

On the Sample Complexity of Privately Learning Unbounded High-Dimensional Gaussians

Ishaq Aden-Ali

Department of Computing and Software, McMaster University

ADENALI@MCMASTER.CA

Hassan Ashtiani

Department of Computing and Software, McMaster University

ZOKAEIAM@MCMASTER.CA

Gautam Kamath

Cheriton School of Computer Science, University of Waterloo.

G@CSAIL.MIT.EDU

Editors: Vitaly Feldman, Katrina Ligett and Sivan Sabato

Abstract

We provide sample complexity upper bounds for agnostically learning multivariate Gaussians under the constraint of approximate differential privacy. These are the first finite sample upper bounds for general Gaussians which do not impose restrictions on the parameters of the distribution. Our bounds are near-optimal in the case when the covariance is known to be the identity, and conjectured to be near-optimal in the general case. From a technical standpoint, we provide analytic tools for arguing the existence of global “locally small” covers from local covers of the space. These are exploited using modifications of recent techniques for differentially private hypothesis selection. Our techniques may prove useful for privately learning other distribution classes which do not possess a finite cover.

1. Introduction

Given samples from a distribution P , can we estimate the underlying distribution? This problem has a long and rich history, culminating in a mature understanding for many settings of interest. However, in many cases the dataset may consist of sensitive data belonging to individuals, and naive execution of classic methods may inadvertently result in private information leakage. See, for instance, privacy attacks described in such estimation settings including [Dinur and Nissim \(2003\)](#); [Homer et al. \(2008\)](#); [Bun et al. \(2014\)](#); [Dwork et al. \(2015\)](#); [Shokri et al. \(2017\)](#), and the survey [Dwork et al. \(2017\)](#).

To address concerns of this nature, in 2006, Dwork, McSherry, Nissim, and Smith introduced the celebrated notion of differential privacy (DP) [Dwork et al. \(2006b\)](#), which provides a strong standard for data privacy. It ensures that no single data point has significant influence on the output of the algorithm, thus masking the contribution of individuals in the dataset. Differential privacy has seen practical adoption in many organizations, including Apple [Differential Privacy Team, Apple \(2017\)](#), Google [Erlingsson et al. \(2014\)](#); [Bittau et al. \(2017\)](#), Microsoft [Ding et al. \(2017\)](#), and the US Census Bureau [Dajani et al. \(2017\)](#). At this point, there is a rich body of literature, giving differentially private algorithms for a wide array of tasks.

There has recently been significant interest in distribution and parameter estimation under differential privacy (see [Section 1.1.2](#) for discussion of related work). Most relevant to our investigation is the work of Bun, Kamath, Steinke, and Wu [Bun et al. \(2019\)](#), which provides a generic framework that, given a cover for a class of distributions, describes a private algorithm for learning said

class with sample complexity logarithmic in the size of the cover. There, the privacy guarantee is the strongest notion of *pure* $(\epsilon, 0)$ -differential privacy.

An obvious drawback of this approach is that it fails to provide sample complexity upper bounds for estimating classes of distributions which do not possess a finite cover. The canonical example is the set of all Gaussian distributions. It turns out that this is inherently impossible – “packing lower bounds” imply that no finite sample algorithm exists for such cases under pure differential privacy. This theoretical limitation can have significant practical implications as well, as it forces the data analyst to choose between having good accuracy and preserving privacy. It turns out that, under pure differential privacy, the only way to avoid this issue is to assume the underlying distribution belongs to a more restricted class – such as Gaussian distributions with bounded mean and covariance.

On the other hand, stronger results are possible if one relaxes the privacy notion to the weaker guarantee of approximate differential privacy [Dwork et al. \(2006a\)](#). In particular, it is known that this relaxation permits “stability-based” approaches, which can avoid issues associated with infinite covers by pinpointing the area where “a lot of the data lies,” see, e.g., the classic example of the stability-based histogram [Korolova et al. \(2009\)](#); [Bun et al. \(2016\)](#).

