgaussian.

4. ADDITIVE NOISE DISTRIBUTIONS FOR PURE PRIVACY

Our goal is to instantiate Algorithm 1 with different choices of the noise Z . To this end, we recall and refine the K -norm mechanisms, introduced by Hardt and Talwar (HT10).

Let K be a convex body (convex bounded set with non-empty interior) in \mathbb{R}^k , which is symmetric around the origin, i.e., $K = -K$. Then K is the unit ball of the norm $\|\cdot\|_K$,

defined by $\|y\|_K = \inf\{t \geq 0 : y \in tK\}$. The K -norm distribution $\mu_{K,\varepsilon}$ is a probability distribution on \mathbb{R}^k with probability density function

$$f_K(y) = \frac{\varepsilon^k e^{\varepsilon\|y\|_K}}{\Gamma(k+1) \text{vol}(K)}.$$

We have the following result, which is a very slight strengthening of a similar result (Theorem 4.3) in (HT10). Hardt and Talwar stated it for the special case $K = K_Q$, but the same proof establishes the lemma stated below. Here, and in the rest of the paper, we use the notation $K - K$ as a shorthand for the Minkowski sum $K + (-K)$, which is defined as the set $\{y - z : y \in K, z \in K\}$.

Lemma 4.1 ((HT10)). *If $K_Q - K_Q \subseteq K$, then the K -norm mechanism, which outputs $Q(X) + \frac{1}{n}Z$ for $Z \sim \mu_{K,\varepsilon}$ on input $X \in \mathcal{X}^n$, satisfies ε -differential privacy. Moreover,*

$$\mathbb{E}\|Z\|_2 \leq \frac{k+1}{\varepsilon} \mathbb{E}\|U\|_2 \quad \mathbb{E}\|Z\|_2^2 \leq \frac{(k+2)(k+1)}{\varepsilon^2} \mathbb{E}\|U\|_2^2, \quad (4.1)$$

where U is a random vector distributed uniformly in K . Therefore, the error of the K -norm mechanism is bounded by $\frac{k+1}{\sqrt{kn\varepsilon}} \mathbb{E}\|U\|_2 \lesssim \frac{\sqrt{k}}{n\varepsilon} \mathbb{E}\|U\|_2$.

The next lemma follows easily from Lemma 4.1 and the observation that $K_Q - K_Q \subseteq \text{diam}(K_Q)B_2^k$.

Lemma 4.2. *For any workload Q of k queries, and $Z \sim \mu_{\text{diam}(K_Q)B_2^k,\varepsilon}$, the mechanism \mathcal{M} that outputs $Q(X) + \frac{1}{n}Z$ is efficient, ε -differentially private and has error*

$$\text{err}(\mathcal{M}, Q, n) \lesssim \frac{\sqrt{k}}{n\varepsilon} \text{diam}(K_Q).$$

In particular, if Q consists of bounded queries, the mechanism has error $\text{err}(\mathcal{M}, Q, n) \lesssim \frac{k}{n\varepsilon}$.

While the error in Lemma 4.2 is optimal in some cases, it can also be far from optimal when the size of the data universe \mathcal{X} is sub-exponential in dimension. In the remainder of this section, we describe a noise mechanism with better error in this latter setting. The mechanism closely follows the one given by Hardt and Talwar, but with a different choice of how to spend the privacy budget, in order to remove unnecessary logarithmic factors. The error bound and efficiency guarantees of this optimized mechanism are given by the following lemma.

Lemma 4.3. *For any workload Q of k queries there exists a noise random variable Z such that the mechanism \mathcal{M} that outputs $Q(X) + \frac{1}{n}Z$ is efficient, ε -differentially private and has error*

$$\text{err}(\mathcal{M}, Q, n) \lesssim \frac{w(K_Q)}{n\varepsilon}.$$

In particular, if Q consists of bounded queries over a universe of size N , then the mechanism has error $\text{err}(\mathcal{M}, Q, n) \lesssim \frac{\sqrt{k} \log N}{n\varepsilon}$.

Towards proving Lemma 4.3, let us recall the volume lower bound from Hardt and Talwar's paper. In formulating it, we use the notion of volume radius, defined for a convex body $K \subseteq \mathbb{R}^k$ as $\text{vrad}(K) = \left(\frac{\text{vol}(K)}{\text{vol}(B_2^k)}\right)^{1/k}$. In other words, the volume radius of K is the radius of a Euclidean ball with the same volume as K . It will also be convenient to use similar terminology for bounded convex sets that lie in a lower dimensional subspace of

\mathbb{R}^k . Let us write $\text{vol}_W(\cdot)$ for volume (i.e., Lebesgue measure) restricted to a subspace W of \mathbb{R}^k . Then for a subspace W of dimension ℓ , and a convex body $K \subseteq W$ (i.e., K is convex and bounded, and its interior relative to W is non-empty), we define the volume radius relative to W as $\text{vrad}_W(K) = \left(\frac{\text{vol}_W(K)}{\text{vol}(B_2^\ell)} \right)^{1/\ell}$. We are now ready to formulate the volume lower bounds for a convex body K in \mathbb{R}^k , given by

$$\text{volLB}(K, \ell) = \ell \sup_{W \in G_{k, \ell}} \text{vrad}_W(P_W K).$$

Above $G_{k, \ell}$ is the set of all ℓ -dimensional subspaces of \mathbb{R}^k , and P_W is the orthogonal projection onto the subspace W . Hardt and Talwar constructed an efficiently sampleable noise distribution whose error can be bounded in terms of $\max_\ell \text{volLB}(K, \ell)$. Using the error guarantees of their noise distribution as a black box, together with classical volume estimates, we can prove an error upper bound weaker than the one in Lemma 4.3 by a polylogarithmic factor. In order to give a tight bound, we open up and optimize Hardt and Talwar's construction.

The next lemma is the key technical tool in proving Lemma 4.3. The lemma restates the key claims in the analysis of Hardt and Talwar's mechanism for non-isotropic bodies (HT10, Section 7). While Hardt and Talwar's result was conditional on the hyperplane conjecture,³ Bhaskara, Dadush, Krishnaswamy, and Talwar showed how to make the result unconditional (BDKT12).

Lemma 4.4 ((HT10; BDKT12)). *For any symmetric convex body $K \subseteq \mathbb{R}^k$, there exists subspaces $W_0 = \mathbb{R}^k, W_1, \dots, W_m$ and symmetric bounded convex sets K_1, \dots, K_m with the following properties:*

- $\bigcup_{i=1}^m W_i$ spans \mathbb{R}^k ;
- for all $i \geq 2$, W_i is contained in the orthogonal complement of W_{i-1} ;
- for all $i \geq 1$, $\dim W_i = \left\lceil \frac{1}{2} \left(k - \sum_{j=1}^{i-1} \dim W_j \right) \right\rceil$;
- for all $i \geq 1$, $P_{W_i} K \subseteq K_i \subset W_i$, and K_i is a convex body inside W_i , i.e., has non-empty interior relative to W_i ;
- for all $1 \leq i \leq m-1$, and for U_i distributed uniformly in K_i ,

$$(\dim W_i)^2 \mathbb{E} \|U_i\|_2^2 \lesssim \text{volLB} \left(K, k - \sum_{j=1}^i \dim W_j \right)^2; \quad (4.2)$$

- moreover, for W_m we have $\dim W_m = 1$ and $\mathbb{E} \|U_m\|_2^2 \lesssim \text{volLB}(K, 1)^2$;
- given a sampling oracle for K , we can sample from each K_i in time polynomial in k .

Next we use Lemma 4.4 to give an optimized version of Hardt and Talwar's noise distribution.

