

Differentially Private Assouad, Fano, and Le Cam

Jayadev Acharya *
Cornell University

ACHARYA@CORNELL.EDU

Ziteng Sun *
Cornell University

ZS335@CORNELL.EDU

Huanyu Zhang *
Cornell University

HZ388@CORNELL.EDU

Editors: Vitaly Feldman, Katrina Ligett and Sivan Sabato

Abstract

Le Cam’s method, Fano’s inequality, and Assouad’s lemma are three widely used techniques to prove lower bounds for statistical estimation tasks. We propose their analogues under central differential privacy. Our results are simple, easy to apply and we use them to establish sample complexity bounds in several estimation tasks.

We establish the optimal sample complexity of discrete distribution estimation under total variation distance and ℓ_2 distance. We also provide lower bounds for several other distribution classes, including product distributions and Gaussian mixtures that are tight up to logarithmic factors. The technical component of our paper relates coupling between distributions to the sample complexity of estimation under differential privacy.

Keywords: Differential privacy, distribution estimation, information theory

1. Introduction

Statistical estimation tasks are often characterized by the optimal trade-off between the sample size and estimation error. There are two steps in establishing tight sample complexity bounds: An information theoretic lower bound on sample complexity and an algorithmic upper bound that achieves it. Several works have developed general tools to obtain the lower bounds (e.g., [Le Cam \(1973\)](#); [Assouad \(1983\)](#); [Ibragimov and Has’ Minskii \(2013\)](#); [Bickel and Ritov \(1988\)](#); [Devroye \(1987\)](#); [Han and Verdú \(1994\)](#); [Cover and Thomas \(2006\)](#); [Scarlett and Cevher \(2019\)](#), and references therein), and three prominent techniques are Le Cam’s method, Fano’s inequality, and Assouad’s lemma. An exposition of these three methods and their connections is presented in [Yu \(1997\)](#)¹.

In several estimation tasks, individual samples have sensitive information that must be protected. This is particularly of concern in applications such as healthcare, finance, geo-location, etc. Privacy-preserving computation has been studied in various fields including database, cryptography, statistics and machine learning ([Warner, 1965](#); [Dalenius, 1977](#); [Dinur and Nissim, 2003](#); [Wasserman and Zhou, 2010](#); [Duchi et al., 2013](#); [Chaudhuri et al., 2011](#)). Differential privacy (DP) ([Dwork et al., 2006](#)), which allows statistical inference while preserving the privacy of the individual samples, has become one of the most popular notions of privacy ([Dwork et al., 2006](#); [Wasserman and Zhou, 2010](#); [Dwork et al., 2010](#); [Blum et al., 2013](#); [McSherry and Talwar, 2007](#); [Dwork and Roth, 2014](#); [Kairouz et al.,](#)

* Research supported by NSF 1815893, NSF 1657471, NSF 1846300 (CAREER), and a Google Faculty Research Fellowship; equal contribution

1. The title of [Yu \(1997\)](#), “Assouad, Fano, and Le Cam” is the inspiration for our title.

2017). Differential privacy has been adopted by the US Census Bureau for the 2020 census and several large technology companies, including Google, Apple, and Microsoft (Erlingsson et al., 2014; Differential Privacy Team, Apple, 2017; Ding et al., 2017).

Differential privacy (Dwork et al., 2006). Let \mathcal{X} denote an underlying data domain of individual data samples and \mathcal{X}^n be the set of all possible length- n sequences over \mathcal{X} . For $x, y \in \mathcal{X}^n$, $d_{\text{Ham}}(x, y)$ is their Hamming distance, the number of coordinates they differ at. A (randomized) estimator $\hat{\theta} : \mathcal{X}^n \rightarrow \Theta$ is (ε, δ) -differentially private (denoted as (ε, δ) -DP) if for any $S \subseteq \Theta$, and all $x, y \in \mathcal{X}^n$ with $d_{\text{Ham}}(x, y) \leq 1$, the following holds:

$$\Pr(\hat{\theta}(x) \in S) \leq e^\varepsilon \cdot \Pr(\hat{\theta}(y) \in S) + \delta. \quad (1)$$

The case $\delta = 0$ is *pure differential privacy* and is denoted by ε -DP. We consider problems of parameter estimation and goodness-of-fit (hypothesis testing) under differential privacy constraints.

Setting. Let \mathcal{P} be any collection of distributions over \mathcal{X}^n , where n denotes the number of samples.² Let $\theta : \mathcal{P} \rightarrow \Theta$ be a parameter of the distribution that we want to estimate. Let $\ell : \Theta \times \Theta \rightarrow \mathbb{R}_+$ be a pseudo-metric which is the loss function for estimating θ . We now describe the minimax framework of parameter estimation, and hypothesis testing.

Parameter estimation. The risk of an estimator $\hat{\theta} : \mathcal{X}^n \rightarrow \Theta$ under loss ℓ is $\max_{p \in \mathcal{P}} \mathbb{E}_{X \sim p} [\ell(\hat{\theta}(X), \theta(p))]$, the worst case expected loss of $\hat{\theta}$ over \mathcal{P} . Note that $X \in \mathcal{X}^n$, since p is a distribution over \mathcal{X}^n . The minimax risk of estimation under ℓ for the class \mathcal{P} is

$$R(\mathcal{P}, \ell) := \min_{\hat{\theta}} \max_{p \in \mathcal{P}} \mathbb{E}_{X \sim p} [\ell(\hat{\theta}(X), \theta(p))].$$

The minimax risk under differentially private protocols is given by restricting $\hat{\theta}$ to be differentially private. For (ε, δ) -DP, we study the following minimax risk:

$$R(\mathcal{P}, \ell, \varepsilon, \delta) := \min_{\hat{\theta} \text{ is } (\varepsilon, \delta)\text{-DP}} \max_{p \in \mathcal{P}} \mathbb{E}_{X \sim p} [\ell(\hat{\theta}(X), \theta(p))]. \quad (2)$$

For $\delta = 0$, the above minimax risk under ε -DP is denoted as $R(\mathcal{P}, \ell, \varepsilon)$.

Hypothesis testing. Hypothesis testing can be cast in the framework of parameter estimation as follows. Let $\mathcal{P}_1 \subset \mathcal{P}$, and $\mathcal{P}_2 \subset \mathcal{P}$ be two disjoint subsets of distributions denoting the two hypothesis classes. Let $\Theta = \{1, 2\}$, such that for $p \in \mathcal{P}_i$, let $\theta(p) = i$. For a test $\hat{\theta} : \mathcal{X}^n \rightarrow \{1, 2\}$, and $\ell(\theta, \theta') = \mathbb{I}\{\theta \neq \theta'\} = |\theta - \theta'|$, the error probability is the worst case risk under this loss function:

$$P_e(\hat{\theta}, \mathcal{P}_1, \mathcal{P}_2) := \max_i \max_{p \in \mathcal{P}_i} \Pr_{X \sim p}(\hat{\theta}(X) \neq i) = \max_i \max_{p \in \mathcal{P}_i} \mathbb{E}_{X \sim p} [|\hat{\theta}(X) - \theta(p)|]. \quad (3)$$

Organization. The remainder of the paper is organized as follows. In Section 2.1, 2.2, and 2.3 we state the privatized versions of Le Cam, Fano, and Assoud’s method respectively. In Section 2.4 we discuss the applications of these results to several estimation tasks. In Section 2.5 we discuss some related and prior work. In Section 3 and 4 we prove the bounds for distribution estimation under pure DP and approximate DP respectively. In Section 5, we prove the main theorems stated in Section 2.

2. In the general setting, we are not assuming i.i.d. distribution over \mathcal{X}^n , although we will specialize to this case later.

2. Results

Le Cam’s method is used to establish lower bounds for hypothesis testing and functional estimation. Fano’s inequality, and Assouad’s lemma prove lower bounds for multiple hypothesis testing problems and can be applied to parameter estimation tasks such as estimating distributions. We develop extensions of these results with differential privacy.

An observation. A coupling between distributions p_1 and p_2 over \mathcal{X}^n is a joint distribution (X, Y) over $\mathcal{X}^n \times \mathcal{X}^n$ whose marginals satisfy $X \sim p_1$ and $Y \sim p_2$ ³. Our lower bounds are based on the following observation. If there is a coupling (X, Y) between distributions p_1 and p_2 over \mathcal{X}^n with $\mathbb{E}[d_{\text{Ham}}(X, Y)] = D$, then a draw from p_1 can be converted to a draw from p_2 by changing D coordinates in expectation. By the group property of differential privacy, roughly speaking, for any (ε, δ) -DP estimator $\hat{\theta}$, it must satisfy

$$\forall S \subseteq \Theta, \Pr_{X \sim p}(\hat{\theta}(X) \in S) \leq e^{D\varepsilon} \cdot \Pr_{X \sim q}(\hat{\theta}(Y) \in S) + \delta D e^{\varepsilon(D-1)}.$$

Hence, if there exists an algorithm that distinguishes between p_1 and p_2 reliably, D must be large ($\Omega(1/\varepsilon + \delta)$).

2.1. DP Le Cam’s method

Le Cam’s method (Lemma 1 of Yu (1997)) is widely used to prove lower bounds for composite hypothesis testings such as uniformity testing (Paninski, 2008), density estimation (Yu, 1997; Ray and Schmidt-Hieber, 2018), and estimating functionals of distributions (Jiao et al., 2015; Wu and Yang, 2016; Polyanskiy and Wu, 2019).

We use the expected Hamming distance between couplings of distributions in the two classes to obtain the following extension of Le Cam’s method with (ε, δ) -DP, which is an adaptation of a similar result in Acharya et al. (2018b). For the hypothesis testing problem described above, let $\text{co}(\mathcal{P}_i)$ be the convex hull of distributions in \mathcal{P}_i , which are also families of distributions over \mathcal{X}^n .

Theorem 1 ((ε, δ)-DP Le Cam’s method) *Let $p_1 \in \text{co}(\mathcal{P}_1)$ and $p_2 \in \text{co}(\mathcal{P}_2)$. Let (X, Y) be a coupling between p_1 and p_2 with $D = \mathbb{E}[d_{\text{Ham}}(X, Y)]$. Then for $\varepsilon \geq 0, \delta \geq 0$, any (ε, δ) -differentially private hypothesis testing algorithm $\hat{\theta}$ must satisfy*

$$P_e(\hat{\theta}, \mathcal{P}_1, \mathcal{P}_2) \geq \frac{1}{2} \max \left\{ 1 - d_{TV}(p_1, p_2), 0.9e^{-10\varepsilon D} - 10D\delta \right\}, \quad (4)$$

where $d_{TV}(p_1, p_2) := \sup_{A \subseteq \mathcal{X}^n} (p_1(A) - p_2(A)) = \frac{1}{2} \|p_1 - p_2\|_1$ is the total variation (TV) distance of p_1 and p_2 .

The first term here is the original Le Cam’s result (Le Cam, 1973, 1986; Yu, 1997; Canonne, 2015) and the second term is a lower bound on the additional error due to privacy. Note that the second term increases when D decreases. Choosing p_1, p_2 with small D makes the RHS of (4) large, hence giving better testing lower bounds. A similar result (Theorem 1 in Acharya et al. (2018b)), along with a suitable coupling was used in Acharya et al. (2018b) to obtain the optimal sample complexity of testing discrete distributions. We defer the proof of this theorem to Section 5.1.

3. We use the term coupling to refer to both the random variable (X, Y) and the joint distribution.

2.2. DP Fano’s inequality

Theorem 1 (DP Le Cam’s method) characterizes lower bounds for binary hypothesis testing. In estimation problems with multiple parameters, it is common to reduce the problem to a multi-way hypothesis testing problem. The following theorem, proved in Section 5.2, provides a lower bound on the risk of multi-way hypothesis testing under ε -DP. Let

$$D_{KL}(p_i, p_j) := \sum_{x \in \mathcal{X}^n} p_i(x) \log \frac{p_i(x)}{p_j(x)}$$

be the KL divergence between (discrete) distributions p_i and p_j .⁴

Theorem 2 (ε -DP Fano’s inequality) *Let $\mathcal{V} = \{p_1, p_2, \dots, p_M\} \subseteq \mathcal{P}$ such that for all $i \neq j$,*

- (a) $\ell(\theta(p_i), \theta(p_j)) \geq \alpha$,
- (b) $D_{KL}(p_i, p_j) \leq \beta$,
- (c) *there exists a coupling (X, Y) between p_i and p_j such that $\mathbb{E}[d_{Ham}(X, Y)] \leq D$, then*

$$R(\mathcal{P}, \ell, \varepsilon) \geq \max \left\{ \frac{\alpha}{2} \left(1 - \frac{\beta + \log 2}{\log M} \right), 0.4\alpha \min \left\{ 1, \frac{M}{e^{10\varepsilon D}} \right\} \right\}. \quad (5)$$

Under pure DP constraints, Theorem 2 extends Theorem 1 to the multiple hypothesis case. Non-private Fano’s inequality (e.g., Lemma 3 of Yu (1997)) requires only conditions (a) and (b) and provides the first term of the risk bound above. Now, if we consider the second term, which is the additional cost due to privacy, we would require $\exp(10\varepsilon D) \geq M$, i.e., $D \geq \log M / (10\varepsilon)$ to achieve a risk less than 0.4α . Therefore, for reliable estimation, the expected Hamming distance between any pair of distributions cannot be too small. In Corollary 4, we provide a corollary of this result to establish sample complexity lower bounds for several distribution estimation tasks.