Using such approaches, [Bun et al. \(2019\)](#) exploit the stability-based GAP-MAX algorithm of [Bun et al. \(2018\)](#), refining their framework to provide algorithms under approximate differential privacy. The caveat is that this time the approach requires construction of a more sophisticated object: a cover for the class which is in a certain technical sense “locally small.” Such locally small covers are much more difficult to construct and analyze. As such, [Bun et al. \(2019\)](#) only provide them for univariate Gaussian and multivariate Gaussians with identity covariance. As the most notable omission, they do not provide a locally small cover for general multivariate Gaussians. Indeed, for Gaussians with identity covariance, it is easy to reason about the local size of covers, as total variation distance between distributions corresponds to the ℓ_2 -distance between their means. However, when the covariance is not fixed, the total variation distance is characterized by the *Mahalanobis distance*, which has a significantly more sophisticated geometry. Analyzing these situations to show local smallness appears to be intractable using current analytic techniques, which involve explicitly constructing and analyzing a cover of the space. Given this challenge, up to now it has not been clear even whether a finite sample algorithm exists at all! And this is only for the fundamental case of Gaussians, raising the question of how one would even approach more complex classes of distributions.

1.1. Results and Techniques

We resolve these issues by providing a simpler method for proving existence of locally small covers. These lead to our main results, sample complexity upper bounds for semi-agnostically learning Gaussian distributions under approximate differential privacy.

Theorem 1 (Informal version of Theorem 40) *The sample complexity of semi-agnostically learning a d -dimensional Gaussian distribution to α -accuracy in total variation distance under (ϵ, δ) -differential privacy is*

$$\tilde{O} \left(\frac{d^2}{\alpha^2} + \frac{d^2}{\alpha\epsilon} + \frac{\log(1/\delta)}{\epsilon} \right).$$

This is the first sample complexity bound for privately learning a multivariate Gaussian with no conditions on the covariance matrix. The first and third terms are known to be tight, and there is strong evidence that the second is as well, see Section 1.1.1. The previous best algorithm was

that of [Kamath et al. \(2019a\)](#), which provided the stronger guarantee of concentrated differential privacy [Dwork and Rothblum \(2016\)](#); [Bun and Steinke \(2016\)](#) (which is intermediate to pure and approximate DP). However, it required the true covariance Σ to be bounded as $I \preceq \Sigma \preceq KI$ for some known parameter K , and the third term in the sample complexity is instead $O\left(\frac{d^{3/2} \log^{1/2} K}{\varepsilon}\right)$, which is prohibitive for large (or unknown) K . In contrast, our result holds for unrestricted Gaussian distributions.

We also provide a better upper bound for the case when the covariance matrix is known.

Theorem 2 (Informal version of Theorem 38) *The sample complexity of semi-agnostically learning a d -dimensional Gaussian distribution with known covariance to α -accuracy in total variation distance under (ε, δ) -differential privacy is*

$$\tilde{O}\left(\frac{d}{\alpha^2} + \frac{d}{\alpha\varepsilon} + \frac{\log(1/\delta)}{\varepsilon}\right).$$

This is the first bound which achieves a near-optimal dependence simultaneously on all parameters, see Section 1.1.1. In particular, it improves upon previous results in which the third term is replaced by $O\left(\frac{\log(1/\delta)}{\alpha\varepsilon}\right)$ [Bun et al. \(2019\)](#) or $O\left(\frac{\sqrt{d} \log^{3/2}(1/\delta)}{\varepsilon}\right)$ [Karwa and Vadhan \(2018\)](#); [Kamath et al. \(2019a\)](#); [Bun et al. \(2019\)](#).

While we apply our approach to multivariate Gaussian estimation, it should more broadly apply to other classes of distributions with no finite-sized cover.