Lemma 4.5. *For any workload Q of k queries there exists a noise random variable Z such that the mechanism \mathcal{M} that outputs $Q(X) + \frac{1}{n}Z$ is efficient, ε -differentially private, and has error*

$$\text{err}(\mathcal{M}, Q, n) = \frac{1}{\sqrt{kn}} \mathbb{E} \|Z\|_2 \lesssim \frac{1}{\sqrt{kn}\varepsilon} \left(\sum_{i=1}^{\lceil \log_2 k \rceil} \text{volLB}(K_Q - K_Q, [2^{-i}k])^{2/3} \right)^{3/2}.$$

³Recently Klartag and Lehec proved the conjecture up to a polylogarithmic factor (KL22).

Proof. Let W_1, \dots, W_m and K_1, \dots, K_m be as in Lemma 4.4, used with $K = K_Q - K_Q$. Let $k_i = \lfloor 2^{-i}k \rfloor$ for $i \in \{0, \dots, m\}$. An easy induction argument shows that $k - \sum_{j=1}^i \dim W_j = k_i$, $k_{m-1} = 1$, and $k_m = 0$. This means that $\dim W_i = k_{i-1} - k_i$.

We define Z to be distributed as $Z_1 + \dots + Z_m$, where Z_1, \dots, Z_m are independent random variables, and each Z_i is supported on W_i and distributed according to μ_{K_i, ε_i} , for a value of ε_i to be chosen shortly. (This choice of ε_i is the only place where our noise distribution differs from Hardt and Talwar's.) Since the subspaces W_1, \dots, W_m are pairwise orthogonal, using Jensen's inequality we have

$$\mathbb{E}\|Z\|_2 \leq (\mathbb{E}\|Z\|_2^2)^{1/2} = \left(\sum_{i=1}^m \mathbb{E}\|Z_i\|_2^2 \right)^{1/2} \lesssim \left(\sum_{i=1}^m \frac{(k_{i-1} - k_i)^2}{\varepsilon_i^2} \mathbb{E}\|U_i\|_2^2 \right)^{1/2}. \quad (4.3)$$

The final inequality above follows from (4.1). We choose $\varepsilon_1, \dots, \varepsilon_m$ so that $\varepsilon_1 + \dots + \varepsilon_m = \varepsilon$, which guarantees that the algorithm is ε -differentially private, as we prove below. The choice that satisfies this constraint while minimizing the right hand side of (4.3) is

$$\varepsilon_i = \frac{\varepsilon((k_{i-1} - k_i)^2 \mathbb{E}\|U_i\|_2^2)^{1/3}}{\sum_{i=1}^m ((k_{i-1} - k_i)^2 \mathbb{E}\|U_i\|_2^2)^{1/3}},$$

and gives

$$\mathbb{E}\|Z\|_2 \lesssim \frac{1}{\varepsilon} \left(\sum_{i=1}^m ((k_{i-1} - k_i)^2 \mathbb{E}\|U_i\|_2^2)^{1/3} \right)^{3/2}. \quad (4.4)$$

Note that this choice of $\varepsilon_1, \dots, \varepsilon_m$ can be computed efficiently by estimating $\mathbb{E}\|U_i\|_2^2$ through sampling. In more detail, let $U_{i,1}, \dots, U_{i,L}$ be independent copies of U_i . By a result of Adamczak, Litvak, Pajor, and Tomczak-Jaegermann (ALPTJ10), if $L \geq L_0$ for some $L_0 \lesssim k \log(2/\beta)^4 \log(\log 2/\beta)^2$, we have that, with probability at least $1 - \beta$, for any $\theta \in \mathbb{R}^k$

$$\frac{1}{2} \mathbb{E}\langle U_i, \theta \rangle^2 \leq \frac{1}{L} \sum_{j=1}^L \langle U_{i,j}, \theta \rangle^2 \leq \frac{3}{2} \mathbb{E}\langle U_i, \theta \rangle^2.$$

Taking θ to range over the standard basis of \mathbb{R}^k and adding up the resulting inequalities, we get that, with probability at least $1 - \beta$,

$$\frac{3}{2} \mathbb{E}\|U_i\|_2^2 \leq \frac{1}{L} \sum_{j=1}^L \|U_{i,j}\|_2^2 \leq \frac{3}{2} \mathbb{E}\|U_i\|_2^2.$$

It follows that, via sampling polynomially many times from the distribution of U_i , we can approximate $\mathbb{E}\|U_i\|_2^2$ up to a multiplicative constant with high probability. Computing such an approximation for each $i \in [m]$ allows us to efficiently approximate the optimal ε_i above up to a multiplicative constant, and, therefore, match the error guarantee (4.4) up to a constant.

Then the whole algorithm is efficient by Lemma 4.4.

Let us return to the analysis of the error of our noise adding mechanism. Using (4.2) for any $1 \leq i \leq m - 1$, we have

$$(k_{i-1} - k_i)^2 \mathbb{E}\|U_i\|_2^2 \lesssim \text{volLB}(K_Q - K_Q, k_i)^2.$$

Moreover, for $i = m$ we get

$$\mathbb{E}\|U_i\|_2^2 \lesssim \text{volLB}(K_Q - K_Q, 1)^2 = \text{volLB}(K_Q - K_Q, k_{m-1})^2.$$

Together with (4.4), these bounds give us

$$\mathbb{E}\|Z\|_2 \lesssim \frac{1}{\varepsilon} \left(\sum_{i=1}^{m-1} \text{volLB}(K_Q - K_Q, k_i)^{2/3} \right)^{3/2}. \quad (4.5)$$

This implies the claimed bound on the error of the mechanism \mathcal{M} .

To argue privacy, notice that $\mathcal{M}(X) = Q(X) + \frac{1}{n}Z$ is distributed identically to

$$\left(P_{W_1}Q(X) + \frac{1}{n}Z_1 \right) + \dots + \left(P_{W_m}Q(X) + \frac{1}{n}Z_m \right).$$

Here we use the fact that $P_{W_1} + \dots + P_{W_m} = I$ since the subspaces W_1, \dots, W_m are pairwise orthogonal and span \mathbb{R}^k . Let us define a mechanism $\mathcal{M}_i(X) = P_{W_i}Q(X) + \frac{1}{n}Z_i$ for every $i \in [m]$. Observe that $K_{P_{W_i}Q} - K_{P_{W_i}Q} = P_{W_i}(K_Q - K_Q)$, and so, by Lemma 4.1, \mathcal{M}_i is ε_i -differentially private. We can see \mathcal{M} as a postprocessing of the composition of $\mathcal{M}_1, \dots, \mathcal{M}_m$, and, therefore, by Lemma 2.1, \mathcal{M} satisfies differential privacy with privacy parameter $\varepsilon_1 + \dots + \varepsilon_m = \varepsilon$. \square

A final tool we need in the proof of Lemma 4.3 is Urysohn's inequality, which states that for any bounded set $K \subseteq \mathbb{R}^k$, we have

$$\text{vrad}(K) \leq \frac{w(K)}{\sqrt{k}}. \quad (4.6)$$

The inequality follows from the Brunn-Minkowski inequality, and the rotational invariance of Gaussian measure. For a proof, see Chapter 1 of Pisier's book (Pis89).

We can now use the inequality (4.6) and the fact that mean width does not increase under projection to give a bound on the volume lower bounds, and, thereby, give an upper bound on the expression for the error in Lemma 4.5.