Theorem 2 (ε -DP Fano’s inequality) can also be seen as a probabilistic generalization of the classic packing lower bound (Vadhan, 2017). The packing argument, with its roots in database theory, considers inputs to be deterministic datasets whose pairwise Hamming distances are bounded with probability one, while Theorem 2 considers randomly generated datasets whose Hamming distances are bounded in expectation. This difference makes Theorem 2 better suited for proving lower bounds for statistical estimation problems. We discuss this difference in details in Section 2.5.

Remark. Theorem 2 is a bound on the risk for pure differential privacy ($\delta = 0$). Our proof extends to (ε, δ) -DP for $\delta = O\left(\frac{1}{M}\right)$, which is not sufficient to establish meaningful bounds since in most problems M will be chosen to be exponential in the problem parameters. To circumvent this difficulty, in the next section we provide a private analogue of Assouad’s method, which also works for (ε, δ) -DP.

2.3. DP Assouad’s method

Our next result is a private version of Assouad’s lemma (Lemma 2 of Yu (1997), and Assouad (1983)). Recall that \mathcal{P} is a set of distributions over \mathcal{X}^n . Let $\mathcal{V} \subseteq \mathcal{P}$ be a set of distributions indexed by the

4. For continuous distributions, the summation is replaced with an integral and the probability mass functions are replaced with densities. We focus on discrete distributions in the proof while the results hold for continuous distributions as well.

hypercube $\mathcal{E}_k := \{\pm 1\}^k$, and the loss ℓ is such that

$$\forall u, v \in \mathcal{E}_k, \ell(\theta(p_u), \theta(p_v)) \geq 2\tau \cdot \sum_{i=1}^k \mathbb{I}(u_i \neq v_i). \quad (6)$$

Assouad's method provides a lower bound on the estimation risk for distributions in \mathcal{V} , which is a lower bound for \mathcal{P} . For each coordinate $i \in [k]$, consider the following mixture distributions obtained by averaging over all distributions with a fixed value at the i th coordinate,

$$p_{+i} = \frac{2}{|\mathcal{E}_k|} \sum_{e \in \mathcal{E}_k: e_i = +1} p_e, \quad p_{-i} = \frac{2}{|\mathcal{E}_k|} \sum_{e \in \mathcal{E}_k: e_i = -1} p_e.$$

Assouad's lemma provides a lower bound on the risk by using (6) and considering the problem of distinguishing p_{+i} and p_{-i} . Analogously, we prove the following privatized version of Assouad's lemma by considering the minimax risk of a private hypothesis testing $\phi : \mathcal{X}^n \rightarrow \{-1, +1\}$ between p_{+i} and p_{-i} . The detailed proof is in Section 5.4.

Theorem 3 (DP Assouad's method) $\forall i \in [k]$, let $\phi_i : \mathcal{X}^n \rightarrow \{-1, +1\}$ be a binary classifier.

$$R(\mathcal{P}, \ell, \varepsilon, \delta) \geq \frac{\tau}{2} \cdot \sum_{i=1}^k \min_{\phi_i \text{ is } (\varepsilon, \delta)\text{-DP}} (\Pr_{X \sim p_{+i}}(\phi_i(X) \neq 1) + \Pr_{X \sim p_{-i}}(\phi_i(X) \neq -1)).$$

Moreover, if $\forall i \in [k]$, there exists a coupling (X, Y) between p_{+i} and p_{-i} with $\mathbb{E}[d_{\text{Ham}}(X, Y)] \leq D$,

$$R(\mathcal{P}, \ell, \varepsilon, \delta) \geq \frac{k\tau}{2} \cdot (0.9e^{-10\varepsilon D} - 10D\delta). \quad (7)$$

The first bound is the classic Assouad's Lemma and (7) is the loss due to privacy constraints. Once again note that (7) grows with decreasing D . Compared to Theorem 2 (DP Fano's inequality), Theorem 3 works under (ε, δ) -DP, which is a less stringent privacy notion.

2.4. Applications

We now describe several applications of the theorems above.

Applications of Theorem 1. Acharya et al. (2018b) developed a result similar to Theorem 1, which is used to establish sample complexity lower bounds for differentially private uniformity testing under total variation distance (Acharya et al., 2018b; Aliakbarpour et al., 2018), and for differentially private entropy and support size estimation (Acharya et al., 2018a). In this paper, we use Theorem 1 as a stepping stone to prove private Assouad's method (Theorem 3).

Distribution Estimation and Applications of Theorem 2 and 3. We will apply Theorem 2 (private Fano's inequality) and Theorem 3 (private Assouad's lemma) to some classic distribution estimation problems. The results are summarized in Table 1 and Table 2. Before presenting the results, we describe the framework of minimax distribution estimation.

Distribution estimation framework. Let \mathcal{Q} be a collection of distributions over \mathcal{X} , and for this \mathcal{Q} , let $\mathcal{P} = \mathcal{Q}^n := \{q^n : q \in \mathcal{Q}\}$ be the collection of n -fold distributions over \mathcal{X}^n induced by i.i.d. draws from a distribution over \mathcal{Q} . The parameter space is $\Theta = \mathcal{Q}$, where $\theta(q^n) = q$, and ℓ is a distance

measure between distributions in \mathcal{Q} . Let $\alpha > 0$ be a fixed parameter. The sample complexity, $S(\mathcal{Q}, \ell, \alpha, \varepsilon, \delta)$ is the smallest number of samples n to make $R(\mathcal{Q}^n, \ell, \varepsilon, \delta) \leq \alpha$, i.e.,

$$S(\mathcal{Q}, \ell, \alpha, \varepsilon, \delta) = \min\{n : R(\mathcal{Q}^n, \ell, \varepsilon, \delta) \leq \alpha\}.$$

When $\delta = 0$, we denote the sample complexity by $S(\mathcal{Q}, \ell, \alpha, \varepsilon)$. We will state our results in terms of sample complexity. The following corollary of Theorem 2 can be used to prove lower bounds on the sample complexity in this distribution estimation framework. The detailed proof of the corollary is in Section 5.3.

Corollary 4 (ε -DP distribution estimation) *Given $\varepsilon > 0$, let $\mathcal{V} = \{q_1, q_2, \dots, q_M\} \subseteq \mathcal{Q}$ be a set of distributions over \mathcal{X} with size M , such that for all $i \neq j$,*

- (a) $\ell(q_i, q_j) \geq 3\tau$,
- (b) $D_{KL}(q_i, q_j) \leq \beta$,
- (c) $d_{TV}(q_i, q_j) \leq \gamma$,

then

$$S(\mathcal{Q}, \ell, \tau, \varepsilon) = \Omega\left(\frac{\log M}{\beta} + \frac{\log M}{\gamma\varepsilon}\right).$$

Remark. With only conditions (a) and (b), we obtain the first term of the sample complexity lower bound which is the original Fano's bound for sample complexity. By Pinsker's inequality, a bound on the KL divergence (Condition (c)) would imply a bound on TV distance (Condition (b)), i.e., $\gamma \leq \sqrt{\beta/2}$. Hence Conditions (a) and (b) can also imply a lower bound on the sample complexity. We include all three conditions here since it is possible that in certain applications $\gamma \ll \sqrt{\beta/2}$, and hence a better bound can be obtained.

Problem	Upper Bounds	Lower Bounds
k-ary	$\Theta\left(\frac{k}{\alpha^2} + \frac{k}{\alpha\varepsilon}\right)$ (Diakonikolas et al. (2015), Theorem 8)	
k-ary, ℓ_2 distance	$O\left(\frac{1}{\alpha^2} + \min\left(\frac{\sqrt{k}}{\alpha\varepsilon}, \frac{\log k}{\alpha^2\varepsilon}\right)\right)$ (Theorem 9)	$\Omega\left(\frac{1}{\alpha^2} + \min\left(\frac{\sqrt{k}}{\alpha\varepsilon}, \frac{\log(k\alpha^2)}{\alpha^2\varepsilon}\right)\right)$ (Theorem 9)
product distribution	$O\left(kd \log\left(\frac{kd}{\alpha}\right)\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)\right)$ (Bun et al., 2019)	$\Omega\left(kd\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)\right)$ (Theorem 11)
Gaussian mixtures	$O\left(kd \log\left(\frac{dR}{\alpha}\right)\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)\right)$ (Bun et al., 2019)	$\Omega\left(kd\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)\right)$ (Theorem 12)

Table 1: Summary of the sample complexity bounds for ε -DP discrete distribution estimation. Unless mentioned, the bounds are all for estimation under total variation distance.

We now present examples of distribution classes we consider.

k -ary discrete distribution estimation. Suppose $\mathcal{X} = [k] := \{1, \dots, k\}$, and $\mathcal{Q} := \Delta_k$ is the simplex of k -ary distributions over $[k]$. We consider estimation in both total variation and ℓ_2 distance.

Problem	Upper Bounds	Lower Bounds
k -ary	$O\left(\frac{k}{\alpha^2} + \frac{k}{\alpha\varepsilon}\right)$ (Diakonikolas et al. (2015), Theorem 13)	$\Omega\left(\frac{k}{\alpha^2} + \frac{k}{\alpha(\varepsilon+\delta)}\right)$ (Theorem 13)
k -ary, ℓ_2 distance	$O\left(\frac{1}{\alpha^2} + \min\left(\frac{\sqrt{k}}{\alpha\varepsilon}, \frac{\log k}{\alpha^2\varepsilon}\right)\right)$ (Theorem 14)	$\Omega\left(\frac{1}{\alpha^2} + \min\left(\frac{\sqrt{k}}{\alpha(\varepsilon+\delta)}, \frac{1}{\alpha^2(\varepsilon+\delta)}\right)\right)$ (Theorem 14)
product distribution ($k = 2$)	$O\left(d \log\left(\frac{d}{\alpha}\right)\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)\right)$ (Kamath et al., 2019; Bun et al., 2019)	$\Omega\left(\frac{d}{\alpha^2} + \frac{d}{\alpha(\varepsilon+\delta)}\right)$ (Theorem 16, Kamath et al. (2019))

Table 2: Summary of the sample complexity bounds for (ε, δ) -DP discrete distribution estimation. Unless mentioned, the bounds are all for estimation under total variation distance.

(k, d) -product distributions. Consider $\mathcal{X} = [k]^d$, and let $\mathcal{Q} := \Delta_k^d$ be the set of product distributions over $[k]^d$, where the marginal distribution on each coordinate is over $[k]$ and independent of the other coordinates. We study estimation under total variation distance. A special case of this is Bernoulli product distributions ($k = 2$), where each of the d coordinates is an independent Bernoulli random variable.

d -dimensional Gaussian mixtures. Suppose $\mathcal{X} = \mathbb{R}^d$, and $\mathcal{G}_d := \{\mathcal{N}(\mu, I_d) : \|\mu\|_2 \leq R\}$ is the set of all Gaussian distributions in \mathbb{R}^d with bounded mean and identity covariance matrix. The bounded mean assumption is unavoidable, since by Bun et al. (2019), it is not possible to learn a single Gaussian distribution under pure DP without this assumption. We consider

$$\mathcal{Q} = \mathcal{G}_{k,d} := \left\{ \sum_{j=1}^k w_j p_j : \forall j \in [k], w_j \geq 0, p_j \in \mathcal{G}_d, w_1 + \dots + w_k = 1 \right\},$$

the collection of mixtures of k distributions from \mathcal{G}_d .

Applications of Theorem 2. We apply Corollary 4 and obtain sample complexity lower bounds for the tasks mentioned above under pure differential privacy.

k -ary distribution estimation. Without privacy constraints, the sample complexity of k -ary discrete distributions under total variation, and ℓ_2 distance is $\Theta(k/\alpha^2)$ and $\Theta(1/\alpha^2)$ respectively, achieved by the empirical estimator. Under ε -DP constraint, an upper bound of $O(k/\alpha^2 + k/\alpha\varepsilon)$ samples for total variation distance is known using Laplace mechanism (Dwork et al., 2006) (e.g. Diakonikolas et al. (2015)). In Theorem 8, we establish the sample complexity of this problem by providing a lower bound that matches this upper bound. The bound shows that when $\varepsilon \ll \alpha$, the cost due to privacy dominates the statistical error and when $\varepsilon \geq \alpha$, the privacy cost is almost negligible. The same break point (up to logarithmic factors) has also been observed for product distributions and mixtures of Gaussian distributions, as listed below.

Under ℓ_2 distance, in Theorem 9 we design estimators and establish their optimality whenever $\alpha < k^{-1/2}$ or $\alpha \geq k^{-0.499}$, which contains almost all the parameter range. Note that under ℓ_2 distance, estimation without privacy has sample complexity independent of k , whereas an unavoidable

logarithmic dependence on k is introduced due to privacy requirements. The results are presented in Section 3.1.

(k, d)-product distribution estimation. For (k, d) -product distribution estimation under ε -DP, [Bun et al. \(2019\)](#) proposed an algorithm that uses $O(kd \log(kd/\alpha)(1/\alpha^2 + 1/\alpha\varepsilon))$ samples. In this paper, we present a lower bound of $\Omega(kd/\alpha^2 + kd/\alpha\varepsilon)$, which matches their upper bound up to logarithmic factors. For Bernoulli product distributions, [Kamath et al. \(2019\)](#) proved a lower bound of $\Omega(d/\alpha\varepsilon)$ under $(\varepsilon, 3/64n)$ -DP, which is naturally a lower bound for pure DP. The details are presented in Section 3.2.