As mentioned before, we build upon the approach of Bun, Kamath, Steinke, and Wu [Bun et al. \(2019\)](#) to provide methods better suited for estimation under the constraint of approximate differential privacy. Their work focuses primarily on pure DP distribution estimation for classes of distributions with a finite cover. Specifically, given a class of distributions with an α -cover of size \mathcal{C}_α , they give a pure DP algorithm for learning said class in total variation distance with sample complexity $O(\log |\mathcal{C}_\alpha|)$. Naturally, this gives vacuous bounds for classes with an infinite cover – indeed, packing lower bounds show that this is inherent under pure DP [Hardt and Talwar \(2010\)](#); [Beimel et al. \(2014\)](#); [Bun et al. \(2019\)](#). To avoid these lower bounds, they show that learning is still possible if one relaxes to approximate DP and considers a “locally small” cover: one that has at most k elements which are within an $O(\alpha)$ -total variation distance ball of any element in the set. The sample complexity of the resulting method does not depend on $|\mathcal{C}_\alpha|$, and instead we pay logarithmically in the parameter k . They apply this framework to provide algorithms for estimating general univariate Gaussians, and multivariate Gaussians with identity covariance. However, their arguments construct explicit covers for these cases, and it appears difficult to construct and analyze covers in situations with a rich geometric structure, such as multivariate Gaussians. Indeed, it seems difficult in these settings to reason that a set is simultaneously a cover (i.e., every distribution in the class has a close element) and locally small (i.e., every distribution does not have *too many* close elements).

We avoid this tension by taking a myopic view: in Lemma 29, we show that if we can construct a cover with few elements for the neighbourhood of each *individual* distribution, then there exists a locally small cover for the *entire space*. This makes it significantly easier to reason about locally small covers, as we only have to consider covering a single distribution at a time, and we do not have to reason about how the elements that cover each distribution overlap with each other. For example: to cover the neighbourhood of a single Gaussian with (full rank) covariance Σ , we can transform

the covariance to the identity by multiplying by $\Sigma^{-1/2}$, cover the neighbourhood of $N(0, I)$ (which is easier), and transform the cover back to the original domain. This is far simpler than trying to understand how to simultaneously cover multiple Gaussians with differently shaped covariance matrices in a locally small manner. Our results for covering are presented in Section 3.

We then go on to apply these locally small covers to derive learning sample complexity upper bounds in Section 4. As mentioned before, this is done in Bun et al. (2019), though we refine their method to achieve stronger bounds. While this refinement is simple, we believe it to be important both technically (as it allows us to achieve likely near-optimal sample complexities) and conceptually (as we believe it clearly identifies what the “hard part” of the problem is). To elaborate, our approach can be divided into two steps;

1. **Coarse Estimation.** Find any distribution which is 0.99-close to the true distribution, using the approximate DP GAP-MAX algorithm in Bun et al. (2019).
2. **Fine Estimation.** Generate an $O(\alpha)$ -cover around the distribution from the previous step, and run the pure DP private hypothesis selection algorithm in Bun et al. (2019).

We are not the first to use this type of two-step approach, as such decomposition has been previously applied, e.g., Karwa and Vadhan (2018); Kamath et al. (2019a, 2020). However, it was not applied in the context of the GAP-MAX algorithm in Bun et al. (2019), preventing them from getting the right dependencies on all parameters – in particular, it was not clear how to disentangle the dependencies on $\log(1/\delta)$ and $1/\alpha$ using their method directly.

Other beneficial features of this two-step approach, which have also been exploited in the past, include its modularity and the fact that the first and second steps involve qualitatively different privacy guarantees. However, we additionally comment how coarse an estimate required in the first step – while the description above states that we require a 0.99-close distribution, we may actually only need one with total variation distance bounded by $1 - \zeta$, where ζ may be exponentially small in the parameters of the problem! See Remark 35. We hope that shining a light on this somewhat unconventional regime for private distribution estimation, which only requires a “whiff” of the true distribution, will inspire further investigation.

As a final contribution, in Section 5, we revisit the generic private hypothesis selection problem. The main result of Bun et al. (2019) is an algorithm for this problem which requires knowledge of the distance to the best hypothesis. They then wrap this algorithm in another procedure which “guesses” the distance to the best hypothesis, resulting in a semi-agnostic algorithm. However, this loses large factors in the agnostic guarantee and is rather indirect. We instead analyze the privatization of a different algorithm, the minimum distance estimate, which gives a semi-agnostic algorithm directly, with an optimal agnostic constant (i.e., providing a tight factor of 3 Devroye and Lugosi (2001)). In our opinion, the algorithm and proof are even simpler than the non-agnostic algorithm of Bun et al. (2019).