Lemma 4.6. *For any convex body $K \subseteq \mathbb{R}^k$, and any $1 \leq \ell \leq k$, we have $\text{volLB}(K, \ell) \leq \sqrt{\ell}w(K)$.*

Proof. The lemma follows directly from (4.6), and the observation that $w(PK) \leq w(K)$ for any orthogonal projection P . The latter fact is well-known, and follows, for example, from Lemma 2.6. \square

Proof of Lemma 4.3. Let us take Z to be the noise distribution guaranteed by Lemma 4.5 for the workload Q . By Lemma 4.5, the mechanism \mathcal{M} that outputs $Q(X) + \frac{1}{n}Z$ is efficient and ε -differentially private. It remains to analyze its error, i.e., to bound $\frac{1}{\sqrt{kn}}\mathbb{E}\|Z\|_2$. By Lemmas 4.5 and 4.6, we have

$$\begin{aligned} \text{err}(\mathcal{M}, Q, n) &= \frac{1}{\sqrt{kn}}\mathbb{E}\|Z\|_2 \lesssim \frac{1}{\sqrt{kn}\varepsilon} \left(\sum_{i=1}^{\lfloor \log_2 k \rfloor} \left(\sqrt{2^{-i}k}w(K_Q - K_Q) \right)^{2/3} \right)^{3/2} \\ &= \frac{w(K_Q - K_Q)}{n\varepsilon} \left(\sum_{i=1}^{\lfloor \log_2 k \rfloor} 2^{-i/3} \right)^{3/2} \\ &\lesssim \frac{w(K_Q - K_Q)}{n\varepsilon}. \end{aligned}$$

In the final line, we used that the sum $\sum_{i=1}^{\infty} 2^{-i/3}$ converges. We then observe that, for a standard Gaussian G in \mathbb{R}^k ,

$$\begin{aligned} w(K_Q - K_Q) &= \mathbb{E} \sup_{y,z \in K_Q} \langle G, y - z \rangle = \mathbb{E} \sup_{y \in K_Q} \langle G, y \rangle + \mathbb{E} \sup_{z \in K_Q} \langle G, -z \rangle \\ &= \mathbb{E} \sup_{y \in K_Q} \langle G, y \rangle + \mathbb{E} \sup_{z \in K_Q} \langle -G, z \rangle \\ &= 2w(K_Q). \end{aligned}$$

In the final equality above we used the fact that G and $-G$ are identically distributed.

The bound for the case when Q is a workload of bounded queries follows from the estimate

$$w(K_Q) = \mathbb{E} \sup_{y \in K_Q} \langle G, y \rangle = \mathbb{E} \sup_{x \in \mathcal{X}} \langle G, Q(x) \rangle \lesssim \sqrt{k \log N}, \quad (4.7)$$

where G is a standard Gaussian in \mathbb{R}^k , the second equality holds because $\sup_{y \in K_Q} \langle G, y \rangle$ is achieved at a vertex of K_Q , and the inequality holds from standard Gaussian process estimates (see, e.g., Example 7.5.10 in (Ver18)) and because for bounded queries $\|Q(x)\|_2 \leq \sqrt{k}$ for all $x \in \mathcal{X}$. \square

5. INSTANTIATION OF THE JL MECHANISM

Theorem 5.1. *For any workload of k queries Q over a universe of size N , and any $0 \leq \alpha \leq \frac{\text{diam}(K_Q)}{\sqrt{k}}$, there exists an efficient ε -differentially private mechanism \mathcal{M} with sample complexity*

$$\text{sc}(\mathcal{M}, Q, \alpha) \lesssim \frac{w(K_Q)^2}{k\alpha^2\varepsilon}.$$

In particular, if Q is a workload of k bounded queries over a universe of size N , we have $\text{sc}(\mathcal{M}, Q, \alpha) \lesssim \frac{\log N}{\varepsilon\alpha^2}$.

Proof. The mechanism \mathcal{M} applies Lemma 3.1 to Q , with the noise Z sampled as in Lemma 4.3 applied to the queries TQ , where T is any $\ell \times k$ JL-matrix with parameter $C \lesssim 1$. By Lemmas 3.1 and 4.3, \mathcal{M} is ε -differentially private, efficient, and its error is bounded as

$$\text{err}(\mathcal{M}, Q, n) \lesssim \frac{w(K_Q)}{\sqrt{k\ell}} + \frac{\sqrt{\ell}}{\sqrt{kn\varepsilon}} \mathbb{E}w(TK_Q) \lesssim \frac{w(K_Q)}{\sqrt{k\ell}} + \frac{\sqrt{\ell}w(K_Q)}{\sqrt{kn\varepsilon}},$$

where we used Lemma 2.8 in the second inequality to conclude that $\mathbb{E}w(TK_Q) \lesssim w(K_Q)$. The right hand side is minimized at $\ell = n\varepsilon$, giving error

$$\text{err}(\mathcal{M}, Q, n) \lesssim \frac{w(K_Q)}{\sqrt{kn\varepsilon}},$$

and the sample complexity bound follows from setting n so that the right hand side is at most α . The bound for the case when Q is a workload of bounded queries follows from the estimate (4.7). \square

Theorem 1.2 follows from Lemma 4.3, and Theorem 5.1. Note that this requires deciding efficiently whether to use the mechanism in Lemma 4.3 or the one in Theorem 5.1. For this, it is enough to approximate $w(K_Q)$. We can get such an approximation by sampling L independent k -dimensional Gaussian vectors G_1, \dots, G_L , and computing $\frac{1}{L} \sum_{i=1}^L \sup_{y \in K_Q} \langle y, G_i \rangle$. This quantity can be computed efficiently given a separation oracle for K_Q by the equivalence of separation and optimization (GLS88). By the Gaussian concentration inequality for Lipschitz functions (Theorem 5.2.2 in (Ver18)), with probability at least $1 - \beta$ for an absolute constant $c > 0$, we have

$$\left| \frac{1}{L} \sum_{i=1}^L \sup_{y \in K_Q} \langle y, G_i \rangle - w(K_Q) \right| \lesssim \frac{\sqrt{\log(2/\beta)} \operatorname{diam}(K_Q)}{\sqrt{L}} \lesssim \frac{\sqrt{\log(2/\beta)} w(K_Q)}{\sqrt{L}}.$$

Thus, setting β to be an arbitrarily small constant, and L to be a large enough multiple of $\log(2/\beta)$, allows us to approximate $w(K_Q)$ within a multiplicative constant with arbitrarily high probability. We can then use this approximation to decide which of the two mechanisms to run, while only losing a constant factor in the error guarantee.

Theorem 1.1 follows from Lemmas 4.2 and 4.3, and Theorem 5.1, since a separation oracle for K_Q can be implemented in time polynomial in N and k using linear programming.

6. EFFICIENT MECHANISMS FOR 2-WAY MARGINALS AND COVARIANCE ESTIMATION

As an application, we consider the problem of estimating the covariance of a distribution μ on B_2^d in Frobenius norm, and prove Theorem 1.3. Note that in this section we do not necessarily normalize our error.

Theorem 1.3, as well as an efficient mechanism for computing 2-way marginals, follows from the following lemma.

Lemma 6.1. *There exists an ε -differentially private mechanism \mathcal{M} such that, given a dataset $X = (x_1 x_1^T, \dots, x_n x_n^T)$ with $x_1, \dots, x_n \in rB_2^d$, \mathcal{M} outputs a $d \times d$ matrix $\mathcal{M}(X)$ satisfying*

$$\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n x_i x_i^T - \mathcal{M}(X) \right\|_F \leq \alpha r$$

as long as $n \geq n_0$, where $n_0 \lesssim \min \left\{ \frac{d^{1.5}}{\alpha \varepsilon}, \frac{d}{\alpha^2 \varepsilon} \right\}$. Moreover, the mechanism runs in time polynomial in $d, n, \frac{1}{\varepsilon}$.