Estimating Gaussian mixtures. [Bun et al. \(2019\)](#) provided an upper bound of $\tilde{O}(kd/\alpha^2 + kd/\alpha\varepsilon)$ samples. Without privacy, a tight bounds of $\Omega(kd/\alpha^2)$ was shown in [Suresh et al. \(2014\)](#); [Daskalakis and Kamath \(2014\)](#); [Ashtiani et al. \(2018\)](#). In this paper, we prove a lower bound of $\Omega(kd/\alpha^2 + kd/\alpha\varepsilon)$, which matches the upper bound up to logarithmic factors. For the special case of estimating a single Gaussian ($k = 1$), a lower bound of $n = \Omega(d/(\alpha\varepsilon \log d))$ was given in [Kamath et al. \(2019\)](#) for $(\varepsilon, 3/64n)$ -DP, which implies a lower bound that is $\log d$ factor weaker than our result under pure DP.

Applications of Theorem 3. As remarked earlier, Theorem 2 only works for pure DP (or approximate DP with very small δ). Assouad’s lemma can be used to obtain lower bounds for distribution estimation under (ε, δ) -DP. For k -ary distribution estimation under TV distance, we get a lower bound of $\Omega(k/\alpha^2 + k/\alpha(\varepsilon + \delta))$. This shows that even up to $\delta = O(\varepsilon)$, the sample complexity for (ε, δ) -DP is the same as that under ε -DP.

For Bernoulli ($k = 2$) product distributions, [Kamath et al. \(2019\)](#) provides an efficient (ε, δ) -DP algorithm that achieves an upper bound of $O(d \log(d/\alpha)(1/\alpha^2 + 1/\alpha\varepsilon))$.⁵ The lower bound $\Omega(d/\alpha^2 + d/\alpha\varepsilon)$ obtained in [Kamath et al. \(2019\)](#) by fingerprinting holds for small values of $\delta = O(1/n)$. Note by the definition of DP (1), if $\delta > 1/n$, a DP algorithm can blatantly disregard the privacy of δn users. Therefore in most of the literature, δ is assumed to be $O(1/n)$. We want to make a complimentary remark that we can obtain the same lower bound all the way up to $\delta = O(\varepsilon)$. This shows that there is no gain even if we compromise the privacy of a δ fraction of users. Therefore, there is no incentive to do it. We describe the details about these applications in Section 4.

2.5. Related and prior work

2.5.1. PRIVATE DISTRIBUTION ESTIMATION

Protecting privacy generally comes at the cost of performance degradation. Previous literature has studied various problems and established utility privacy trade-off bounds, including distribution estimation, hypothesis testing, property estimation, empirical risk minimization, etc ([Chaudhuri et al., 2011](#); [Lei, 2011](#); [Bassily et al., 2014](#); [Diakonikolas et al., 2015](#); [Cai et al., 2017](#); [Acharya et al., 2018b](#); [Kamath et al., 2019](#); [Aliakbarpour et al., 2018](#); [Acharya et al., 2018a](#)).

There has been significant recent interest in differentially private distribution estimation. [Diakonikolas et al. \(2015\)](#) gives upper bounds for privately learning k -ary distributions under total variation distance. [Kamath et al. \(2019\)](#); [Bun et al. \(2019\)](#); [Karwa and Vadhan \(2018\)](#) focus on high-dimensional distributions, including product distributions and Gaussian distributions. As discussed in the previous section, our proposed lower bounds improve upon their lower bounds in various settings.

5. The algorithm in [Bun et al. \(2019\)](#) works for ε -DP and general k but it is not computationally efficient.

Bun et al. (2015) studies the problem of privately estimating a distribution in Kolmogorov distance, which is weaker than total variation distance. Upper and lower bounds for private estimation of the mean of product distributions in ℓ_∞ distance, heavy tailed distributions, and Markov Random fields are studied in Blum et al. (2005); Dwork et al. (2006); Steinke and Ullman (2017a); Bun et al. (2018); Kamath et al. (2020); Zhang et al. (2020).

Several estimation tasks including distribution estimation and hypothesis testing have also been considered under the distributed notion of local differential privacy, e.g., Warner (1965); Kasiviswanathan et al. (2011); Erlingsson et al. (2014); Duchi et al. (2013); Kairouz et al. (2016); Wang et al. (2019); Sheffet (2017); Ye and Barg (2018); Gaboardi and Rogers (2018); Acharya et al. (2019c); Acharya and Sun (2019); Acharya et al. (2019a).

2.5.2. LOWER BOUNDS IN DIFFERENTIAL PRIVACY

Several methods have been proposed in the literature to prove lower bounds under DP constraints. These include packing argument (Hardt and Talwar, 2010; Vadhan, 2017; Beimel et al., 2010), fingerprinting (Bun et al., 2015; Steinke and Ullman, 2017a, 2015; Bun et al., 2017, 2018; Kamath et al., 2019) and coupling based arguments (Acharya et al., 2018b; Karwa and Vadhan, 2018).

Binary Testing and Coupling. Coupling based arguments have been recently used to prove lower bounds for binary hypothesis testing, including the independent works of Acharya et al. (2018b); Karwa and Vadhan (2018). Acharya et al. (2018b) establishes a very similar result to Theorem 1 and uses it to obtain lower bounds for a composite hypothesis testing problem on discrete distributions. Karwa and Vadhan (2018) proves a similar result for simple hypothesis testing and uses it to lower bound the sample complexity of estimating the mean of a one-dimensional Gaussian distribution. For both papers, the coupling argument implies that it is hard to differentially privately distinguish between two distributions, supposing there exists a coupling with small expected Hamming distance. This method can be viewed as another form of private Le Cam’s method (Theorem 1) and it can only be applied where binary hypothesis testing is involved. Barber and Duchi (2014) also uses a private version of Le Cam’s method to prove lower bounds for differentially private mean estimation. However, instead of the expected Hamming distance between any couplings, their method only depends on the TV distance between the distributions, which corresponds to the naive independent coupling. Canonne et al. (2019) uses coupling bounds in Acharya et al. (2018b) to derive instance-optimal bounds for simple binary hypothesis testing under pure DP. They consider a coupling only for symbols whose likelihood ratio between the two hypothesis distributions is large, which results in better bounds for certain instances. The argument only considers pure DP and the case where samples are i.i.d generated while Theorem 1 and Acharya et al. (2018b) can handle approximate DP and arbitrary distributions (e.g. mixtures of i.i.d distributions considered in this paper).

Pure DP Estimation and Packing. Packing argument (Hardt and Talwar, 2010; Beimel et al., 2010; Vadhan, 2017) is a geometric approach to prove lower bounds for estimation under pure DP. We state a form of the packing bound below:

Lemma 5 (Packing lower bound (Vadhan, 2017)) *Let $\mathcal{V} = \{x_1, x_2, \dots, x_M\}$ be a set of M datasets over \mathcal{X}^n . For any pair of datasets x_i and x_j , we have $d_{\text{Ham}}(x_i, x_j) \leq d$. Let $\{S_i\}_{i \in [M]}$ be a collection of disjoint subsets of \mathcal{S} . If there exists an ε -DP algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{S}$ such that for every*

$i \in [M]$, $\Pr(\mathcal{A}(x_i) \in S_i) \geq 9/10$, then

$$\varepsilon = \Omega\left(\frac{\log M}{d}\right).$$

Our ε -DP Fano’s inequality (Theorem 2) and its corollary for distribution estimation (Corollary 4) can be viewed as a probabilistic packing argument which generalizes Lemma 5 to the case where \mathcal{V} consists of distributions over \mathcal{X}^n instead of deterministic datasets. The distances between distributions are measured in the minimum expected hamming distance between random datasets generated from a coupling between the distributions. Lemma 5 can be obtained from ε -DP Fano’s inequality by setting the distributions to be point masses over \mathcal{X}^n .

Note that d in Lemma 5 is an upper bound on the worst-case Hamming distance while D is a bound on the expected Hamming distance and therefore $D \leq d$. In statistical applications where $D \ll d$, we can obtain stronger lower bounds by replacing d with D . For example, in the k -ary distribution estimation problem, a naive application of the packing argument can only give a lower bound of $n = \Omega(k \log(1/\alpha)/\varepsilon)$ instead of the optimal $n = \Omega(k/\alpha\varepsilon)$ lower bound, where there is an exponential gap in the parameter $1/\alpha$.

Approximate DP and Fingerprinting. Fingerprinting (Steinke and Ullman, 2015; Bun et al., 2015; Dwork et al., 2015; Steinke and Ullman, 2017a; Bun et al., 2017, 2018; Kamath et al., 2019; Cai et al., 2019) is a versatile lower bounding method for (ε, δ) -DP for $\delta = O(1/n)$. It has been used to prove lower bounds for several problems, including attribute mean estimation in databases (Steinke and Ullman, 2017a), lower bounds on the number of online statistical queries (Bun et al., 2017), and private selection problem (Steinke and Ullman, 2017b). Kamath et al. (2019) uses fingerprinting to prove lower bounds on estimating Bernoulli product distributions and Gaussian distributions. We believe fingerprinting and DP Assouad’s lemma are both powerful tools for proving lower bounds under approximate DP. In estimating Gaussian distributions, fingerprinting provides strong lower bounds under approximate DP, whereas private Assouad’s method gives an additional polynomial blow-up compared to fingerprinting. However, for discrete distribution estimation, private Assouad’s method provides tight lower bounds, and we do not know how to obtain such bounds from the fingerprinting lemma.

Duchi et al. (2013) derives analogues of Le Cam, Assouad, and Fano in the local model of differential privacy, and uses them to establish lower bounds for several problems under local differential privacy. Acharya et al. (2019b, 2020) proves lower bounds for various testing and estimation problems under local differential privacy using a notion of chi-squared contractions based on Le Cam’s method and Fano’s inequality.

3. ε -DP distribution estimation

In this section, we use Corollary 4 to prove sample complexity lower bounds for various ε -DP distribution estimation problems. The general idea is to construct a subset of distributions in \mathcal{Q} such that they are close in both TV distance and KL divergence while being separated in the loss function ℓ . The larger the subsets we construct, the better the lower bounds we can get. In Section 3.1, we derive sample complexity lower bounds for k -ary distribution estimation under both TV and ℓ_2 distance that are tight up to constant factors. Tight sample complexity lower bounds up to logarithmic factors for (k, d) -product distributions and d -dimensional Gaussian mixtures are derived in Section 3.2 and 3.3 respectively.

Corollary 4 requires a *packing* of distributions with pairwise distance at least 3τ apart in ℓ . A standard method to construct such distributions is using results from coding theory.

We start with some definitions. An h -ary code of length k is a set $\mathcal{C} \subseteq \{0, 1, \dots, h-1\}^k$, and each $c \in \mathcal{C}$ is a *codeword*. The *minimum distance* of a code \mathcal{C} is the smallest Hamming distance between two codewords in \mathcal{C} . The code is called binary when $h = 2$. The weight of a binary codeword $c \in \mathcal{C}$ is $wt(c) = |\{i : c_i = 1\}|$, the number of 1's in c . A binary code \mathcal{C} is a *constant weight code* if each $c \in \mathcal{C}$ has the same weight. We now present some useful variants of the classic Gilbert Varshamov bounds on the existence of codes with certain properties. We prove these in Section 6.

Lemma 6 *Let l be an integer at most $k/2$ and at least 20. There exists a constant weight binary code \mathcal{C} which has code length k , weight l , minimum distance $l/4$ with $|\mathcal{C}| \geq \left(\frac{k}{2^{7/8}l}\right)^{7l/8}$.*

Lemma 7 *There exists an h -ary code \mathcal{H} with code length d and minimum Hamming distance $\frac{d}{2}$, which satisfies that $|\mathcal{H}| \geq \left(\frac{h}{16}\right)^{\frac{d}{2}}$.*

3.1. k -ary distribution estimation

We establish the sample complexity of ε -DP k -ary distribution estimation under TV and ℓ_2 distance.

Theorem 8 *The sample complexity of ε -DP k -ary distribution estimation under TV distance is*

$$S(\Delta_k, d_{TV}, \alpha, \varepsilon) = \Theta\left(\frac{k}{\alpha^2} + \frac{k}{\alpha\varepsilon}\right). \quad (8)$$

Theorem 9 *The sample complexity of ε -DP k -ary distribution estimation under ℓ_2 distance is*

$$S(\Delta_k, \ell_2, \alpha, \varepsilon) = \Theta\left(\frac{1}{\alpha^2} + \frac{\sqrt{k}}{\alpha\varepsilon}\right), \text{ for } \alpha < \frac{1}{\sqrt{k}}, \text{ and} \quad (9)$$

$$\Omega\left(\frac{1}{\alpha^2} + \frac{\log(k\alpha^2)}{\alpha^2\varepsilon}\right) \leq S(\Delta_k, \ell_2, \alpha, \varepsilon) \leq O\left(\frac{1}{\alpha^2} + \frac{\log k}{\alpha^2\varepsilon}\right) \text{ for } \alpha > \frac{1}{\sqrt{k}}. \quad (10)$$

For ℓ_2 loss, our bounds are tight within constant factors when $\alpha < \frac{1}{\sqrt{k}}$ or $\alpha > k^{-(\frac{1}{2}-0.001)}$.

3.1.1. TOTAL VARIATION DISTANCE

In this section, we derive the sample complexity of ε -DP k -ary distribution estimation under TV distance, which is stated in Theorem 8.

Upper bound: [Diakonikolas et al. \(2015\)](#) provides an upper bound based on Laplace mechanism ([Dwork et al., 2006](#)). We state the algorithm and a proof for completeness and we will use it for estimation under ℓ_2 distance.