Theorem 3 (Informal version of Theorem 27) *There exists an ε -differentially private algorithm for semi-agnostic hypothesis selection from a set of m distributions $\mathcal{H} = \{H_1, \dots, H_m\}$. The algorithm requires n samples from some distribution P and returns a distribution $\hat{H} \in \mathcal{H}$ where $\text{TV}(P, \hat{H}) \leq 3 \cdot \text{OPT} + \alpha$, where*

$$n = O\left(\frac{\log m}{\alpha^2} + \frac{\log m}{\alpha\varepsilon}\right).$$

1.1.1. COMPARISON WITH LOWER BOUNDS

It is folklore that the non-private sample complexity of estimating a single d -dimensional Gaussian to accuracy α in total variation distance is $\Theta(d^2/\alpha^2)$, or, in the case when the covariance is the identity, $\Theta(d/\alpha^2)$. Therefore the leading terms in the sample complexity bounds of Theorems 1 and 2 are tight.

Lower bounds for private statistical estimation are comparatively less explored. Karwa and Vadhan [Karwa and Vadhan \(2018\)](#) showed a lower bound of $\Omega(\log(1/\delta)/\varepsilon)$, even for the simple case of estimating the mean of a univariate Gaussian with known variance, thus matching the third terms in Theorems 1 and 2.

However, approximate DP lower bounds in the multivariate setting are notoriously hard to come by, with the predominant technique being the ‘‘fingerprinting’’ approach [Bun et al. \(2014\)](#); [Steinke and Ullman \(2015\)](#); [Dwork et al. \(2015\)](#); [Kamath et al. \(2019a\)](#); [Cai et al. \(2019\)](#). Using this technique, Kamath, Li, Singhal, and Ullman [Kamath et al. \(2019a\)](#) show a lower bound of $\tilde{\Omega}(d/\alpha\varepsilon)$ for Gaussian estimation with identity covariance, thus nearly-matching the second and final term in Theorem 2. We note that this sample complexity does not change when we convert to the stronger notion of *pure* differential privacy [Kamath et al. \(2019a\)](#); [Bun et al. \(2019\)](#).

[Kamath et al. \(2019a\)](#) also proves a lower bound of $\Omega(d^2/\alpha\varepsilon)$ for general Gaussian estimation under *pure* differential privacy. Using the aforementioned invariance of the complexity of estimation with identity covariance under pure and approximate DP, we take this as strong evidence that there exists a lower bound of $\Omega(d^2/\alpha\varepsilon)$ for estimation of general Gaussians under approximate DP as well.

1.1.2. ADDITIONAL RELATED WORK

The most relevant works are those on private distribution and parameter estimation, particularly in multivariate settings [Nissim et al. \(2007\)](#); [Bun et al. \(2014\)](#); [Steinke and Ullman \(2017a,b\)](#); [Dwork et al. \(2015\)](#); [Bun et al. \(2017\)](#); [Karwa and Vadhan \(2018\)](#); [Kamath et al. \(2019a\)](#); [Cai et al. \(2019\)](#); [Bun et al. \(2019\)](#); [Acharya et al. \(2020\)](#); [Kamath et al. \(2020\)](#); [Biswas et al. \(2020\)](#). While many of these focus on settings with parameters bounded by some constant, some pay particular attention to the cost in terms of this bound, including [Karwa and Vadhan \(2018\)](#); [Kamath et al. \(2019a\)](#); [Bun et al. \(2019\)](#); [Biswas et al. \(2020\)](#); [Du et al. \(2020\)](#). Private ball-finding algorithms [Nissim et al. \(2016\)](#); [Nissim and Stemmer \(2018\)](#) provide approximate DP approaches for finding small ℓ_2 -balls containing many points, which can be applied in sample-and-aggregate settings [Kamath et al. \(2019b\)](#). However, these too appear to be unable to exploit the sophisticated geometry that arises with non-spherical covariance matrices of general Gaussians. Broadly speaking, we are not aware of any existing approach that is able to entirely avoid dependence on bounds on the parameters. Other works on differentially private estimation include [Dwork and Lei \(2009\)](#); [Smith \(2011\)](#); [Diakonikolas et al. \(2015\)](#); [Bun and Steinke \(2019\)](#); [Zhang et al. \(2020\)](#). See [Kamath and Ullman \(2020\)](#) for more coverage of recent works in private statistics.