Proof. Let us denote by $M_2(X) = \frac{1}{n} \sum_{i=1}^n x_i x_i^T$ the second moment matrix of X . We notice that $M_2(X)$ is equivalent to a workload of d^2 statistical queries on the universe B_2^d . The sensitivity polytope (or, more appropriately in this case, sensitivity body) of M_2 is $K_{M_2} = \operatorname{conv}\{xx^T : x \in rB_2^d\}$. To estimate M_2 , we use the better mechanism among the one in Lemma 4.3, and the one in Theorem 5.1. Since the Frobenius norm is simply the ℓ_2 norm on matrices, seen as elements in \mathbb{R}^{d^2} , Lemma 4.3 and Theorem 5.1 imply that we can achieve $\mathbb{E} \|M_2(X) - \mathcal{M}(X)\|_F \leq \alpha r$ as long as

$$n \geq C \min \left\{ \frac{dw(K_{M_2})}{\alpha r \varepsilon}, \frac{w(K_{M_2})^2}{\alpha^2 r^2 \varepsilon} \right\}, \quad (6.1)$$

for a sufficiently large constant C . To show that this implies the sample complexity bound in the statement of the theorem, we need to bound the Gaussian width $w(K_{M_2})$. Noticing

that the standard inner product $\langle A, B \rangle$ of two matrices A and B can be written as $\text{tr}(AB^T)$, we get

$$w(K_{M_2}) = \mathbb{E} \sup_{x \in rB_2^d} \text{tr}(Gxx^T) = r^2 \mathbb{E} \sup_{x \in B_2^d} x^T G x = r^2 \mathbb{E} \|G\|,$$

with G a $d \times d$ matrix of independent standard Gaussians, and $\|G\|$ denoting its operator norm. It is a classical fact (see, e.g., Theorem 4.4.5 in (Ver18)) that $\mathbb{E} \|G\| \lesssim \sqrt{d}$, and this, together with (6.1) finishes the proof of the sample complexity bound.

Finally, to prove efficiency, we need to exhibit a separation oracle for $K_{M_2} = \text{conv}\{xx^T : x \in rB_2^d\}$. We claim the alternative characterization $K_{M_2} = \{A \in \mathbb{R}^{d \times d} : A \succeq 0, \text{tr} A \leq r^2\}$, where the notation $A \succeq 0$ means that A is positive semidefinite. On the one hand, it is clear that for any $A \in K_{M_2}$, $A \succeq 0$ and $\text{tr} A \leq r^2$, since these conditions hold for any extreme point xx^T , $x \in rB_2^d$, of K_{M_2} , and are maintained under taking convex combinations. On the other hand, if $A \succeq 0$ and $\text{tr} A \leq r^2$, then, by the spectral theorem, $A = \sum_{i=1}^d \lambda_i (ra_i)(ra_i)^T$, where the $\lambda_i \geq 0$ are the eigenvalues of $\frac{1}{r^2}A$, and $a_i \in B_2^d$ are the eigenvectors. Moreover $\sum_i \lambda_i = \frac{\text{tr}(A)}{r^2} \leq 1$, so this spectral decomposition in fact gives a representation of A as convex combination of points in $\{xx^T : x \in rB_2^d\}$. With this alternative characterization, we see that K_{M_2} can be written as a feasibility SDP, so it has a polynomial time separation oracle. \square

First we use Lemma 6.1 to prove Theorem 1.3.

Proof of Theorem 1.3. A standard symmetrization argument shows that, if the samples from \mathcal{D} are $X = (X_1, \dots, X_n)$, and $\Sigma(X)$ is their empirical covariance matrix, then

$$\mathbb{E} \|\Sigma(X) - \Sigma\|_F \lesssim \frac{1}{\sqrt{n}}.$$

This means that the non-private sample complexity of the problem is at most on the order of $\frac{1}{\alpha^2}$, and it suffices to privately estimate the empirical covariance within error α with sample complexity $\min\left\{\frac{d^{1.5}}{\varepsilon\alpha}, \frac{d}{\varepsilon\alpha^2}\right\}$. To do so, we estimate the empirical second moment matrix $M_2(X) = \frac{1}{n} \sum_{i=1}^n X_i X_i^T$, within Frobenius error $\frac{\alpha}{2}$, and the empirical mean $\mu(X) = \frac{1}{n} \sum X_i$ within ℓ_2 error $\frac{\alpha}{4}$, each under $\frac{\varepsilon}{2}$ -differential privacy. Let us denote the estimates \widetilde{M}_2 and $\widetilde{\mu}$, respectively. Let us also assume that $\widetilde{\mu} \in B_2^d$. To see that this is without loss of generality, recall that \mathcal{D} is supported on B_2^d , and, therefore $\|\mu\|_2 \leq 1$. This means that replacing $\widetilde{\mu}$ with its projection onto B_2^d , given by $\frac{1}{\max\{1, \|\widetilde{\mu}\|_2\}} \widetilde{\mu}$, does not increase the error $\|\widetilde{\mu} - \mu\|_2$. Finally, let us denote $\widetilde{\Sigma} = \widetilde{M}_2 - \widetilde{\mu}\widetilde{\mu}^T$. We have

$$\begin{aligned} \|\widetilde{\Sigma} - \Sigma\|_F &\leq \|\widetilde{M}_2 - M_2\|_F + \|\widetilde{\mu}\widetilde{\mu}^T - \mu\mu^T\|_F \\ &\leq \|\widetilde{M}_2 - M_2\|_F + \|\widetilde{\mu}(\widetilde{\mu}^T - \mu^T)\|_F + \|(\widetilde{\mu} - \mu)\mu^T\|_F \\ &= \|\widetilde{M}_2 - M_2\|_F + (\|\widetilde{\mu}\|_2 + \|\mu\|_2)\|\widetilde{\mu}^T - \mu^T\|_2 \\ &\leq \|\widetilde{M}_2 - M_2\|_F + 2\|\widetilde{\mu} - \mu\|_2, \end{aligned}$$

where in the last inequality we used that $\|\mu\|_2 \leq 1$ and $\|\widetilde{\mu}\|_2 \leq 1$. Therefore, as long as $\mathbb{E}\|\widetilde{M}_2 - M_2\|_F \leq \frac{\alpha}{2}$, and $\mathbb{E}\|\widetilde{\mu} - \mu\|_2 \leq \frac{\alpha}{4}$, our mechanism can achieve the desired error and privacy guarantees if it outputs $\widetilde{\Sigma}$.

It follows from Lemma 6.1 (with $r = 1$) that the estimate \widetilde{M}_2 of $M_2(X)$ can be computed under $\frac{\varepsilon}{2}$ -differential privacy with error $\mathbb{E}\|\widetilde{M}_2 - M_2\|_F \leq \frac{\alpha}{2}$ and with the required sample complexity $\min\left\{\frac{d^{1.5}}{\varepsilon\alpha}, \frac{d}{\varepsilon\alpha^2}\right\}$. We then notice that μ is a workload of d statistical queries on the universe, B_2^d , and its sensitivity polytope (i.e., body) is $K_\mu = B_2^d$. We use the mechanism in Lemma 4.2 to compute $\widetilde{\mu}$. The lemma guarantees that we achieve $\frac{\varepsilon}{2}$ -differential privacy and error $\mathbb{E}\|\widetilde{\mu} - \mu\|_2 \leq \frac{\alpha}{4}$ as long as $n \geq \frac{Cd}{\varepsilon\alpha}$ for a sufficiently large constant C . The mechanism is also efficient, as it only needs a separation oracle for B_2^d . \square

Finally, we observe that Lemma 6.1 implies an efficient algorithm for estimating 2-way marginals with optimal sample complexity. Recall that $Q_{2,d}$ is the workload of 2-way marginals on the universe $\{0, 1\}^d$, consisting of the $\binom{d}{2}$ queries $q_{\{i,j\}}(x) = x_i x_j$, $i, j \in [d], i \neq j$.

Theorem 6.2. *For any integer $d \geq 1$, there exists an ε -differentially private algorithm \mathcal{M} whose sample complexity for the workload of 2-way marginals is*

$$\text{sc}(\mathcal{M}, Q_{2,d}, \alpha) \lesssim \min\left\{\frac{d^{1.5}}{\alpha\varepsilon}, \frac{d}{\alpha^2\varepsilon}\right\}.$$

Moreover, the mechanism runs in time polynomial in $d, n, \frac{1}{\varepsilon}$.