Given a X^n from an unknown distribution p over $[k]$. Let $M_x(X^n)$ be the number of appearances of x in X^n . Let p^{erm} be the empirical estimator where $p^{\text{erm}}(x) := \frac{M_x(X^n)}{n}$. We note that changing one X_i in X^n can change at most two coordinates of p^{erm} , each by at most $\frac{1}{n}$, and thus changing

one X_i changes the p^{erm} by at most $2/n$ in ℓ_1 distance. Therefore, by [Dwork et al. \(2006\)](#), adding a Laplace noise of parameter $2/n\varepsilon$ to each coordinate of p^{erm} makes it ε -DP. For $x \in [k]$, let

$$h(x) = p^{\text{erm}}(x) + \text{Lap}\left(\frac{2}{n\varepsilon}\right),$$

where $\text{Lap}(\beta)$ is a Laplace random variable with parameter β . The final output \hat{p} is the projection of h on the simplex Δ_k in ℓ_2 distance. The expected ℓ_2 loss between h and p can be upper bounded by

$$(\mathbb{E}[\|h - p\|_2])^2 \leq \mathbb{E}[\|h - p\|_2^2] \leq \mathbb{E}[\|p^{\text{erm}} - p\|_2^2] + \mathbb{E}[\|h - p^{\text{erm}}\|_2^2],$$

where the first inequality comes from the Jensen's inequality and the second inequality comes from the triangle inequality.

The first term $\mathbb{E}[\|p^{\text{erm}} - p\|_2^2]$ is upper bounded by $\frac{1}{n}$ by an elementary analysis of the empirical estimator. For the second term, note that $\mathbb{E}[\|h - p^{\text{erm}}\|_2^2] = \sum_{i=1}^k \mathbb{E}[Z_i^2]$, where $\forall i, Z_i \sim \text{Lap}\left(\frac{2}{n\varepsilon}\right)$. By the variance of Laplace distribution, we have $\mathbb{E}[\|p^{\text{erm}} - h\|_2^2] = O\left(\frac{k}{n^2\varepsilon^2}\right)$. Therefore $\mathbb{E}[\|h - p\|_2] \leq O\left(\frac{1}{\sqrt{n}} + \frac{\sqrt{k}}{n\varepsilon}\right)$.

Note that since Δ_k is convex, $\|\hat{p} - p\|_2 \leq \|h - p\|_2$. Finally, by Cauchy-Schwarz Inequality, $\mathbb{E}[\|\hat{p} - p\|_1] \leq \sqrt{k} \cdot \mathbb{E}[\|\hat{p} - p\|_2] \leq \sqrt{k} \cdot \mathbb{E}[\|h - p\|_2] = O\left(\sqrt{\frac{k}{n}} + \frac{k}{n\varepsilon}\right)$. Therefore $\mathbb{E}[\|\hat{p} - p\|_1] \leq \alpha$ when $n = O\left(\frac{k}{\alpha^2} + \frac{k}{\alpha\varepsilon}\right)$.

Lower bound. We will construct a large set of distributions such that the conditions of [Corollary 4](#) hold. Suppose $\alpha < 1/48$. Applying [Lemma 6](#) with $l = k/2$, there exists a constant weight binary code \mathcal{C} of weight $k/2$, and minimum distance $k/8$, and $|\mathcal{C}| > 2^{7k/128}$. For each codeword $c \in \mathcal{C}$, a distribution p_c over $[k]$ is defined as follows:

$$p_c(i) = \begin{cases} \frac{1+24\alpha}{k}, & \text{if } c_i = 1, \\ \frac{1-24\alpha}{k}, & \text{if } c_i = 0. \end{cases}$$

We choose $\mathcal{V} = \{p_c : c \in \mathcal{C}\}$ to apply [Corollary 4](#). By the minimum distance property, any two distributions in \mathcal{V} have a total variation distance of at least $24\alpha/k \cdot k/8 = 3\alpha$, and at most 24α . Furthermore, by using $\log(1+x) \leq x$, we can bound the KL divergence between distributions by their χ^2 distance,

$$d_{KL}(p, q) \leq \chi^2(p, q) = \sum_{x=1}^k \frac{(p(x) - q(x))^2}{q(x)} < 10000\alpha^2.$$

Setting $\tau = \alpha$, $\gamma = 24\alpha$, and $\beta = 10000\alpha^2$, and using $\log M > 7k/64$ in [Corollary 4](#), we obtain $S(\Delta_k, d_{TV}, \alpha, \varepsilon) = \Omega\left(\frac{k}{\alpha^2} + \frac{k}{\alpha\varepsilon}\right)$.

3.1.2. ℓ_2 DISTANCE

In this section, we derive the sample complexity of ε -DP k -ary distribution estimation under ℓ_2 distance, which is stated in [Theorem 9](#).

Upper bound: We use the same algorithm as in Section 3.1.1. Following the same argument as in Section 3.1.1, the square of expected ℓ_2 loss of \hat{p} can be upper bounded by

$$(\mathbb{E} [\|\hat{p} - p\|_2])^2 \leq \mathbb{E} [\|h - p\|_2^2] \leq \mathbb{E} [\|p^{\text{erm}} - p\|_2^2] + \mathbb{E} [\|h - p^{\text{erm}}\|_2^2] = O\left(\frac{1}{n} + \frac{k}{n^2\varepsilon^2}\right).$$

Since Δ_k is convex, we have $\|\hat{p} - p\|_2 \leq \|h - p\|_2$. Moreover, the following lemma gives another bound for $\|\hat{p} - p\|_2$ (See Corollary 2.3 in Bassily (2019)).

Lemma 10 *Let $L \subset \mathbb{R}^d$ be a symmetric convex body of k vertices $\{a_j\}_{j=1}^k$, and let $y \in L$ and $\bar{y} = y + z$ for some $z \in \mathbb{R}^d$. Let $\hat{y} = \arg \min_{w \in L} \|w - \bar{y}\|_2$. Then we must have*

$$\|y - \hat{y}\|_2^2 \leq 4 \max_{j \in [k]} \langle z, a_j \rangle.$$

From the lemma, we have $\mathbb{E} [\|\hat{p} - h\|_2^2] \leq 4 \cdot \mathbb{E} [\max_{j \in [k]} |Z_j|]$, where $\forall j \in [k]$, $Z_j \sim \text{Lap}(\frac{2}{n\varepsilon})$. Note that $\mathbb{E} [\max |Z_j|] = O\left(\frac{\log k}{n\varepsilon}\right)$ due to the tail bound of Laplace distribution. We have $(\mathbb{E} [\|\hat{p} - p\|_2])^2 = O\left(\frac{1}{n} + \frac{\log k}{n\varepsilon}\right)$. Combined with the previous analysis, $(\mathbb{E} [\|\hat{p} - p\|_2])^2 = O\left(\frac{1}{n} + \min\left(\frac{k}{n^2\varepsilon^2}, \frac{\log k}{n\varepsilon}\right)\right)$. Therefore $\mathbb{E} [\|\hat{p} - p\|_2] \leq \frac{1}{10}\alpha$ when $n = O\left(\frac{1}{\alpha^2} + \min\left(\frac{\sqrt{k}}{\alpha\varepsilon}, \frac{\log k}{\alpha^2\varepsilon}\right)\right)$.

Lower bound: We first consider the case when $\alpha < \frac{1}{\sqrt{k}}$, where we can derive the lower bound simply by a reduction. By Cauchy-Schwarz inequality, for any estimator \hat{p} , $\mathbb{E} [\|\hat{p} - p\|_1] \leq \sqrt{k} \cdot \mathbb{E} [\|\hat{p} - p\|_2]$. Therefore $S(\Delta_k, \ell_2, \alpha, \varepsilon) \geq S(\Delta_k, d_{TV}, \sqrt{k}\alpha, \varepsilon)$, which gives us $S(\Delta_k, \ell_2, \alpha, \varepsilon) = \Omega\left(\frac{1}{\alpha^2} + \frac{\sqrt{k}}{\alpha\varepsilon}\right)$.

Now we consider $\alpha \geq \frac{1}{\sqrt{k}}$. Note that it is enough if we prove the lower bound of $\Omega\left(\frac{\log(\alpha^2 k)}{\alpha^2\varepsilon}\right)$, since $\Omega\left(\frac{1}{\alpha^2}\right)$ is the sample complexity of non-private estimation problem for all range of α . Similarly, we follow Corollary 4, except that we need to construct a different set of distributions.

Without loss of generality, we assume $\alpha < 0.1$. Now we use the codebook in Lemma 6 to construct our distribution set. We fix weight $l = \lfloor \frac{1}{50\alpha^2} \rfloor$. Note that for any $x > 2$, $\lfloor x \rfloor > \frac{x}{2}$. Then we have $\frac{1}{100\alpha^2} < \lfloor l \rfloor \leq \frac{1}{50\alpha^2}$ since $\frac{1}{50\alpha^2} > 2$. Therefore we get a codebook \mathcal{C} with $|\mathcal{C}| \geq (k\alpha^2)^{\frac{1}{200\alpha^2}}$. Given $c \in \mathcal{C}$, we construct the following distribution p_c in Δ_k :

$$p_c(i) = \frac{1}{l} c_i.$$

We use $\mathcal{V}_k = \{p_c : c \in \mathcal{C}\}$ to denote the set of all these distributions. It is easy to check that $\forall p \in \mathcal{V}_k$ is a valid distribution. Moreover, for any pair of distributions $p, q \in \mathcal{V}_k$, we have $\|p - q\|_2 > \frac{1}{2\sqrt{l}} = \Omega(\alpha)$.

For any pair $p, q \in \mathcal{V}_k$, $d_{TV}(p, q) \leq 1$, which is a naive upper bound for TV distance. Finally by setting ℓ in Corollary 4 to be ℓ_2 distance, we have $S(\Delta_k, \ell_2, \alpha, \varepsilon) = \Omega\left(\frac{\log|\mathcal{C}|}{\varepsilon}\right) = \Omega\left(\frac{\log(k\alpha^2)}{\alpha^2\varepsilon}\right)$.

3.2. Product distribution estimation

Recall that $\Delta_{k,d}$ is the set of all (k, d) -product distributions. Bun et al. (2019) proves an upper bound of $O\left(kd \log\left(\frac{kd}{\alpha}\right)\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)\right)$. We prove a sample complexity lower bound for ε -DP (k, d) -product distribution estimation in Theorem 11, which is optimal up to logarithmic factors.

Theorem 11 *The sample complexity of ε -DP (k, d) -product distribution estimation satisfies*

$$S(\Delta_{k,d}, d_{TV}, \alpha, \varepsilon) = \Omega\left(\frac{kd}{\alpha^2} + \frac{kd}{\alpha\varepsilon}\right).$$

Proof We start with the construction of the distribution set. First we use the same binary code as in Lemma 6 with weight $l = \frac{k}{2}$. Let $h := |\mathcal{C}|$ denote the size of the codebook. Given $j \in [h]$, we construct the following k -ary distribution p_j based on $c_j \in \mathcal{C}$:

$$p_j(i) = \frac{1}{k} + \frac{\alpha}{k\sqrt{d}} \cdot \mathbb{I}(c_{j,i} = 1),$$

where $c_{j,i}$ denotes the i -th coordinate of c_j .

Now we have designed a set of k -ary distributions of size $h = \Omega\left(2^{\frac{7k}{128}}\right)$. To construct a set of product distributions, we use the codebook construction in Lemma 7 to get an h -ary codebook \mathcal{H} with length d and minimum hamming distance $d/2$. Moreover, $|\mathcal{H}| \geq \left(\frac{h}{16}\right)^{\frac{d}{2}}$.

Now we can construct the distribution set of (k, d) -product distributions. Given $b \in \mathcal{H}$, define

$$p_b = p_{b_1} \times p_{b_2} \times \cdots \times p_{b_d}.$$

Let $\mathcal{V}_{k,d}$ denote the set of distributions induced by \mathcal{H} . We want to prove that $\forall p \neq q \in \mathcal{V}_{k,d}$,

$$d_{TV}(p, q) \geq C\alpha, \tag{11}$$

$$D_{KL}(p, q) \leq 4\alpha^2, \tag{12}$$

for some constant C . Suppose these two inequalities hold, using (12), by Pinsker's Inequality, we get $d_{TV}(p, q) \leq \sqrt{2D_{KL}(p, q)} \leq 2\sqrt{2}\alpha$. Then using Corollary 4, we can get

$$S(\Delta_{k,d}, d_{TV}, \alpha, \varepsilon) = \Omega\left(\frac{kd}{\alpha^2} + \frac{kd}{\alpha\varepsilon}\right).$$

Now it remains to prove (11) and (12). For (12), note that for any distribution pair $p, q \in \mathcal{V}_{k,d}$,

$$D_{KL}(p, q) \leq d \cdot \max_{i,j \in [h]} d_{KL}(p_i, p_j) \leq 4\alpha^2,$$

where the first inequality comes from the additivity of KL divergence for independent distributions and $\forall i, j \in [h]$,

$$d_{KL}(p_i, p_j) = \sum_{x \in [k]} p_i(x) \log \frac{p_i(x)}{p_j(x)} \leq \sum_{x \in [k]} \frac{(p_i(x) - p_j(x))^2}{p_j(x)} \leq k \left(\frac{\alpha}{k\sqrt{d}}\right)^2 / \frac{1}{k} = \frac{\alpha^2}{d}.$$