The work of Bun, Kamath, Steinke, and Wu [Bun et al. \(2019\)](#) is built upon classic results in hypothesis selection, combined with the exponential mechanism [McSherry and Talwar \(2007\)](#). The underlying non-private approach was pioneered by Yatracos [Yatracos \(1985\)](#), and refined in subsequent work by Devroye and Lugosi [Devroye and Lugosi \(1996, 1997, 2001\)](#). After this, additional considerations have been taken into account, such as computation, approximation factor, robustness, and more [Mahalanabis and Stefankovic \(2008\)](#); [Daskalakis et al. \(2012\)](#); [Daskalakis](#)

and Kamath (2014); Suresh et al. (2014); Acharya et al. (2014); Diakonikolas et al. (2016); Ashtiani et al. (2018b,a); Acharya et al. (2018); Bousquet et al. (2019); Aden-Ali and Ashtiani (2020). Notably, these primitives have also been translated to the more restrictive setting of *local* differential privacy Gopi et al. (2020). Similar techniques have also been exploited in a federated setting Liu et al. (2020).

2. Preliminaries

2.1. Notation

For any $m \in \mathbb{N}$, $[m]$ denotes the set $\{1, 2, \dots, m\}$. Let $X \sim P$ denote a random variable X sampled from distribution P . Let $(X_i)_{i=1}^m \sim P^m$ denote an i.i.d. random sample of size m from distribution P . For a distribution Q over a domain \mathcal{X} and a set $A \subseteq \mathcal{X}$, we define $Q(A)$ as the probability Q assigns to the event A .

For a positive integer d let $\mathbb{S}_d \subset \mathbb{R}^{d \times d}$ be the set of all d -by- d positive semi-definite real matrices. For a matrix $A \in \mathbb{R}^{m \times n}$, define $\|A\|_{1,1} = \sum_{i=1}^m \sum_{j=1}^n |A_{ij}|$ and $\|A\|_{\infty,\infty} = \max_{i,j} |A_{ij}|$. The determinant of a square matrix A is given by $\det(A)$.

Definition 4 A d -dimensional Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ with mean $\mu \in \mathbb{R}^d$ and covariance matrix $\Sigma \in \mathbb{S}_d$ is a distribution with density function:

$$p(x) = \frac{\exp\left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu)\right)}{\sqrt{(2\pi)^d \cdot \det(\Sigma)}}.$$

We define the set of d -dimensional *location Gaussians* as $\mathcal{G}_d^L := \{\mathcal{N}(\mu, I) : \mu \in \mathbb{R}^d\}$, and the set of d -dimensional *scale Gaussians* as $\mathcal{G}_d^S := \{\mathcal{N}(0, \Sigma) : \Sigma \in \mathbb{S}_d\}$. Define the set of (all) Gaussians as $\mathcal{G}_d := \{\mathcal{N}(\mu, \Sigma) : \mu \in \mathbb{R}^d, \Sigma \in \mathbb{S}_d\}$. A useful property of Gaussian distributions is that any linear transformation of a Gaussian random vector is also a Gaussian random vector. In particular, if $X \sim \mathcal{N}(\mu, \Sigma)$ is a d -dimensional Gaussian random vector and A and b are a d -dimensional square matrix and vector respectively, it follows that

$$AX + b \sim \mathcal{N}(A\mu + b, A\Sigma A^T). \tag{1}$$

2.2. Distribution Learning

A *distribution learning method* is an algorithm that, given a sequence of i.i.d. samples from a distribution P , outputs a distribution \hat{H} as an estimate of P . The focus of this paper is on absolutely continuous probability distributions (distributions that have a density with respect to the Lebesgue measure), so we will refer to a probability distribution and its probability density function interchangeably. The specific measure of “closeness” between distributions that we use is the *total variation distance*:

Definition 5 Let P and Q be two probability distributions defined over \mathcal{X} and let Ω be the Borel sigma-algebra on \mathcal{X} . The total variation distance between P and Q is defined as

$$\text{TV}(P, Q) = \sup_{S \in \Omega} |P(S) - Q(S)| = \frac{1}{2} \int_{x \in \mathcal{X}} |P(x) - Q(x)| dx = \frac{1}{2} \|P - Q\|_1 \in [0, 1].$$

Moreover, if \mathcal{H} is a set of distributions over a common domain, we define $\text{TV}(P, \mathcal{H}) = \inf_{H \in \mathcal{H}} \text{TV}(P, H)$.