Proof. Note that for a dataset $X = (x_1, \dots, x_n)$, and any distinct $i, j \in [d]$,

$$q_{\{i,j\}}(X) = \left(\frac{1}{n} \sum_{i=1}^n x_i x_i^T\right)_{i,j}.$$

Therefore, to approximate $Q_{2,d}$ it suffices to run the mechanism \mathcal{M} from Lemma 6.1, and the (normalized) error we get is

$$\mathbb{E} \left(\frac{1}{\binom{d}{2}} \sum_{1 \leq i < j \leq d} (q_{i,j}(X) - \mathcal{M}(X)_{i,j})^2 \right)^{1/2} \lesssim \mathbb{E} \frac{1}{d} \|M_2(X) - \mathcal{M}(X)\|_F,$$

where $M_2(X) = \frac{1}{n} \sum_{i=1}^n x_i x_i^T$. Since $\{0, 1\}^d \subseteq \sqrt{d}B_2^d$, Lemma 6.1 implies the required sample complexity of $\min\left\{\frac{d^{1.5}}{\alpha\varepsilon}, \frac{d}{\alpha^2\varepsilon}\right\}$, as well as the running time guarantee. \square

Note that the sample complexity in Theorem 6.2 matches the bound in (1.2) with a mechanism running in time polynomial in d . Moreover, as we will see in Section 7.2, this sample complexity bound is optimal among all ε -differentially private mechanisms. There is no differentially private mechanism known to achieve the optimal sample complexity for w -way marginals in time polynomial in d^w for any $w \geq 3$, and also no complexity theoretic evidence that this is impossible.

7. LOWER BOUNDS

In this section we show that the sample complexity bound in Theorem 1.1 is tight both for random query workloads, and for w -way marginals.

7.1. Worst Case Queries. The following lower bound shows that Theorem 1.1 is tight up to constants in the worst case. It is worth noting that the lower bound in Theorem 7.1 below holds with high probability for random ± 1 -valued queries, i.e., queries where for each $x \in \mathcal{X}$, $Q(x)$ is an independently drawn uniformly random vector in $\{-1, +1\}^k$.

Theorem 7.1. *There exists a constant $c > 0$ such that for any $\alpha \in (0, c)$, $\varepsilon \in (0, 1)$, and any k and N such that $k \leq \frac{1}{2}N$ and $N \geq \frac{2}{\alpha^2}$, there exists a workload Q of k bounded queries on a universe of size N for which*

$$\text{sc}_\varepsilon(Q, \alpha) \gtrsim \min \left\{ \frac{k}{\varepsilon\alpha}, \frac{\sqrt{k \log(N/k)}}{\varepsilon\alpha}, \frac{\log(\alpha^2 N)}{\varepsilon\alpha^2} \right\}. \quad (7.1)$$

Proof. The first two terms on the right hand side of (7.1) follow from results of Hardt and Talwar (HT10), and De (De12): see, in particular, Theorem 3.2 in (De12). To establish the third term, we slightly modify De's construction. By Claim 3.1 in (De12), and Lemma 2.2, it is enough to construct datasets $X_1, \dots, X_M \in \mathcal{X}^n$ such that

- $n = \frac{\log(M/2)}{\varepsilon} \gtrsim \frac{\log(\alpha^2 N)}{\varepsilon\alpha^2}$;
- any two distinct X_i, X_j differ in at least $\Delta \gtrsim n$ elements;
- there is an $a \gtrsim \alpha^2 n$ such that for any X_i and any $x \in \mathcal{X}$, x appears either a or 0 times in X_i .

Similarly to Claim A.1 in (De12), the existence of such X_1, \dots, X_M follows from classical lower bounds on the maximum achievable sizes of packings. We let $n = \frac{\log(M/2)}{\varepsilon}$, and define $s = \frac{n}{a}$ and $t = \frac{n-\Delta}{a}$. Then, for appropriate choices of $a \gtrsim \alpha^2 n$ and $\Delta \gtrsim n$, Proposition 2.1 in (EFF85) shows that there exists a family S_1, \dots, S_M of subsets of \mathcal{X} such that $|S_i| = s$ for all i , $|S_i \cap S_j| \leq t$ for all $i \neq j$, and

$$\log(M/2) \gtrsim t \log \left(\frac{Nt}{s^2} \right) \gtrsim \frac{\log(\alpha^2 N)}{\alpha^2}.$$

To finish the construction of the datasets, we define X_i to contain a copies of each element of S_i . It is then straightforward to verify that $X_1, \dots, X_M \in \mathcal{X}^n$ satisfy the required properties. \square

7.2. Marginals. Except for random queries, Theorem 1.1 also gives tight bounds for the important class of constant width marginal queries. Before stating our lower bound proving (1.2) is tight, we recall a useful standard lemma about packing points in Hamming distance. The lemma is a simple consequence of the classical Gilbert-Varshamov bound, but we give a short probabilistic proof for the convenience of the reader.

Lemma 7.2. *For any integers $d \geq 1$ and $M \geq 2$, and any set Σ of cardinality M , there exists a set $P \subseteq \Sigma^d$ such that, for any $x, y \in P$, $|\{i : x_i \neq y_i\}| \geq \frac{d}{4}$, and $\log |P| \gtrsim d \log M$.*

Proof. Before we construct P , let us bound the probability that two independent, uniformly random points X and Y in Σ^d satisfy $|\{i : X_i \neq Y_i\}| < \frac{d}{4}$. Note that the events $X_i = Y_i$ are independent for different i , and each of them happens with probability $\frac{1}{M}$. Therefore, the expected value of $|\{i : X_i = Y_i\}|$ is $\frac{d}{M}$, and, by a Chernoff bound,

$$\Pr \left[|\{i : X_i \neq Y_i\}| < \frac{d}{4} \right] < \Pr \left[|\{i : X_i = Y_i\}| \geq \frac{3d}{4} \right] \leq \left(\frac{e^{(3M/4)-1}}{(3M/4)^{3M/4}} \right)^{\frac{d}{M}}.$$

Let us pick a sequence of points $P_0 = (X_1, \dots, X_N)$ by independently and uniformly sampling N points from Σ^d , for an integer N to be chosen later. Let us call X_i and X_j a *bad pair* if $|\{\ell : X_{i,\ell} \neq X_{j,\ell}\}| < \frac{d}{4}$, where $X_{i,\ell}$ is the ℓ -th coordinate of X_i , and similarly for $X_{j,\ell}$. By the discussion above, the expected number of bad pairs is at most

$$\frac{N(N-1)}{2} \left(\frac{e^{(3M/4)-1}}{(3M/4)^{3M/4}} \right)^{\frac{d}{M}} \leq \frac{N^2}{2} \left(\frac{e^{(3/4)-(1/M)}}{(3M/4)^{3/4}} \right)^d.$$

Let us pick N to be the largest integer less than or equal to $\frac{1}{4} \left(\frac{(3M/4)^{3/4}}{e^{(3/4)-(1/M)}} \right)^d$, so that the expected number of bad pairs is at most $\frac{N}{8}$. Then, by Markov's inequality, with probability at least $\frac{1}{2}$, the number of bad pairs is at most $\frac{N}{4}$. We remove all X_i from P_0 involved in a bad pair, and define P to equal the set of the remaining points in P_0 . Clearly P satisfies the requirement of the lemma. Moreover, since we remove at most two points for each bad pair, and, with probability at least $\frac{1}{2}$ there are at most $\frac{N}{4}$ bad pairs, this means that, with probability at least $\frac{1}{2}$, P has size at least $\frac{N}{2}$. It is straightforward to verify that $\log \frac{N}{2} \gtrsim d \log M$ for any large enough d and for $M \geq 2$. For d smaller than a fixed constant, the lemma follows by a trivial construction, e.g., by taking $P = \{(\sigma, \dots, \sigma) : \sigma \in \Sigma\}$. \square

The next theorem uses a packing argument to show that the sample complexity bound (1.2) is tight. While the first term in the minimum on the right hand side dominates the second term, we keep it in the statement of the theorem to make the comparison with Theorem 1.1 more clear.