Next we prove (11). For any $b \in \mathcal{H}$ and $\forall i \in [k]$, define set

$$S_i = \{j \in [k] : c_{b_i,j} = 1\},$$

which contains the locations of +1's in the code at the i th coordinate of b . Based on this, we define a product distribution

$$p'_b = \prod_{i=1}^d \mathcal{B}(\mu_i),$$

where $\mu_i = \sum_{j \in S_i} p_{b_i}(j)$ and $\mathcal{B}(t)$ is a Bernoulli distribution with mean t . For any $b' \neq b \in \mathcal{H}$, we define

$$p'_{b'} = \prod_{i=1}^d \mathcal{B}(\mu'_i),$$

where $\mu'_i = \sum_{j \in S_i} p_{b'_i}(j)$. Then we have:

$$d_{TV}(p'_b, p'_{b'}) \leq d_{TV}(p_b, p_{b'}),$$

since p'_b and $p'_{b'}$ can be viewed as a post processing of p_b and $p_{b'}$ by mapping elements in S_i to 1 and others to 0 at the i -th coordinate. Moreover, we have $d_{\text{Ham}}(b, b') \geq \frac{d}{2}$, and $\forall i$, if $b_i \neq b'_i$, we have $d_H(c_{b_i}, c_{b'_i}) > \frac{k}{8}$. By the definition of p_i 's, we have

$$\|\mu_1 - \mu_2\|_2^2 \geq \frac{d}{2} \times \left(\frac{k}{8} \times \frac{\alpha}{k\sqrt{d}} \right)^2 = \frac{\alpha^2}{128}.$$

By Lemma 6.4 in [Kamath et al. \(2019\)](#), there exists a constant C such that $d_{TV}(p'_b, p'_{b'}) \geq C\alpha$, proving (11). ■

3.3. Gaussian mixtures estimation

Recall $\mathcal{G}_d = \{\mathcal{N}(\mu, I_d) : \|\mu\|_2 \leq R\}$ is the set of d -dimensional spherical Gaussians with unit variance and bounded mean and $\mathcal{G}_{k,d} = \{p : p \text{ is a } k\text{-mixture of } \mathcal{G}_d\}$ consists of mixtures of k distributions in \mathcal{G}_d . [Bun et al. \(2019\)](#) proves an upper bound of $\tilde{O}\left(\frac{kd}{\alpha^2} + \frac{d}{\alpha\varepsilon}\right)$ for estimating k -mixtures of Gaussians. We provide a sample complexity lower bound for estimating mixtures of Gaussians in [Theorem 12](#), which matches the upper bound up to logarithmic factors.

Theorem 12 *Given $k \leq d$ and $R \geq \sqrt{64 \log\left(\frac{8k}{\alpha}\right)}$, or $k \geq d$ and $R \geq (k)^{\frac{1}{d}} \cdot \sqrt{64d \log\left(\frac{8k}{\alpha}\right)}$,*

$$S(\mathcal{G}_{k,d}, d_{TV}, \alpha, \varepsilon) = \Omega\left(\frac{kd}{\alpha^2} + \frac{kd}{\alpha\varepsilon}\right).$$

Proof

We first consider the case when $k \leq d$ and $R \geq \sqrt{64 \log\left(\frac{8k}{\alpha}\right)}$. Let \mathcal{C} denote the codebook in [Lemma 6](#) with weight $l = \frac{d}{2}$. Then we have $|\mathcal{C}| \geq 2^{\frac{7k}{128}}$. Given c_i in codebook \mathcal{C} , we construct the following d -dimensional Gaussian distribution p_i , with identity covariance matrix and mean μ_i satisfying

$$\mu_{i,j} = \frac{\alpha}{\sqrt{d}} c_{i,j},$$

where $\mu_{i,j}$ denotes the j -th coordinate of μ_i .

Let $h = |\mathcal{C}|$. Similar to the product distribution case, using [Lemma 7](#), we can get an h -ary codebook \mathcal{H} with length d and minimum hamming distance $d/2$. Moreover, $|\mathcal{H}| \geq \left(\frac{h}{16}\right)^{\frac{d}{2}}$.

$\forall i \in [h]$ and $j \in k$, define $p_i^{(j)} = \mathcal{N}(\mu_i + \frac{R}{2}e_j, I_d)$, where e_j is the j th standard basis vector. It is easy to verify their means satisfy the norm bound. For a codeword $b \in \mathcal{H}$, let

$$p_b = \frac{1}{k} \left(p_{b_1}^{(1)} + p_{b_2}^{(2)} + \dots + p_{b_k}^{(k)} \right).$$

Let $\mathcal{V}_{\mathcal{G}} = \{p_b : b \in \mathcal{H}\}$ be the set of the distributions defined above. Next we prove that $\forall p_b \neq p_{b'} \in \mathcal{V}_{\mathcal{G}}$,

$$d_{TV}(p_b, p_{b'}) \geq C\alpha, \quad (13)$$

$$d_{KL}(p_b, p_{b'}) \leq 4\alpha^2. \quad (14)$$

where C is a constant. If these two inequalities hold, using (14), by Pinsker's Inequality, we get $d_{TV}(p_b, p_{b'}) \leq \sqrt{2d_{KL}(p_b, p_{b'})} \leq 2\sqrt{2}\alpha$. Using Corollary 4, we get

$$S(\mathcal{G}_{k,d}, d_{TV}, \alpha, \varepsilon) = \Omega\left(\frac{kd}{\alpha^2} + \frac{kd}{\alpha\varepsilon}\right).$$

It remains to prove (13) and (14).

For (14), note that for any distribution pair $p_b \neq p_{b'} \in \mathcal{V}_{\mathcal{G}}$,

$$d_{KL}(p_b, p_{b'}) \leq \frac{1}{k} \sum_{t=1}^k d_{KL}(p_{b_t}^{(t)}, p_{b'_t}^{(t)}) \leq \max_{i,j \in [h]} d_{KL}(p_i, p_j) \leq 4\alpha^2,$$

where the first inequality comes from the convexity of KL divergence and the last inequality uses the fact that the KL divergence between two Gaussians with identity covariance is at most the ℓ_2^2 distance between their means.

Next we prove (13). Let $B_j = B_{j,1} \times \dots \times B_{j,d}$, where

$$B_{j,i} = \begin{cases} [\frac{R}{4}, \frac{3R}{4}], & \text{when } i = j, \\ [-\frac{R}{4}, \frac{R}{4}], & \text{when } i \neq j \text{ and } i \leq k, \\ [-\infty, \infty], & \text{when } k < i \leq d. \end{cases}$$

Then by Gaussian tail bound and union bound, for any $p \in \mathcal{V}_{\mathcal{G}}$, the mass of the j -th Gaussian component outside B_j is at most $2ke^{-\frac{1}{2} \cdot (\frac{1}{4}R)^2}$. And the mass of other Gaussian components inside B_j is at most $e^{-\frac{1}{2} \cdot (\frac{1}{4}R)^2}$. Hence we have:

$$\begin{aligned} d_{TV}(p_b, p_{b'}) &= \frac{1}{2k} \int_{z \in \mathbb{R}^d} \left| p_{b_1}^{(1)}(z) + \dots + p_{b_k}^{(k)}(z) - p_{b'_1}^{(1)}(z) - \dots - p_{b'_k}^{(k)}(z) \right| dz \\ &\geq \frac{1}{2k} \sum_{j=1}^k \int_{z \in B_j} \left| p_{b_1}^{(1)}(z) + \dots + p_{b_k}^{(k)}(z) - p_{b'_1}^{(1)}(z) - \dots - p_{b'_k}^{(k)}(z) \right| dz \\ &\geq \frac{1}{2k} \cdot \sum_{j=1}^k \left(\int_{z \in B_j} \left| p_{b_j}^{(j)}(z) - p_{b'_j}^{(j)}(z) \right| dz - (k-1) \cdot e^{-\frac{1}{2} \cdot (\frac{1}{4}R)^2} \right) \\ &\geq \frac{1}{2k} \cdot \sum_{j=1}^k \left(\int_{z \in \mathbb{R}^d} \left| p_{b_j}^{(j)}(z) - p_{b'_j}^{(j)}(z) \right| dz - 3k \cdot e^{-\frac{1}{2} \cdot (\frac{1}{4}R)^2} \right) \\ &= \frac{1}{2k} \cdot \sum_{j=1}^k d_{TV}(p_{b_j}, p_{b'_j}) - \frac{3\alpha^2}{64k}. \end{aligned}$$

By Fact 6.6 in [Kamath et al. \(2019\)](#), there exists a constant C_1 such that for any pair $i \neq j \in [h]$,

$$d_{TV}(p_i, p_j) \geq C_1 \alpha.$$

Hence we have

$$\frac{1}{2k} \cdot \sum_{j=1}^k d_{TV}(p_{b_j}, p_{b'_j}) \geq \frac{C_1 \alpha}{2k} d_{\text{Ham}}(b, b') \geq \frac{C_1 \alpha}{4},$$

where the last inequality comes from the property of the codebook. WLOG, we can assume $\frac{3\alpha}{64k} < C_1/8$. Taking $C = \frac{C_1}{8}$ completes the proof of (13).

Now we consider the case when $k \geq d$ and $R \geq (k)^{\frac{1}{d}} \cdot \sqrt{64d \log\left(\frac{8k}{\alpha}\right)}$. Let $r = \sqrt{16d \log\left(\frac{8k}{\alpha}\right)}$, we note that there exists a packing set $S = \{v_1, v_2, \dots, v_k\} \subset \mathbb{R}^d$ which satisfies $\forall u, v \in S$,

$$\|u - v\|_2 > r, \quad \|u\|_2 \leq R, \quad \|v\|_2 \leq R/3,$$

and $|S| = k$ since $R \geq 2(k)^{\frac{1}{d}}r$. Consider the set of mixture distributions as following: For a codeword $b \in \mathcal{H}$, let

$$p'_b = \frac{1}{k} \left(p_{b_1}^{(1)'} + p_{b_2}^{(2)'} + \dots + p_{b_k}^{(k)'} \right),$$

where $\forall i \in [k], p_{b_j}^{(j)'} = \mathcal{N}(\mu_{b_j} + v_j, I_d)$. Let B'_j denote the ℓ_2 ball centering at the v_j with radius $\frac{r}{2}$. We note that by similar analysis using the tail bound of the Gaussian distribution, the mass of the j -th Gaussian component outside B'_j is at most $\frac{\alpha^2}{64k^2}$. Meanwhile, the mass of other Gaussian components inside B'_j is also at most $\frac{\alpha^2}{64k^2}$. Hence the remaining analysis follows from the previous case. \blacksquare

4. (ε, δ) -DP distribution estimation

In the previous section we used Theorem 2 to obtain sample complexity lower bounds for pure differential privacy. We will now use Theorem 3 to prove sample complexity lower bounds under (ε, δ) -DP.

4.1. k -ary distribution estimation

Theorem 13 *The sample complexity of (ε, δ) -DP k -ary distribution estimation under total variation distance is*

$$S(\Delta_k, d_{TV}, \alpha, \varepsilon, \delta) = \Omega\left(\frac{k}{\alpha^2} + \frac{k}{\alpha(\varepsilon + \delta)}\right).$$

In practice, δ is chosen to be $\delta = O\left(\frac{1}{n}\right)$, and the privacy parameter is chosen as a small constant, $\varepsilon = \Theta(1)$. In particular, when $\delta \leq \varepsilon$, the theorem above shows

$$S(\Delta_k, d_{TV}, \alpha, \varepsilon, \delta) = \Omega\left(\frac{k}{\alpha^2} + \frac{k}{\alpha\varepsilon}\right).$$

Since the sample complexity of ε -DP is at most the sample complexity of (ε, δ) -DP, this shows that the bound above is tight for $\delta \leq \varepsilon$. The lower bound part is proved using Theorem 3 in Section 4.1.1.

Theorem 14 *The sample complexity of (ε, δ) -DP discrete distribution estimation under ℓ_2 distance,*

$$\Omega\left(\frac{1}{\alpha^2} + \frac{\sqrt{k}}{\alpha(\varepsilon + \delta)}\right) \leq S(\Delta_k, \ell_2, \alpha, \varepsilon, \delta) \leq O\left(\frac{1}{\alpha^2} + \frac{\sqrt{k}}{\alpha\varepsilon}\right), \quad \text{for } \alpha < \frac{1}{\sqrt{k}},$$

$$\Omega\left(\frac{1}{\alpha^2} + \frac{1}{\alpha^2(\varepsilon + \delta)}\right) \leq S(\Delta_k, \ell_2, \alpha, \varepsilon, \delta) \leq O\left(\frac{1}{\alpha^2} + \frac{\log k}{\alpha^2\varepsilon}\right), \quad \text{for } \alpha > \frac{1}{\sqrt{k}}.$$

When $\delta = O(\varepsilon)$, the bounds are tight when $\alpha < 1/\sqrt{k}$ and differ by a factor of $\log k$ when $\alpha \geq 1/\sqrt{k}$. We prove this result in Section 4.1.2.

4.1.1. PROOF OF THEOREM 13.

The first term k/α^2 is the tight sample complexity without privacy. We prove that $S(\Delta_k, d_{TV}, \alpha, \varepsilon, \delta) = \Omega\left(\frac{k}{\alpha(\varepsilon + \delta)}\right)$.