Given $X \sim P$ and $Y \sim Q$, it is often useful for us to overload notation and define $\text{TV}(X, Y) = \text{TV}(P, Q)$. We say two distributions P and H are γ -close if $\text{TV}(P, H) \leq \gamma$. We also say a distribution P is γ -close to a set of distributions \mathcal{H} if $\min_{H \in \mathcal{H}} \text{TV}(P, H) \leq \gamma$. We can now formally define a distribution learner:

Definition 6 (Realizable PAC learner) *An algorithm is said to be a (realizable) PAC learner for a set of distributions \mathcal{H} with sample complexity $n_{\mathcal{H}}(\alpha, \beta)$ if given parameters $\alpha, \beta \in (0, 1)$ and any $P \in \mathcal{H}$, the algorithm takes as input α, β and a sequence of $n_{\mathcal{H}}(\alpha, \beta)$ i.i.d. samples from P , and outputs $\hat{H} \in \mathcal{H}$ such that $\text{TV}(P, \hat{H}) \leq \alpha$ with probability at least $1 - \beta$.¹*

The following two definitions handle the case when we have model misspecification: we receive samples from a distribution P which is not in the class \mathcal{H} . The difference between the two is that in the robust definition (Definition 7) the algorithm is provided with an upper bound on the distance between P and \mathcal{H} , while it is not in the agnostic setting (Definition 8).

Definition 7 ((ξ, C) -robust PAC learner) *An algorithm is said to be a (ξ, C) -robust PAC learner for a set of distributions \mathcal{H} with sample complexity $\tilde{n}_{\mathcal{H}}^C(\alpha, \beta)$ if given parameters $\alpha, \beta, \xi \in (0, 1)$ and any distribution P such that $\text{TV}(P, \mathcal{H}) \leq \xi$, the algorithm takes as input α, β, ξ and a sequence of $\tilde{n}_{\mathcal{H}}^C(\alpha, \beta)$ i.i.d. samples from P , and outputs $\hat{H} \in \mathcal{H}$ such that $\text{TV}(H^*, \hat{H}) \leq C \cdot \xi + \alpha$ with probability at least $1 - \beta$.*

Definition 8 (C -agnostic PAC learner) *An algorithm is said to be a C -agnostic PAC learner for a set of distributions \mathcal{H} with sample complexity $n_{\mathcal{H}}^C(\alpha, \beta)$ if for any $\alpha, \beta \in (0, 1)$ and distribution P such that $\text{TV}(P, \mathcal{H}) = \text{OPT}$, given α, β and a sequence of $n_{\mathcal{H}}^C(\alpha, \beta)$ i.i.d. samples from P , the algorithm outputs $\hat{H} \in \mathcal{H}$ such that $\text{TV}(P, \hat{H}) \leq C \cdot \text{OPT} + \alpha$ with probability at least $1 - \beta$. If $C = 1$, the algorithm is said to be agnostic.*

We will sometimes refer to a C -agnostic PAC learner as a *semi-agnostic* PAC learner for $C > 1$, as is standard in learning theory. A useful object for us to define is the total variation ball.

Definition 9 (TV ball) *The total variation ball of radius $\gamma \in [0, 1]$, centered at a distribution P with respect to a set of distributions \mathcal{H} , is the following subset of \mathcal{H} :*

$$\mathcal{B}(\gamma, P, \mathcal{H}) = \{H \in \mathcal{H} : \text{TV}(P, H) \leq \gamma\}.$$

In this paper we consider coverings and packings of sets of distributions with respect to the total variation distance.