Theorem 7.3. *For any integer $\omega \geq 1$ there exists a constant $c > 0$ such that for any integer $d \geq 10\omega$ and for any $\alpha \in (0, c]$, we have*

$$\text{sc}_\varepsilon(Q_{\omega,d}, \alpha) \geq c \min \left\{ \frac{d^\omega}{\varepsilon \alpha}, \frac{d^{(\omega+1)/2}}{\varepsilon \alpha}, \frac{d}{\varepsilon \alpha^2} \right\}.$$

Proof. The first term in the lower bound only achieves the minimum when $\omega = 1$, so we only have to prove it holds in the case of 1-way marginals. In fact, this term is never smaller than the second term, but we handle the $\omega = 1$ case of the lower bound separately since it is simpler. This sample complexity lower bound for 1-way marginals shown by Hardt and Talwar (HT10) (see also (SU17)), and here we only briefly sketch the proof. Note that the set $S_{1,d} = \{Q_{1,d}(x) : x \in \{0,1\}^d\}$ is simply the boolean cube $\{0,1\}^d$. By Lemma 7.2, $\log P(\{0,1\}^d, \frac{\sqrt{d}}{2} B_2^d) \gtrsim d$. Then the lower bound on sample complexity follows from Lemma 2.3.

For the other two terms, we build packings of datasets and use Lemma 2.2. We start with the second term, which achieves the minimum when $\omega \geq 2$, and $\alpha \leq d^{-(\omega-1)/2}$. Let $h_0, h_1, \dots, h_{d-1} \in \{-1,1\}^d$ be columns of a Hadamard matrix (i.e., any two distinct h_i, h_j are orthogonal), where h_0 is the all-ones vector. Let us define, for $i \in [d-1]$, $g_i = \frac{1}{2}(h_0 + h_i) \in \{0,1\}^d$. Let, $y_1, \dots, y_M \in \{0,1\}^d$ be such that for all distinct y_i, y_j we have $\|y_i - y_j\|_2 \geq \frac{\sqrt{d}}{2}$. As noted above, by Lemma 7.2, we can take $\log M \gtrsim d$. Let us now define, for $i \in [M]$, and $s \in [d-1]^{\omega-1}$, $x_{i,s} \in \{0,1\}^{\omega d}$ to be the concatenation of y_i and $g_{s_1}, \dots, g_{s_{\omega-1}}$. For any $f : [d-1]^{\omega-1} \rightarrow [M]$, we define a dataset X_f that consists of $a \gtrsim \frac{d}{\varepsilon}$ copies of each of the elements $x_{f(s),s}$ for each sequence $s \in [d-1]^{\omega-1}$, together with $n - a(d-1)^{\omega-1}$ copies of the zero vector, where $n \geq a(d-1)^{\omega-1}$ will be chosen later.

We consider a family \mathcal{F} of functions from $[d-1]^{\omega-1}$ to $[M]$ such that for any two distinct $f, f' \in \mathcal{F}$ we have $|\{s \in [d-1]^{\omega-1} : f(s) \neq f'(s)\}| \geq \frac{1}{4}(d-1)^{\omega-1}$. Again by Lemma 7.2, we can find such a family with $\log |\mathcal{F}| \gtrsim (d-1)^{\omega-1} \log M$. The packing of datasets we consider consists of X_f for $f \in \mathcal{F}$.

An important observation we make is that if we project $Q_{\omega, \omega d}(x_{i,s})$ onto coordinates corresponding to queries q_S for which S has 1 element from the first d coordinates, one from the second d coordinates, etc., then we get the vector $y_i \otimes \bigotimes_{j=1}^{\omega-1} g_{s_j}$. We then have

$$\begin{aligned} \|Q_{\omega, \omega d}(X_f) - Q_{\omega, \omega d}(X_{f'})\|_2^2 &\geq \frac{a^2}{n^2} \left\| \sum_{s \in [d-1]^{\omega-1}} (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} g_{s_j} \right\|_2^2 \\ &= \frac{a^2}{2^{2\omega-2} n^2} \left\| \sum_{s \in [d-1]^{\omega-1}} (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} (h_0 + h_{s_j}) \right\|_2^2. \end{aligned} \quad (7.2)$$

Let Π be the orthogonal projection matrix onto the hyperplane orthogonal to h_0 (i.e., to the all-ones vector). Then multiplying $(y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} (h_0 + h_{s_j})$ by the orthogonal projection matrix $I \otimes \Pi^{\otimes \omega-1}$ gives

$$(y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} \Pi(h_0 + h_{s_j}) = (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} h_{s_j}.$$

Since orthogonal projection cannot increase the ℓ_2 norm, this implies

$$\begin{aligned} \left\| \sum_{s \in [d-1]^{\omega-1}} (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} (h_0 + h_{s_j}) \right\|_2^2 &\geq \left\| \sum_{s \in [d-1]^{\omega-1}} (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} h_{s_j} \right\|_2^2 \\ &= \sum_{s \in [d-1]^{\omega-1}} \left\| (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} h_{s_j} \right\|_2^2 \end{aligned} \quad (7.3)$$

$$= \sum_{s \in [d-1]^{\omega-1}} d^{\omega-1} \|y_{f(s)} - y_{f'(s)}\|_2^2. \quad (7.4)$$

Equality (7.3) holds since for any two distinct $s, s' \in [d-1]^{\omega-1}$

$$\left\langle (y_{f(s)} - y_{f'(s)}) \otimes \bigotimes_{j=1}^{\omega-1} h_{s_j}, (y_{f(s')} - y_{f'(s')}) \otimes \bigotimes_{j=1}^{\omega-1} h_{s'_j} \right\rangle = \langle y_{f(s)} - y_{f'(s)}, y_{f(s')} - y_{f'(s')} \rangle \prod_{j=1}^{\omega-1} \langle h_{s_j}, h_{s'_j} \rangle = 0,$$

with the final equality due to the pairwise orthogonality of h_1, \dots, h_{d-1} . Going back to (7.4), we recall that for any $f, f' \in \mathcal{F}$, the number of sequences s for which $f(s) \neq f'(s)$ is at least $\frac{1}{4}(d-1)^{\omega-1}$. Since any two distinct y_i, y_j satisfy $\|y_i - y_j\|_2^2 \geq \frac{d}{4}$, this gives us

$$\sum_{s \in [d-1]^{\omega-1}} d^{\omega-1} \|y_{f(s)} - y_{f'(s)}\|_2^2 \gtrsim d^{2\omega-1}. \quad (7.5)$$

Combining (7.2), (7.4), and (7.5), we get for any two distinct $f, f' \in \mathcal{F}$,

$$\frac{1}{\sqrt{\binom{\omega d}{\omega}}} \|Q_{\omega, \omega d}(X_f) - Q_{\omega, \omega d}(X_{f'})\|_2 \gtrsim \frac{cad^{(\omega-1)/2}}{n}.$$

where c is a constant that may depend on ω . Towards applying Lemma 2.2, let us choose $a \gtrsim \frac{d}{\varepsilon}$ so that $a(d-1)^{\omega-1} \leq \frac{\log(|\mathcal{F}|/2)}{\varepsilon}$. Then,

$$\frac{1}{\sqrt{\binom{\omega d}{\omega}}} \|Q_{\omega, \omega d}(X_f) - Q_{\omega, \omega d}(X_{f'})\|_2 \gtrsim \frac{cd^{(\omega+1)/2}}{\varepsilon n},$$

and, since any two distinct datasets $X_f, X_{f'}$ differ in at most $a(d-1)^{\omega-1}$ elements, we can apply Lemma 2.2 with $n = \frac{cd^{(\omega+1)/2}}{\alpha\varepsilon}$ to get the lower bound. Recall that we need $n \geq a(d-1)^{\omega-1}$, which happens when $\alpha \leq c'd^{-(\omega-1)/2}$ for a sufficiently small $c' > 0$. This is exactly the parameter regime in which we need the second term of the lower bound to hold. Note that we proved the lower bound against $Q_{\omega, \omega d}$, but this implies a lower bound for $Q_{\omega, d}$ for $d \geq \omega$.