Suppose k is even and $\alpha < 1/10$. Let $\mathcal{E}_{k/2} = \{-1, +1\}^{k/2}$, for $e \in \mathcal{E}_{k/2}$, we define $p_e \in \Delta_k$ as follows.

$$\text{For } i = 1, \dots, k/2 \quad p_e(2i-1) = \frac{1 + 10e_i \cdot \alpha}{k}, \quad p_e(2i) = \frac{1 - 10e_i \cdot \alpha}{k}. \quad (15)$$

To apply Theorem 3, let $\mathcal{V}_{k/2} = \{p_e^n, e \in \mathcal{E}_{k/2}\}$. p_e^n is the distribution of n i.i.d. samples from distribution p_e , and $\theta(p_e^n) = p_e$. For $u, v \in \mathcal{E}_{k/2}$,

$$\ell(\theta(p_u^n), \theta(p_v^n)) = d_{TV}(p_u, p_v) = \frac{20\alpha}{k} \cdot \sum_{i=1}^{k/2} \mathbb{I}(u_i \neq v_i),$$

thus obeying (6) with $\tau = 10\alpha/k$.

Recall the mixture distributions p_{+i} and p_{-i} ,

$$p_{+i} = \frac{2}{|\mathcal{E}_{k/2}|} \sum_{e \in \mathcal{E}_{k/2}: e_i = +1} p_e^n, \quad p_{-i} = \frac{2}{|\mathcal{E}_{k/2}|} \sum_{e \in \mathcal{E}_{k/2}: e_i = -1} p_e^n.$$

To apply Theorem 3, we prove the following bound on the Hamming distance between a coupling between p_{+i} and p_{-i} .

Lemma 15 *For any i , there is a coupling (X, Y) between p_{+i} and p_{-i} , such that*

$$\mathbb{E}[d_{Ham}(X, Y)] \leq \frac{20\alpha n}{k}.$$

Proof By the construction in (15), note that the distributions p_{+i} and p_{-i} only have a difference in the number of times $2i-1$ and $2i$ appear. To generate $Y \sim p_{-i}$ from from $X \sim p_{+i}$, we scan through X and independently change the coordinates that have the symbol $2i-1$ to the symbol $2i$ with probability $\frac{20\alpha}{1+10\alpha}$. The expected Hamming distance is bounded by $\frac{20\alpha}{1+10\alpha} \cdot \frac{1+10\alpha}{k} \cdot n = \frac{20\alpha n}{k}$. ■

Note that $\mathcal{V} \subset \mathcal{P} := \{p^n | p \in \Delta_k\}$. By Theorem 3, using the bound on D from Lemma 18, and $\tau = 10\alpha/k$,

$$R(\mathcal{P}, d_{TV}, \varepsilon, \delta) \geq \frac{5\alpha}{k} \cdot k \cdot \left(0.9e^{-10\varepsilon D} - 10D\delta\right) \geq 5\alpha \cdot \left(0.9e^{-200n\varepsilon\alpha/k} - 200\frac{n\varepsilon\alpha\delta}{k}\right).$$

To achieve $R(\mathcal{P}, d_{TV}, \varepsilon, \delta) \leq \alpha$, either $n\varepsilon\alpha/k = \Omega(1)$ or $n\varepsilon\alpha\delta/k = \Omega(1)$, which implies that $n = \Omega\left(\frac{k}{\alpha(\varepsilon + \delta)}\right)$.

4.1.2. PROOF OF THEOREM 14

We first consider the case where $\alpha < \frac{1}{\sqrt{k}}$. By Cauchy-Schwarz inequality, $S(\Delta_k, \ell_2, \alpha, \varepsilon, \delta) \geq S(\Delta_k, d_{TV}, \sqrt{k}\alpha, \varepsilon, \delta)$, and therefore $S(\Delta_k, \ell_2, \alpha, \varepsilon, \delta) = \Omega\left(\frac{1}{\alpha^2} + \frac{\sqrt{k}}{\alpha(\varepsilon+\delta)}\right)$ by Theorem 13.

For $\alpha \geq \frac{1}{\sqrt{k}}$, we have $l = \lfloor \frac{1}{16\alpha^2} \rfloor \leq k$. Therefore, $\Delta_l \subset \Delta_k$ and $\alpha < \frac{1}{\sqrt{l}}$. Hence,

$$S(\Delta_k, \ell_2, \alpha, \varepsilon, \delta) \geq S(\Delta_l, \ell_2, \alpha, \varepsilon, \delta) = \Omega\left(\frac{1}{\alpha^2} + \frac{1}{\alpha^2(\varepsilon + \delta)}\right).$$

4.2. Binary product distribution estimation

We now consider estimation of Bernoulli product distributions under total variation distance. A Bernoulli product distribution in d dimensions is a distribution over $\{0, 1\}^d$ parameterized by $\mu \in [0, 1]^d$, where the i th coordinate is distributed $\mathcal{B}(\mu_i)$, where $\mathcal{B}(\cdot)$ is a Bernoulli distribution. Let $\Delta_{2,d}$ be the class of Bernoulli product distributions in d dimensions.

Theorem 16 *The sample complexity of (ε, δ) -DP binary product distribution estimation satisfies*

$$S(\Delta_{2,d}, d_{TV}, \alpha, \varepsilon, \delta) = \Omega\left(\frac{d}{\alpha^2} + \frac{d}{\alpha(\varepsilon + \delta)}\right).$$

Compared to the upper bound of $O(d \log(d/\alpha)(1/\alpha^2 + 1/\alpha\varepsilon))$ in Bun et al. (2019); Kamath et al. (2019), our bound is tight up to logarithmic factors when $\delta \leq \varepsilon$. Kamath et al. (2019) also presents a lower bound of $\Omega\left(\frac{d}{\alpha^2} + \frac{d}{\alpha(\varepsilon+\delta)}\right)$ under (ε, δ) -DP when $\delta = O(1/n)$. Although $\delta = O(1/n)$ is the more interesting regime in practice, our bound complements the result by stating that the utility will not improve even if δ can be as large as ε .

Proof Since $\Theta(d/\varepsilon^2)$ is an established tight bound for non-private estimation, we only prove the second term.

We start by constructing a set of Bernoulli product distributions indexed by $\mathcal{E}_d = \{\pm 1\}^d$. For all $e \in \mathcal{E}_d$, let $p_e = \mathcal{B}(\mu_1^e) \times \mathcal{B}(\mu_2^e) \times \cdots \times \mathcal{B}(\mu_d^e)$, where

$$\mu_i^e = \frac{1 + e_i \cdot 20\alpha}{d}.$$

Let $\mathcal{V} = \{p_e^n, e \in \mathcal{E}_d\}$, the set of distributions of n i.i.d samples from p_e , and $\theta(p_e^n) = p_e$. For $u, v \in \mathcal{E}_d$, $\ell(\theta(p_u^n), \theta(p_v^n)) = d_{TV}(p_u, p_v)$. We first prove that (6) holds under total variation distance for an appropriate τ .

Lemma 17 *There exists a constant $C_1 > 5$ such that $\forall u, v \in \mathcal{E}_d$,*

$$d_{TV}(p_u, p_v) \geq \frac{C_1\alpha}{d} \cdot \sum_{i=1}^d \mathbb{I}(u_i \neq v_i).$$

Proof Let $S = \{i \in [d] : u_i \neq v_i\}$, and $S' = \{i \in S : u_i = 1\}$. WLOG, let $|S'| \geq \frac{1}{2}|S|$ (or else we can define $S' = \{i \in S : u_i = -1\}$). Given a random sample $Z \in \{\pm 1\}^d$, we define an event

$A = \{\forall i \in S', Z_i = 0\}$. Now we consider the difference between the following two probabilities, which is a lower bound of the total variation distance between p_u and p_v .

$$\begin{aligned} d_{TV}(p_u, p_v) &\geq |\Pr_{Z \sim p_u}(A) - \Pr_{Z \sim p_v}(A)| \\ &= \left(1 - \frac{1 - 20\alpha}{d}\right)^{|S'|} - \left(1 - \frac{1 + 20\alpha}{d}\right)^{|S'|} \\ &\geq \frac{40\alpha}{d} \cdot |S'| \cdot \left(1 - \frac{1 + 20\alpha}{d}\right)^{|S'|} \\ &\geq \frac{40\alpha}{d} \cdot |S'| e^{-(1+20\alpha)} \geq \frac{C_1\alpha}{d} \cdot d_{\text{Ham}}(u, v), \end{aligned}$$

where in the last two inequalities, we assume $d \geq 1000$ and $\alpha < 0.01$. \blacksquare

Let D be an upper bound on the expected Hamming distance for a coupling between p_{+i} and p_{-i} over all i . Since $\mathcal{V}_d \subset \Delta_{2,d}$, applying Theorem 3 with Lemma 17 we have

$$R(\mathcal{P}, d_{TV}, \varepsilon, \delta) \geq \frac{C_1\alpha}{2d} \cdot d \cdot (0.9e^{-10\varepsilon D} - 10D\delta) = \frac{C_1\alpha}{2} \cdot (0.9e^{-10\varepsilon D} - 10D\delta).$$

Setting $R(\mathcal{P}, d_{TV}, \varepsilon, \delta) \leq \alpha$, we get $D = \Omega\left(\frac{1}{\varepsilon}\right)$ or $D = \Omega\left(\frac{1}{\delta}\right)$, or equivalently, $D = \Omega\left(\frac{1}{\varepsilon + \delta}\right)$. Lemma 18 below shows that we can take $D = \frac{40\alpha n}{d}$, which proves the result. \blacksquare

Lemma 18 *There is a coupling between (X, Y) between p_{+i} and p_{-i} , such that $\mathbb{E}[d_{\text{Ham}}(X, Y)] \leq \frac{40\alpha n}{d}$.*

Proof We generate $Y \sim p_{-i}$ from $X \sim p_{+i}$ as follows. If the i th coordinate of a sample X is $+1$, we independently flip it to -1 with probability $\frac{40\alpha}{1+20\alpha}$ to obtain a sample Y . The expected Hamming distance is bounded by $\frac{40\alpha}{1+20\alpha} \cdot \frac{1+20\alpha}{d} \cdot n = \frac{40\alpha n}{d}$. \blacksquare

5. Proof of Theorems

5.1. Proof of DP Le Cam's method (Theorem 1)

The proof technique is similar to the proof of coupling lemma in Acharya et al. (2018b). However, we directly characterize the error probability in Theorem 1, which we then use to prove Theorem 3 (DP Assoud's method).

Theorem 1 ((ε, δ)-DP Le Cam's method) *Let $p_1 \in \text{co}(\mathcal{P}_1)$ and $p_2 \in \text{co}(\mathcal{P}_2)$. Let (X, Y) be a coupling between p_1 and p_2 with $D = \mathbb{E}[d_{\text{Ham}}(X, Y)]$. Then for $\varepsilon \geq 0, \delta \geq 0$, any (ε, δ) -differentially private hypothesis testing algorithm $\hat{\theta}$ must satisfy*

$$P_e(\hat{\theta}, \mathcal{P}_1, \mathcal{P}_2) \geq \frac{1}{2} \max \left\{ 1 - d_{TV}(p_1, p_2), 0.9e^{-10\varepsilon D} - 10D\delta \right\}, \quad (4)$$

where $d_{TV}(p_1, p_2) := \sup_{A \subseteq \mathcal{X}^n} (p_1(A) - p_2(A)) = \frac{1}{2} \|p_1 - p_2\|_1$ is the total variation (TV) distance of p_1 and p_2 .

Proof From (3),

$$P_e(\hat{\theta}, \mathcal{P}_1, \mathcal{P}_2) \geq \frac{1}{2} \left(\Pr_{X \sim p_1} (\hat{\theta}(X) \neq p_1) + \Pr_{X \sim p_2} (\hat{\theta}(X) \neq p_2) \right).$$

The first term in Theorem 1 follows from the classic Le Cam's method (Lemma 1 in Yu (1997)). For the second term, let (X, Y) be distributed according to a coupling of p_1 and p_2 with $\mathbb{E}[d_{\text{Ham}}(X, Y)] \leq D$. By Markov's inequality, $\Pr(d_{\text{Ham}}(X, Y) > 10D) < 0.1$. Let x and y be the realization of X and Y . $W := \{(x, y) \in \mathcal{X}^n \times \mathcal{X}^n \mid d_{\text{Ham}}(x, y) \leq 10D\}$ be the set of pairs of realizations with Hamming distance at most $10D$. Then we have

$$\begin{aligned} \Pr(\hat{\theta}(X) = p_2) &= \sum_{x, y} \Pr(X = x, Y = y) \cdot \Pr(\hat{\theta}(x) = p_2) \\ &\geq \sum_{(x, y) \in W} \Pr(X = x, Y = y) \cdot \Pr(\hat{\theta}(x) = p_2). \end{aligned} \quad (16)$$

Let $\beta_1 = \Pr(\hat{\theta}(X) = p_2)$, so we have

$$\sum_{(x, y) \in W} \Pr(X = x, Y = y) \cdot \Pr(\hat{\theta}(x) = p_2) \leq \beta_1$$

Next, we need the following group property of differential privacy.