Definition 10 (γ -covers and γ -packings) *For any $\gamma \in [0, 1]$ a γ -cover of a set of distributions \mathcal{H} is a set of distributions \mathcal{C}_{γ} , such that for every $H \in \mathcal{H}$, there exists some $P \in \mathcal{C}_{\gamma}$ such that $\text{TV}(P, H) \leq \gamma$.*

A γ -packing of a set of distributions \mathcal{H} is a set of distributions $\mathcal{P}_{\gamma} \subseteq \mathcal{H}$, such that for every pair of distributions $P, Q \in \mathcal{P}_{\gamma}$, we have that $\text{TV}(P, Q) \geq \gamma$.

1. The probability is over the random samples drawn from P^n and the randomness of the algorithm.

Definition 11 (γ -covering and γ -packing number) For any $\gamma \in [0, 1]$, the γ -covering number of a set of distributions \mathcal{H} , $N(\mathcal{H}, \gamma) := \min\{n \in \mathbb{N} : \exists \mathcal{C}_\gamma \text{ s.t. } |\mathcal{C}_\gamma| = n\}$, is the size of the smallest possible γ -covering of \mathcal{H} . Similarly, the γ -packing number of a set of distributions \mathcal{H} , $M(\mathcal{H}, \gamma) := \max\{n \in \mathbb{N} : \exists \mathcal{P}_\gamma \text{ s.t. } |\mathcal{P}_\gamma| = n\}$, is the size of the largest subset of \mathcal{H} that forms a packing for \mathcal{H} ,

The following is a well known relation between covers and packings of a set of distribution. We defer the proof to Section B.1.

Lemma 12 For a set of distributions \mathcal{H} with γ -covering number $M(\mathcal{H}, \gamma)$ and γ -packing number $N(\mathcal{H}, \gamma)$, the following holds:

$$M(\mathcal{H}, 2\gamma) \leq N(\mathcal{H}, \gamma) \leq M(\mathcal{H}, \gamma).$$

An important property of a set of distributions we will need to quantify in this paper is how small they are ‘‘locally’’. The following definition formalizes this:

Definition 13 ((k, γ) -locally small) Fix some $\gamma \in [0, 1]$. We say a set of distributions \mathcal{H} is (k, γ) -locally small if

$$\sup_{H' \in \mathcal{H}} |\mathcal{B}(\gamma, H', \mathcal{H})| \leq k,$$

for some $k \in \mathbb{N}$. If no such k exists, we say \mathcal{H} is not (k, γ) -locally small.

2.3. VC Dimension and Uniform Convergence

An important property of a set of binary functions is its Vapnik-Chervonenkis (VC) dimension, which has the following definition:

Definition 14 (VC dimension Vapnik and Chervonenkis (1971)) Let \mathcal{F} be a set of binary functions $f : \mathcal{X} \rightarrow \{0, 1\}$. The VC dimension of \mathcal{F} is defined to be the largest d such that there exist $x_1, \dots, x_d \in \mathcal{X}$ and $f_1, \dots, f_{2^d} \in \mathcal{F}$ such that for all $i, j \in [2^d]$ where $i < j$, there exists $k \in [d]$ such that $f_i(x_k) \neq f_j(x_k)$.

The most important application of the VC dimension is the following celebrated uniform convergence bound:

Theorem 15 (Uniform Convergence Talagrand (1994)) Let \mathcal{F} be a set of binary functions $f : \mathcal{X} \rightarrow \{0, 1\}$ with VC dimension d . For any distribution P defined on \mathcal{X} , we have

$$\Pr_{D \sim P^n} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{x \in D} f(x) - \mathbf{E}_{X \sim P} [f(X)] \right| \leq \alpha \right] \geq 1 - \beta,$$

whenever $n = O\left(\frac{d + \log(1/\beta)}{\alpha^2}\right)$.

We can define the VC dimension of a set of distributions \mathcal{H} by looking at the VC dimension of a set of binary functions that is defined with respect to \mathcal{H} . More precisely:

Definition 16 (VC dimension of a set of distributions) Let \mathcal{H} be a set of probability distributions on a space \mathcal{X} . Define the set of binary functions $\mathcal{F}(\mathcal{H}) = \{f_{H_i, H_j} : H_i, H_j \in \mathcal{H}\}$ where $\forall x \in \mathcal{X}$, $f_{H_i, H_j}(x) = 1 \iff H_i(x) > H_j(x)$. We define the VC dimension of \mathcal{H} to be the VC dimension of $\mathcal{F}(\mathcal{H})$.²

Lemma 17 The set of location Gaussians \mathcal{G}_d^L has VC dimension $d + 1$. Furthermore, the set of Gaussians \mathcal{G}_d has VC dimension $O(d^2)$.