The third term of the lower bound is proved analogously to the second, with some modifications. We choose some subset $S \subseteq [d-1]^{\omega-1}$ of size $|S| \gtrsim \frac{1}{\alpha^2}$ and only consider functions $f : S \rightarrow [M]$. When defining a dataset X_f , we take $a \gtrsim \frac{d}{\varepsilon}$ copies of each of the elements $x_{f(s),s}$ for each sequence $s \in S$, and we do not take any additional copies of the zero vector. The rest of the proof proceeds mutatis mutandis as the proof of the second term of the lower bound. \square

8. OPTIMALITY OF THE MECHANISMS

In the previous section we argued that the mechanisms in this paper are optimal for worst-case query workloads. We can in fact show that in the constant error regime the sample complexity of our mechanisms is also tight with respect to the optimal sample complexity for the given workload. The next lemma follows from combining Theorem 1.2 with the methods of Błasiok, Bun, Nikolov, and Steinke (BBNS19), and is the main ingredient in proving Theorem 1.4

Lemma 8.1. *For any workload Q of k queries over a universe of size N , and any $0 \leq \alpha \leq \frac{\text{diam}(K_Q)}{\sqrt{k}}$, there exists an ε -differentially private mechanism \mathcal{M} with running time polynomial in $k, N, n, \frac{1}{\varepsilon}$, and with sample complexity*

$$\text{sc}(\mathcal{M}, Q, \alpha) \lesssim \frac{1}{\varepsilon\alpha^2} \log \left(\frac{\text{diam}(S_Q)}{\alpha\sqrt{k}} \right)^2 \sup_{t \geq \alpha/4} t^2 \log P(S_Q, t\sqrt{k}B_2^k), \quad (8.1)$$

where $S_Q = \{Q(x) : x \in \mathcal{X}\}$.

In the proof of Lemma 8.1, we use the following lemma from (BBNS19).

Lemma 8.2. *Let $S \subseteq \mathbb{R}^k$ be a set such that for any distinct $y, y' \in S$ we have $\|y - y'\|_2 > \frac{\alpha\sqrt{k}}{2}$. Then*

$$w(S) \lesssim \sqrt{k} \log \left(\frac{\text{diam}(S)}{\alpha\sqrt{k}} \right) \sup \left\{ t \sqrt{\log P(S, t\sqrt{k}B_2^k)} : t \geq \frac{\alpha}{4} \right\}.$$

The mechanism in Lemma 8.1 is given by Algorithm 2.

Algorithm 2: The Coarse JL-Release mechanism

Let $S \subseteq S_Q = \{Q(x) : x \in \mathcal{X}\}$ be inclusion-maximal s.t. for all distinct $y, y' \in S$,
 $\|y - y'\|_2 > \frac{\alpha\sqrt{k}}{2}$;
 Define workload \tilde{Q} s.t. for each $x \in \mathcal{X}$, $\tilde{Q}(x) \in \arg \min\{\|y - Q(x)\|_2 : y \in S\}$.;
 Apply the mechanism from Theorem 1.2 to $\tilde{Q}(X)$, and output answer;

Proof of Lemma 8.1. As already mentioned, the required mechanism \mathcal{M} is given by Algorithm 2. The first key claim is that, for all $X \in \mathcal{X}^n$,

$$\frac{1}{\sqrt{k}} \|\tilde{Q}(X) - Q(X)\|_2 \leq \frac{\alpha}{2}. \quad (8.2)$$

By the triangle inequality, in order to verify (8.2), it suffices to check it for single-element datasets, i.e., to check that, for every $x \in \mathcal{X}$, $\frac{1}{\sqrt{k}} \|\tilde{Q}(x) - Q(x)\|_2 \leq \frac{\alpha}{2}$. This follows by maximality of S : if there were any $x \in \mathcal{X}$ such that $\frac{1}{\sqrt{k}} \|\tilde{Q}(x) - Q(x)\|_2 > \frac{\alpha}{2}$, then, by the construction of \tilde{Q} , it would follow that $\frac{1}{\sqrt{k}} \|y - Q(x)\|_2 > \frac{\alpha}{2}$ for all $y \in S$, and we could then add $Q(x)$ to S , contradicting its maximality. Therefore, no such x exists, and this proves (8.2).

By Theorem 1.2, we have $\mathbb{E} \frac{1}{\sqrt{k}} \|\mathcal{M}(X) - \tilde{Q}(X)\|_2 \leq \frac{\alpha}{2}$ as long as $n \geq \frac{Cw(\text{conv}(S))^2}{k\varepsilon\alpha^2} = \frac{Cw(S)^2}{k\varepsilon\alpha^2}$ for a constant $C > 0$. Together with Lemma 8.2 and (8.2), this means that

$$\text{sc}(\mathcal{M}, Q, \alpha) \lesssim \frac{w(S)^2}{k\alpha\varepsilon^2} \lesssim \frac{1}{\varepsilon\alpha^2} \log \left(\frac{\text{diam}(S)}{\alpha\sqrt{k}} \right)^2 \sup \left\{ t^2 \log P(S, t\sqrt{k}B_2^k) : t \geq \frac{\alpha}{4} \right\}.$$

This establishes the required sample complexity, since $P(S, t\sqrt{k}B_2^k) \leq P(S_Q, t\sqrt{k}B_2^k)$

The running time guarantee holds since S can be computed greedily by making a single pass over \mathcal{X} , and the rest of the algorithm is just the mechanism from Theorem 1.2. Privacy follows from Theorem 1.2. \square

Proof of Theorem 1.4. By Lemma 2.3, we have

$$\text{sc}_\varepsilon(Q, \alpha/8) \gtrsim \frac{1}{\varepsilon\alpha} \sup \left\{ t \log P(K_Q, t\sqrt{k}B_2^k) : t \geq \frac{\alpha}{4} \right\}. \quad (8.3)$$

Note that the suprema in (8.1) and (8.3) can both be taken over $t \in \left[\frac{\alpha}{4}, \frac{\text{diam}(S_Q)}{\sqrt{k}} \right)$. Therefore, we have

$$\frac{\text{sc}(\mathcal{M}, Q, \alpha)}{\text{sc}_\varepsilon(Q, \alpha/8)} \lesssim \frac{\text{diam}(S_Q)}{\alpha\sqrt{k}} \log \left(\frac{\text{diam}(S_Q)}{\alpha\sqrt{k}} \right)^2 = \frac{\text{diam}(K_Q)}{\alpha\sqrt{k}} \log \left(\frac{\text{diam}(K_Q)}{\alpha\sqrt{k}} \right)^2.$$

Together with the running time and privacy guarantees in Lemma 8.1, this proves the theorem. \square

ACKNOWLEDGMENT

The author acknowledges support from an NSERC Discovery Grant (RGPIN-2021-03206).