Lemma 19 *Let $\hat{\theta}$ be a (ε, δ) -DP algorithm, then for sequences x , and y with $d_{\text{Ham}}(x, y) \leq t$, and $\forall S$, $\Pr(\hat{\theta}(x) \in S) \leq e^{t\varepsilon} \cdot \Pr(\hat{\theta}(y) \in S) + \delta t e^{\varepsilon(t-1)}$.*

By Lemma 19, and $\Pr(d_{\text{Ham}}(X, Y) > 10D) < 0.1$, let $\Pr(\hat{\theta}(Y) = p_2) = 1 - \beta_2$,

$$\begin{aligned} 1 - \beta_2 &= \sum_{(x, y) \in W} \Pr(x, y) \cdot \Pr(\hat{\theta}(y) = p_2) + \sum_{(x, y) \notin W} \Pr(x, y) \cdot \Pr(\hat{\theta}(y) = p_2) \\ &\leq \sum_{(x, y) \in W} \Pr(x, y) \cdot \left(e^{\varepsilon \cdot 10D} \Pr(\hat{\theta}(x) = p_2) + 10D\delta \cdot e^{\varepsilon \cdot 10(D-1)} \right) + 0.1 \\ &\leq \beta_1 \cdot e^{\varepsilon \cdot 10D} + 10D\delta \cdot e^{\varepsilon \cdot 10D} + 0.1. \end{aligned}$$

Similarly, we get

$$1 - \beta_1 \leq \beta_2 \cdot e^{\varepsilon \cdot 10D} + 10D\delta \cdot e^{\varepsilon \cdot 10D} + 0.1.$$

Adding the two inequalities and rearranging terms,

$$\beta_1 + \beta_2 \geq \frac{1.8 - 20D\delta e^{\varepsilon \cdot 10D}}{1 + e^{\varepsilon \cdot 10D}} \geq 0.9e^{-10\varepsilon D} - 10D\delta.$$

■

5.2. Proof of private Fano's inequality (Theorem 2)

In this section, we prove ε -DP Fano's inequality (Theorem 2), restated below.

Theorem 2 (ε -DP Fano's inequality) *Let $\mathcal{V} = \{p_1, p_2, \dots, p_M\} \subseteq \mathcal{P}$ such that for all $i \neq j$,*

- (a) $\ell(\theta(p_i), \theta(p_j)) \geq \alpha$,
- (b) $D_{KL}(p_i, p_j) \leq \beta$,
- (c) *there exists a coupling (X, Y) between p_i and p_j such that $\mathbb{E}[d_{\text{Ham}}(X, Y)] \leq D$, then*

$$R(\mathcal{P}, \ell, \varepsilon) \geq \max \left\{ \frac{\alpha}{2} \left(1 - \frac{\beta + \log 2}{\log M} \right), 0.4\alpha \min \left\{ 1, \frac{M}{e^{10\varepsilon D}} \right\} \right\}. \quad (5)$$

The proof is based on the observation that if you can change a sample from p_i to p_j by changing D coordinates in expectation, then an algorithm that correctly outputs a sample as from p_i has to output p_j with probability roughly $e^{-\varepsilon D}$. With a total of M distributions in total, we show that the error probability is large as long as $\frac{M}{e^{\varepsilon D}}$ is large.

Proof The first term in (5) follows from the non-private Fano's inequality (Lemma 3 in Yu (1997)). For an observation $X \in \mathcal{X}^n$,

$$\hat{p}(X) := \arg \min_{p \in \mathcal{V}} \ell(\theta(p), \hat{\theta}(X)),$$

be the distribution in \mathcal{P} closest in parameters to an ε -DP estimate $\hat{\theta}(X)$. Therefore, $\hat{p}(X)$ is also ε -DP. By the triangle inequality,

$$\ell(\theta(\hat{p}), \theta(p)) \leq \ell(\theta(\hat{p}), \hat{\theta}(X)) + \ell(\theta(p), \hat{\theta}(X)) \leq 2\ell(\theta(p), \hat{\theta}(X)).$$

Hence,

$$\begin{aligned} \max_{p \in \mathcal{P}} \mathbb{E}_{X \sim p} [\ell(\hat{\theta}(X), \theta(p))] &\geq \max_{p \in \mathcal{V}} \mathbb{E}_{X \sim p} [\ell(\hat{\theta}(X), \theta(p))] \geq \frac{1}{2} \max_{p \in \mathcal{V}} \mathbb{E}_{X \sim p} [\ell(\theta(\hat{p}), \theta(p))] \\ &\geq \max_{p \in \mathcal{V}} \frac{\alpha}{2} \Pr_{X \sim p} (\hat{p}(X) \neq p) \\ &\geq \frac{\alpha}{2M} \sum_{p \in \mathcal{V}} \Pr_{X \sim p} (\hat{p}(X) \neq p). \end{aligned} \quad (17)$$

Let $\beta_i = \Pr_{X \sim p_i} (\hat{p}(X) \neq p_i)$ be the probability that $\hat{p}(X) \neq p_i$ when the underlying distribution generating X is p_i . For $p_i, p_j \in \mathcal{V}$, let (X, Y) be the coupling in condition (c). By Markov's inequality $\Pr(d_{\text{Ham}}(X, Y) > 10D) < 1/10$.

Similar to the proof of Theorem 1 in the previous section, let $W := \{(x, y) | d_{\text{Ham}}(x, y) \leq 10D\}$ and $\Pr(x, y) = \Pr(X = x, Y = y)$. Then

$$1 - \beta_j = \Pr(\hat{p}(Y) = p_j) \leq \sum_{(x, y) \in W} \Pr(x, y) \cdot \Pr(\hat{p}(y) = p_j) + \sum_{(x, y) \notin W} \Pr(x, y) \cdot 1.$$

Therefore,

$$\sum_{(x, y) \in W} \Pr(x, y) \cdot \Pr(\hat{p}(y) = p_j) \geq 0.9 - \beta_j.$$

Then, we have

$$\begin{aligned}
 \Pr(\hat{p}(X) = p_j) &\geq \sum_{(x,y) \in W} \Pr(x,y) \cdot \Pr(\hat{p}(x) = p_j) \\
 &\geq \sum_{(x,y) \in W} \Pr(x,y) e^{-10\varepsilon D} \Pr(\hat{p}(y) = p_j) \\
 &\geq (0.9 - \beta_j) e^{-10\varepsilon D},
 \end{aligned} \tag{18}$$

where (18) uses that \hat{p} is ε -DP and $d_{\text{Ham}}(x, y) \leq 10D$. Similarly, for all $j' \neq i$,

$$\Pr(\hat{p}(X) = p_{j'}) \geq (0.9 - \beta_{j'}) e^{-10\varepsilon D}.$$

Summing over $j' \neq i$, we obtain

$$\beta_i = \sum_{j' \neq i} \Pr(\hat{p}(X) = p_{j'}) \geq \left(0.9(M-1) - \sum_{j' \neq i} \beta_{j'} \right) e^{-10\varepsilon D}.$$

Summing over $i \in [M]$,

$$\sum_{i \in [M]} \beta_i \geq \left(0.9M(M-1) - (M-1) \sum_{i \in [M]} \beta_i \right) e^{-10\varepsilon D}.$$

Rearranging the terms

$$\sum_{i \in [M]} \beta_i \geq \frac{0.9M(M-1)}{M-1 + e^{10\varepsilon D}} \geq 0.8M \min \left\{ 1, \frac{M}{e^{10\varepsilon D}} \right\}.$$

Combining this with (17) completes the proof. ■

5.3. Proof of Corollary 4

Corollary 4 (ε -DP distribution estimation) *Given $\varepsilon > 0$, let $\mathcal{V} = \{q_1, q_2, \dots, q_M\} \subseteq \mathcal{Q}$ be a set of distributions over \mathcal{X} with size M , such that for all $i \neq j$,*

- (a) $\ell(q_i, q_j) \geq 3\tau$,
- (b) $D_{KL}(q_i, q_j) \leq \beta$,
- (c) $d_{TV}(q_i, q_j) \leq \gamma$,

then

$$S(\mathcal{Q}, \ell, \tau, \varepsilon) = \Omega \left(\frac{\log M}{\beta} + \frac{\log M}{\gamma \varepsilon} \right).$$

Proof Recall that $\mathcal{Q}^n := \{q^n | q \in \mathcal{Q}\}$ is the set of induced distributions over \mathcal{X}^n and $q^n \in \mathcal{Q}^n$, $\theta(q^n) = q$. Then, $\forall i \neq j \in [M]$, $\ell(\theta(q_i^n), \theta(q_j^n)) \geq 3\tau$, and $D_{KL}(q_i^n, q_j^n) = nD_{KL}(q_i, q_j) \leq n\beta$.

The following lemma is a corollary of maximal coupling ([den Hollander, 2012](#)), which states that for two distributions there is a coupling of their n fold product distributions with an expected Hamming distance n times their total variation distance.

Lemma 20 *Given distributions q_1, q_2 over \mathcal{X} , there exists a coupling (X, Y) between q_1^n and q_2^n such that*

$$\mathbb{E} [d_{\text{Ham}}(X, Y)] = n \cdot d_{TV}(q_1, q_2),$$

where $X \sim q_1^n$ and $Y \sim q_2^n$.

By Lemma 20, $\forall i, j \in [M]$, there exists a coupling (X, Y) between q_i^n and q_j^n such that $\mathbb{E} [d_{\text{Ham}}(X, Y)] \leq n\gamma$. Now by Theorem 2,

$$R(\mathcal{Q}^*, \ell, \varepsilon) \geq \max \left\{ \frac{3\tau}{2} \left(1 - \frac{n\beta + \log 2}{\log M} \right), 1.2\tau \min \left\{ 1, \frac{M}{e^{10\varepsilon n\gamma}} \right\} \right\}. \quad (19)$$

Therefore, for $R(\mathcal{Q}^n, \ell, \varepsilon) \leq \tau$,

$$S(\mathcal{Q}, \ell, \tau, \varepsilon) = \Omega \left(\frac{\log M}{\beta} + \frac{\log M}{\gamma\varepsilon} \right).$$

■

5.4. Proof of private Assouad's method (Theorem 3)

We restate the theorem below and the notions are the same as defined in Section 2.3.

Theorem 3 (DP Assouad's method) $\forall i \in [k]$, let $\phi_i : \mathcal{X}^n \rightarrow \{-1, +1\}$ be a binary classifier.

$$R(\mathcal{P}, \ell, \varepsilon, \delta) \geq \frac{\tau}{2} \cdot \sum_{i=1}^k \min_{\phi_i \text{ is } (\varepsilon, \delta)\text{-DP}} (\Pr_{X \sim p_{+i}} (\phi_i(X) \neq 1) + \Pr_{X \sim p_{-i}} (\phi_i(X) \neq -1)).$$

Moreover, if $\forall i \in [k]$, there exists a coupling (X, Y) between p_{+i} and p_{-i} with $\mathbb{E} [d_{\text{Ham}}(X, Y)] \leq D$,

$$R(\mathcal{P}, \ell, \varepsilon, \delta) \geq \frac{k\tau}{2} \cdot (0.9e^{-10\varepsilon D} - 10D\delta). \quad (7)$$

Proof The first part is from the non-private Assouad's lemma, which we include here for completeness. Let $p \in \mathcal{V} \subset \mathcal{P}$ and $X \sim p$. For an estimator $\hat{\theta}(X)$, consider an estimator $\hat{E}(X) = \arg \min_{e \in \mathcal{E}_k} \ell(\hat{\theta}(X), \theta(p_e))$. Then, by the triangle inequality,

$$\ell(\theta(p_{\hat{E}}), \theta(p)) \leq \ell(\hat{\theta}, \theta(p_{\hat{E}})) + \ell(\hat{\theta}, \theta(p)) \leq 2\ell(\hat{\theta}, \theta(p)).$$

Hence,

$$R(\mathcal{V}, \ell, \varepsilon, \delta) = \min_{\hat{\theta} \text{ is } (\varepsilon, \delta)\text{-DP}} \max_{p \in \mathcal{V}} \mathbb{E}_{X \sim p} [\ell(\hat{\theta}(X), \theta(p))] \geq \frac{1}{2} \min_{\hat{E} \text{ is } (\varepsilon, \delta)\text{-DP}} \max_{p \in \mathcal{V}} \mathbb{E}_{X \sim p} [\ell(\theta(p_{\hat{E}(X)}), \theta(p))] \quad (20)$$

For any (ε, δ) -DP index estimator \hat{E} , and by (6),

$$\max_{p \in \mathcal{V}} \mathbb{E}_{X \sim p} [\ell(\theta(p_{\hat{E}}), \theta(p))] \geq \frac{1}{|\mathcal{E}_k|} \sum_{e \in \mathcal{E}_k} \mathbb{E}_{X \sim p_e} [\ell(\theta(p_{\hat{E}}), \theta(p_e))] \geq \frac{2\tau}{|\mathcal{E}_k|} \sum_{i=1}^k \sum_{e \in \mathcal{E}_k} \Pr(\hat{E}_i \neq e_i | E = e).$$

For each i , we divide $\mathcal{E}_k = \{\pm 1\}^k$ into two sets according to the value of i -th position,

$$\begin{aligned} \max_{p \in \mathcal{V}} \mathbb{E}_{X \sim p} [\ell(\theta(p_{\hat{E}}), \theta(p))] &\geq \frac{2\tau}{|\mathcal{E}_k|} \sum_{i=1}^k \left[\sum_{e: e_i=1} \Pr(\hat{E}_i \neq 1 | E = e) + \sum_{e: e_i=-1} \Pr(\hat{E}_i \neq -1 | E = e) \right] \\ &= \tau \cdot \sum_{i=1}^k \left(\Pr_{X \sim p_{+i}}(\hat{E}_i \neq 1) + \Pr_{X \sim p_{-i}}(\hat{E}_i \neq -1) \right) \\ &\geq \tau \cdot \sum_{i=1}^k \min_{\phi_i: \phi_i \text{ is DP}} \left(\Pr_{X \sim p_{+i}}(\phi_i(X) \neq 1) + \Pr_{X \sim p_{-i}}(\phi_i(X) \neq -1) \right). \end{aligned}$$

Combining with (20), we have

$$R(\mathcal{P}, \ell, \varepsilon, \delta) \geq R(\mathcal{V}, \ell, \varepsilon, \delta) \geq \frac{\tau}{2} \cdot \sum_{i=1}^k \min_{\phi_i: \phi_i \text{ is DP}} \left(\Pr_{X \sim p_{+i}}(\phi_i(X) \neq 1) + \Pr_{X \sim p_{-i}}(\phi_i(X) \neq -1) \right),$$

proving the first part.

For the second part. Note that for each $i \in [k]$, the summand above is the error probability of hypothesis testing between the mixture distributions p_{+i} and p_{-i} . Hence, using Theorem 1,

$$R(\mathcal{P}, \ell, \varepsilon, \delta) \geq \frac{k\tau}{2} \cdot \left(0.9e^{-10\varepsilon D} - 10D\delta \right).$$

■

6. Proofs of existence of codes (Lemma 6 and Lemma 7)

Proof [Proof of Lemma 6] This proof is a standard argument for Gilbert-Varshamov bound applied to constant weight codes. We use the following version (Theorem 7 in [Graham and Sloane \(1980\)](#)).

Lemma 21 *There exists a length- k constant weight binary code \mathcal{C} with weight l and minimum Hamming distance 2δ , with*

$$|\mathcal{C}| \geq \frac{\binom{k}{l}}{\sum_{i=0}^{\delta} \binom{l}{i} \binom{k-l}{i}}.$$

Applying this Lemma with $2\delta = \frac{l}{4}$, we have

$$\begin{aligned} |\mathcal{C}| &\geq \frac{\binom{k}{l}}{\sum_{j=0}^{l/8} \binom{l}{j} \cdot \binom{k-l}{j}} \geq \frac{\binom{k}{l}}{\frac{l}{8} \cdot \binom{l}{\frac{l}{8}} \cdot \binom{k-l}{\frac{l}{8}}} = \frac{1}{\frac{l}{8} \cdot \binom{l}{\frac{l}{8}}} \cdot \prod_{i=0}^{\frac{7l}{8}-1} \frac{k - \frac{l}{8} - i}{l - i} \\ &\geq \frac{2\sqrt{7}\pi}{e} \cdot (0.59)^{\frac{7l}{8}} \cdot \left(\frac{k - \frac{l}{8}}{l} \right)^{\frac{7l}{8}} \\ &\geq \left(\frac{k}{2^{7/8} l} \right)^{\frac{7l}{8}}, \end{aligned} \tag{21}$$

In (21), we note that $\frac{k-\frac{l}{8}-i}{l-i}$ is monotonically increasing as i increases. And the first part is obtained by the Stirling's approximation $\sqrt{2\pi} \cdot l^{l+\frac{1}{2}} \cdot e^{-l} \leq l! \leq e \cdot l^{l+\frac{1}{2}} \cdot e^{-l}$ and the fact that $1.1^l \geq \sqrt{l}$ when $l \geq 20$. The last inequality comes from $l \leq k/2$ and $15/16 \times 0.59 > 1/2^{7/8}$. ■

Proof [Proof of Lemma 7] By the Gilbert-Varshamov bound (Lemma 21),

$$|\mathcal{H}| \geq \frac{h^d}{\sum_{j=0}^{\frac{d}{2}-1} \binom{d}{j} (h-1)^j} \geq \frac{h^d}{\frac{d}{2} \cdot \binom{d}{\frac{d}{2}} \cdot h^{\frac{d}{2}}} \geq \frac{h^{\frac{d}{2}}}{d \cdot 2^d} \geq \left(\frac{h}{16}\right)^{\frac{d}{2}}.$$

■

Acknowledgments

The authors thank Gautam Kamath and Ananda Theertha Suresh for their thoughtful comments and insights that helped improve the paper.

References

- Jayadev Acharya and Ziteng Sun. Communication complexity in locally private distribution estimation and heavy hitters. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- Jayadev Acharya, Gautam Kamath, Ziteng Sun, and Huanyu Zhang. INSPECTRE: Privately estimating the unseen. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 30–39, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018a. PMLR.
- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. In *Advances in Neural Information Processing Systems*, pages 6879–6891, 2018b.
- Jayadev Acharya, Clément Canonne, Cody Freitag, and Himanshu Tyagi. Test without trust: Optimal locally private distribution testing. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2067–2076, 2019a.
- Jayadev Acharya, Clément L Canonne, and Himanshu Tyagi. Inference under information constraints: Lower bounds from chi-square contraction. *Proceedings of Machine Learning Research* vol, 99: 1–15, 2019b.
- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Hadamard response: Estimating distributions privately, efficiently, and with little communication. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1120–1129, 2019c.
- Jayadev Acharya, Clément L Canonne, Yuhan Liu, Ziteng Sun, and Himanshu Tyagi. Interactive inference under information constraints. *arXiv preprint arXiv:2007.10976*, 2020.

- Maryam Aliakbarpour, Ilias Diakonikolas, and Ronitt Rubinfeld. Differentially private identity and equivalence testing of discrete distributions. In *Proceedings of the 35th International Conference on Machine Learning*, pages 169–178, 2018.
- Hassan Ashtiani, Shai Ben-David, Nicholas Harvey, Christopher Liaw, Abbas Mehrabian, and Yaniv Plan. Nearly tight sample complexity bounds for learning mixtures of gaussians via sample compression schemes. In *Advances in Neural Information Processing Systems*, pages 3412–3421, 2018.
- Patrice Assouad. Deux remarques sur l’estimation. *Comptes rendus des séances de l’Académie des sciences. Série I, Mathématique*, 296(23):1021–1024, 1983.
- Rina Foygel Barber and John C Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv preprint arXiv:1412.4451*, 2014.
- Raef Bassily. Linear queries estimation with local differential privacy. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 721–729, 2019.
- Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science, FOCS ’14*, pages 464–473, Washington, DC, USA, 2014. IEEE Computer Society.
- Amos Beimel, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. In *Theory of Cryptography Conference*, pages 437–454. Springer, 2010.
- Peter J Bickel and Yaacov Ritov. Estimating integrated squared density derivatives: sharp best order of convergence estimates. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 381–393, 1988.
- Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: The SuLQ framework. In *Proceedings of the 24th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS ’05*, pages 128–138, New York, NY, USA, 2005. ACM.
- Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *Journal of the ACM (JACM)*, 60(2):12, 2013.
- Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Differentially private release and learning of threshold functions. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science, FOCS ’15*, pages 634–649, Washington, DC, USA, 2015. IEEE Computer Society.
- Mark Bun, Thomas Steinke, and Jonathan Ullman. Make up your mind: The price of online queries in differential privacy. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1306–1325. SIAM, 2017.
- Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. *SIAM Journal on Computing*, 47(5):1888–1938, 2018.

- Mark Bun, Gautam Kamath, Thomas Steinke, and Zhiwei Steven Wu. Private hypothesis selection. In *Advances in Neural Information Processing Systems*, 2019.
- Bryan Cai, Constantinos Daskalakis, and Gautam Kamath. Priv’it: Private and sample efficient identity testing. In *Proceedings of the 34th International Conference on Machine Learning, ICML ’17*, pages 635–644. JMLR, Inc., 2017.
- T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *arXiv preprint arXiv:1902.04495*, 2019.
- Clément L. Canonne. A survey on distribution testing: Your data is big, but is it blue? *Electronic Colloquium on Computational Complexity (ECCC)*, 22(63):63, 2015.
- Clément L Canonne, Gautam Kamath, Audra McMillan, Adam Smith, and Jonathan Ullman. The structure of optimal private tests for simple hypotheses. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 310–321. ACM, 2019.
- Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (2. ed.)*. Wiley, 2006.
- Tore Dalenius. Towards a methodology for statistical disclosure control. *Statistisk Tidskrift*, 15: 429–444, 1977.
- Constantinos Daskalakis and Gautam Kamath. Faster and sample near-optimal algorithms for proper learning mixtures of gaussians. In *Proceedings of the 27th Annual Conference on Learning Theory, COLT ’14*, pages 1183–1213, 2014.
- Frank den Hollander. Probability theory: The coupling method. *Lecture notes available online (<http://websites.math.leidenuniv.nl/probability/lecturenotes/CouplingLectures.pdf>)*, 2012.
- Luc Devroye. *A course in density estimation*. Birkhauser Boston Inc., 1987.
- Ilias Diakonikolas, Moritz Hardt, and Ludwig Schmidt. Differentially private learning of structured discrete distributions. In *Advances in Neural Information Processing Systems 28, NIPS ’15*, pages 2566–2574. Curran Associates, Inc., 2015.
- Differential Privacy Team, Apple. Learning with privacy at scale. <https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appliedifferentialprivacysystem.pdf>, December 2017.
- Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems*, pages 3571–3580, 2017.
- Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS ’03*, pages 202–210, New York, NY, USA, 2003. ACM.

- John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '13, pages 429–438. IEEE, 2013.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC '06, pages 265–284, Berlin, Heidelberg, 2006. Springer.
- Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '10, pages 51–60, Washington, DC, USA, 2010. IEEE Computer Society.
- Cynthia Dwork, Adam Smith, Thomas Steinke, Jonathan Ullman, and Salil Vadhan. Robust traceability from trace amounts. In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '15, pages 650–669, Washington, DC, USA, 2015. IEEE Computer Society.
- Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM Conference on Computer and Communications Security*, CCS '14, pages 1054–1067, New York, NY, USA, 2014. ACM.
- Marco Gaboardi and Ryan Rogers. Local private hypothesis testing: Chi-square tests. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- Ron Graham and Neil Sloane. Lower bounds for constant weight codes. *IEEE Transactions on Information Theory*, 26(1):37–43, 1980.
- Te Sun Han and Sergio Verdú. Generalizing the fano inequality. *IEEE Transactions on Information Theory*, 40(4):1247–1251, 1994.
- Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, pages 705–714. ACM, 2010.
- Il'dar Abdulovich Ibragimov and Rafail Zalmanovich Has' Minskii. *Statistical estimation: asymptotic theory*, volume 16. Springer Science & Business Media, 2013.
- Jiantao Jiao, Kartik Venkat, Yanjun Han, and Tsachy Weissman. Minimax estimation of functionals of discrete distributions. *IEEE Transactions on Information Theory*, 61(5):2835–2885, May 2015.
- Peter Kairouz, Keith Bonawitz, and Daniel Ramage. Discrete distribution estimation under local privacy. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pages 2436–2444, 2016.
- Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017.
- Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Proceedings of the 32nd Annual Conference on Learning Theory*, 2019.

- Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Conference on Learning Theory*, pages 2204–2235. PMLR, 2020.
- Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. In *Proceedings of the 9th Conference on Innovations in Theoretical Computer Science*, ITCS '18, pages 44:1–44:9. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- Lucien Le Cam. Convergence of estimates under dimensionality restrictions. *The Annals of Statistics*, 1(1):38–53, 1973.
- Lucien Marie Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer series in statistics. Springer, 1986.
- Jing Lei. Differentially private m-estimators. In *Advances in Neural Information Processing Systems*, pages 361–369, 2011.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science*, pages 94–103. IEEE, 2007.
- Liam Paninski. A coincidence-based test for uniformity given very sparsely sampled discrete data. *IEEE Transactions on Information Theory*, 54(10):4750–4755, 2008.
- Yury Polyanskiy and Yihong Wu. Dualizing le cam’s method, with applications to estimating the unseens. *arXiv preprint arXiv:1902.05616*, 2019.
- Kolyan Ray and Johannes Schmidt-Hieber. The le cam distance between density estimation, poisson processes and gaussian white noise. *Mathematical Statistics and Learning*, 1(2):101–170, 2018.
- Jonathan Scarlett and Volkan Cevher. An introductory guide to fano’s inequality with applications in statistical estimation. *CoRR*, abs/1901.00555, 2019.
- Or Sheffet. Differentially private ordinary least squares. In *Proceedings of the 34th International Conference on Machine Learning*, ICML '17, pages 3105–3114. JMLR, Inc., 2017.
- Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *Proceedings of the 28th Annual Conference on Learning Theory*, COLT '15, pages 1588–1628, 2015.
- Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *The Journal of Privacy and Confidentiality*, 7(2):3–22, 2017a.
- Thomas Steinke and Jonathan Ullman. Tight lower bounds for differentially private selection. In *Proceedings of the 58th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '17, pages 552–563. IEEE Computer Society, 2017b.
- Ananda Theertha Suresh, Alon Orlitsky, Jayadev Acharya, and Ashkan Jafarpour. Near-optimal-sample estimators for spherical Gaussian mixtures. In *Advances in Neural Information Processing Systems 27*, NIPS '14, pages 1395–1403. Curran Associates, Inc., 2014.

- Salil Vadhan. The complexity of differential privacy. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich*, chapter 7, pages 347–450. Springer International Publishing AG, Cham, Switzerland, 2017.
- S. Wang, L. Huang, Y. Nie, X. Zhang, P. Wang, H. Xu, and W. Yang. Local differential private data aggregation for discrete distribution estimation. *IEEE Transactions on Parallel and Distributed Systems*, 30(9):2046–2059, 2019. doi: 10.1109/TPDS.2019.2899097.
- Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- Larry Wasserman and Shuheng Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389, 2010.
- Yihong Wu and Pengkun Yang. Minimax rates of entropy estimation on large alphabets via best polynomial approximation. *IEEE Trans. Information Theory*, 62(6):3702–3720, 2016.
- Min Ye and Alexander Barg. Optimal schemes for discrete distribution estimation under locally differential privacy. *IEEE Transactions on Information Theory*, 64:5662–5676, 2018.
- Bin Yu. Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pages 423–435. Springer New York, 1997.
- Huanyu Zhang, Gautam Kamath, Janardhan Kulkarni, and Steven Wu. Privately learning markov random fields. In *International Conference on Machine Learning*, pages 11129–11140. PMLR, 2020.