REFERENCES

- [ALPTJ10] Radosław Adamczak, Alexander E. Litvak, Alain Pajor, and Nicole Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *J. Amer. Math. Soc.*, 23(2):535–561, 2010. <https://doi.org/10.1090/S0894-0347-09-00650-X>.
- [BBDS12] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The johnson-lindenstrauss transform itself preserves differential privacy. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 410–419. IEEE Computer Society, 2012, pages 410–419. <https://doi.org/10.1109/FOCS.2012.67>.
- [BBNS19] Jarosław Błasiok, Mark Bun, Aleksandar Nikolov, and Thomas Steinke. Towards instance-optimal private query release. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2019*, pages 2480–2497. SIAM, Philadelphia, PA, 2019, pages 2480–2497. <https://doi.org/10.1137/1.9781611975482.152>.
- [BCD⁺07] Boaz Barak, Kamalika Chaudhuri, Cynthia Dwork, Satyen Kale, Frank McSherry, and Kunal Talwar. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In *Proceedings of the Twenty-Sixth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 11-13, 2007, Beijing, China*, pages 273–282. ACM, 2007, pages 273–282. <https://doi.org/10.1145/1265530.1265569>.
- [BDKT12] Aditya Bhaskara, Daniel Dadush, Ravishankar Krishnaswamy, and Kunal Talwar. Unconditional differentially private mechanisms for linear queries. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 1269–1284. ACM, 2012, pages 1269–1284. <https://doi.org/10.1145/2213977.2214089>.
- [BLR13] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *J. ACM*, 60(2):12:1–12:25, 2013. <https://doi.org/10.1145/2450142.2450148>.
- [BS16] Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 635–658, 2016, 9985:635–658. https://doi.org/10.1007/978-3-662-53641-4_24.
- [BUV14] Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 1–10. ACM, 2014, pages 1–10. <https://doi.org/10.1145/2591796.2591877>.
- [CKKL12] Mahdi Cheraghchi, Adam Klivans, Pravesh Kothari, and Homin K. Lee. Sub-modular functions are noise stable. In *SODA*, pages 1586–1592, 2012, pages 1586–1592. <https://doi.org/10.1137/1.9781611973099.126>.
- [CTUW14] Karthekeyan Chandrasekaran, Justin Thaler, Jonathan R. Ullman, and Andrew Wan. Faster private release of marginals on small databases. In *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014*, pages 387–402. ACM, 2014, pages 387–402. <https://doi.org/10.1145/>

2554797.2554833.

- [De12] Anindya De. Lower bounds in differential privacy. In *Proceedings of the 9th International Conference on Theory of Cryptography*, TCC'12, pages 321–338, Berlin, Heidelberg, 2012. Springer-Verlag, pages 321–338. https://doi.org/10.1007/978-3-642-28914-9_18.
- [DLY22] Wei Dong, Yuting Liang, and Ke Yi. Differentially private covariance revisited. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/057405fd73dd7ba7f32a7cb34fb7c7f5-Paper-Conference.pdf.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Conference on Theory of Cryptography*, TCC'06, pages 265–284, Berlin, Heidelberg, 2006. Springer-Verlag, pages 265–284. https://doi.org/10.1007/11681878_14.
- [DNT15] Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Efficient algorithms for privately releasing marginals via convex relaxations. *Discrete Comput. Geom.*, 53(3):650–673, 2015. <https://doi.org/10.1007/s00454-015-9678-x>.
- [DR14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014. <https://doi.org/10.1561/04000000042>.
- [DR16] Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016. URL: <http://arxiv.org/abs/1603.01887>.
- [DRS22] Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84(1):3–37, 2022. <https://doi.org/10.1111/rssb.12454>.
- [EFF85] P. Erdős, P. Frankl, and Z. Füredi. Families of finite sets in which no set is covered by the union of r others. *Israel J. Math.*, 51(1-2):79–89, 1985. <https://doi.org/10.1007/BF02772959>.
- [GLS88] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics: Study and Research Texts*. Springer-Verlag, Berlin, 1988. <https://doi.org/10.1007/978-3-642-97881-4>.
- [Har11] Moritz Hardt. *A study of privacy and fairness in sensitive data analysis*. PhD thesis, Princeton University, 2011. URL: <http://arks.princeton.edu/ark:/88435/dsp01vq27zn422> , .
- [HRS12] Moritz Hardt, Guy N. Rothblum, and Rocco A. Servedio. Private data release via learning thresholds. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 168–187. SIAM, 2012, pages 168–187. <https://doi.org/10.1137/1.9781611973099.15>.
- [HT10] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 705–714. ACM, 2010, pages 705–714. <https://doi.org/10.1145/1806689.1806786>.

- [Kas21] Shiva Prasad Kasiviswanathan. SGD with low-dimensional gradients with applications to private and distributed learning. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, pages 1905–1915. AUAI Press, 2021, 161:1905–1915. URL: <https://proceedings.mlr.press/v161/kasiviswanathan21a.html>.
- [KKMM13] Krishnaram Kenthapadi, Aleksandra Korolova, Ilya Mironov, and Nina Mishra. Privacy via the johnson-lindenstrauss transform. *J. Priv. Confidentiality*, 5(1), 2013. <https://doi.org/10.29012/jpc.v5i1.625>.
- [KL22] Bo’az Klartag and Joseph Lehec. Bourgain’s slicing problem and KLS isoperimetry up to polylog. *Geom. Funct. Anal.*, 32(5):1134–1159, 2022. <https://doi.org/10.1007/s00039-022-00612-9>.
- [LMPV17] Christopher Liaw, Abbas Mehrabian, Yaniv Plan, and Roman Vershynin. A simple tool for bounding the deviation of random matrices on geometric sets. In *Geometric aspects of functional analysis*, volume 2169 of *Lecture Notes in Math.*, pages 277–299. Springer, Cham, 2017. <https://doi.org/10.1007/978-3-319-45282-1>.
- [LT11] Michel Ledoux and Michel Talagrand. *Probability in Banach spaces*. Classics in Mathematics. Springer-Verlag, Berlin, 2011. Isoperimetry and processes, Reprint of the 1991 edition. <https://doi.org/10.1007/978-3-642-20212-4>.
- [Mir17] Ilya Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 263–275. IEEE Computer Society, 2017, pages 263–275. <https://doi.org/10.1109/CSF.2017.11>.
- [MMHM18] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. Optimizing error of high-dimensional statistical queries under differential privacy. *Proc. VLDB Endow.*, 11(10):1206–1219, 2018. <https://doi.org/10.14778/3231751.3231769>.
- [Nik15] Aleksandar Nikolov. An improved private mechanism for small databases. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 1010–1021. Springer, 2015, 9134:1010–1021. https://doi.org/10.1007/978-3-662-47672-7_82.
- [NTZ16] Aleksandar Nikolov, Kunal Talwar, and Li Zhang. The geometry of differential privacy: The small database and approximate cases. *SIAM J. Comput.*, 45(2):575–616, 2016. <https://doi.org/10.1137/130938943>.
- [NU21] Aleksandar Nikolov and Jonathan Ullman. Open problem - optimal query release for pure differential privacy. *DifferentialPrivacy.org*, 07 2021. <https://differentialprivacy.org/open-problem-optimal-query-release/>.
- [Pis89] Gilles Pisier. *The volume of convex bodies and Banach space geometry*, volume 94 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1989. URL: <https://doi-org.myaccess.library.utoronto.ca/10.1017/CB09780511662454>, <https://doi.org/10.1017/CB09780511662454>.
- [Rot11] Aaron Roth. The algorithmic foundations of differential privacy (lecture notes), Fall 2011.

- [Sta21] Nina Mesing Stausholm. Improved differentially private euclidean distance approximation. In *PODS'21: Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Virtual Event, China, June 20-25, 2021*, pages 42–56. ACM, 2021, pages 42–56. <https://doi.org/10.1145/3452021.3458328>.
- [SU17] Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *Journal of Privacy and Confidentiality*, 7(2), Jan. 2017. <https://doi.org/10.29012/jpc.v7i2.648>.
- [TUV12] Justin Thaler, Jonathan R. Ullman, and Salil P. Vadhan. Faster algorithms for privately releasing marginals. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*, volume 7391 of *Lecture Notes in Computer Science*, pages 810–821. Springer, 2012, 7391:810–821. https://doi.org/10.1007/978-3-642-31594-7_68.
- [Vad17] Salil Vadhan. *The Complexity of Differential Privacy*, pages 347–450. Springer International Publishing, Cham. pages 347–450, 2017. https://doi.org/10.1007/978-3-319-57048-8_7.
- [Ver18] Roman Vershynin. *High-dimensional probability*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018. An introduction with applications in data science, With a foreword by Sara van de Geer. URL: <https://doi-org.myaccess.library.utoronto.ca/10.1017/9781108231596>, <https://doi.org/10.1017/9781108231596>.