Proof For location Gaussians, $\mathcal{F}(\mathcal{G}_d^L)$ corresponds to linear threshold functions (i.e., half-spaces), which have VC dimension $d + 1$. Similarly $\mathcal{F}(\mathcal{G}_d)$ corresponds to quadratic threshold functions, which have VC dimension $\binom{d+2}{2} = O(d^2)$ [Anthony \(1995\)](#). ■

2.4. Differential Privacy

Let $X^* = \cup_{i=1}^{\infty} X^i$ be the set of possible datasets. We say that two datasets $D, D' \in X^*$ are neighbours if D and D' differ by at most one data point. Informally, an algorithm that receives a dataset and outputs a value in \mathcal{R} is differentially private if it outputs similar values on (any) two neighboring datasets. Formally:

Definition 18 (Dwork et al. (2006b,a)) A randomized algorithm $T : X^* \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private if for all $n \geq 1$, for all neighbouring datasets $D, D' \in X^n$, and for all measurable subsets $S \subseteq \mathcal{R}$,

$$\Pr [T(D) \in S] \leq e^\epsilon \Pr [T(D') \in S] + \delta.$$

If $\delta = 0$, we say that T is ϵ -differentially private.

We will refer to ϵ -DP as *pure DP*, and (ϵ, δ) -DP for $\delta > 0$ as *approximate DP*. A fundamental building block of differential privacy is the exponential mechanism [McSherry and Talwar \(2007\)](#). It is used to privately select an approximate “best” candidate from a (finite) set of candidates. The quality of a candidate with respect to the dataset is measured by a score function. Let \mathcal{R} be the set of possible candidates. A score function $S : X^* \times \mathcal{R} \rightarrow \mathbb{R}$ maps each pair consisting of a dataset and a candidate to a real-valued score. The *exponential mechanism* \mathcal{M}_E takes as input a dataset D , a set of candidates \mathcal{R} , a score function S , a privacy parameter ϵ and outputs a candidate $r \in \mathcal{R}$ with probability proportional to $\exp\left(\frac{\epsilon S(D, r)}{2\Delta(S)}\right)$, where $\Delta(S)$ is the sensitivity of the score function which is defined as

$$\Delta(S) = \max_{r \in \mathcal{R}, D \sim D'} |S(D, r) - S(D', r)|.$$

Theorem 19 (McSherry and Talwar (2007)) For any dataset D , score function S and privacy parameter $\epsilon > 0$, the exponential mechanism $\mathcal{M}_E(D, S, \epsilon)$ is an ϵ -differentially private algorithm, and with probability at least $1 - \beta$, it selects an outcome $r \in \mathcal{R}$ such that

$$S(D, r) \geq \max_{r' \in \mathcal{R}} S(D, r') - \frac{2\Delta(S) \log(|\mathcal{R}|/\beta)}{\epsilon}.$$

2. To avoid measurability issues we assume the preimage of 0 is measurable for any $f \in \mathcal{F}(\mathcal{H})$.

One of the most useful properties of differentially private algorithms is that they can be composed adaptively while promising a graceful degradation of privacy.

Lemma 20 (Composition of DP) *If M is an adaptive composition of differentially private algorithms M_1, \dots, M_T , then if M_1, \dots, M_T are $(\varepsilon_1, \delta_1), \dots, (\varepsilon_T, \delta_T)$ -differentially private then M is $(\sum_t \varepsilon_t, \sum_t \delta_t)$ -differentially private.*

Another strength of differential privacy is that it is closed under post-processing:

Lemma 21 (Post Processing) *If $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ε, δ) -differentially private, and $P : \mathcal{Y} \rightarrow \mathcal{Z}$ is any randomized function, then the algorithm $P \circ M$ is (ε, δ) -differentially private.*

We now define (ε, δ) -DP learners.

Definition 22 ((ε, δ) -DP PAC learner)

An algorithm is said to be an (ε, δ) -DP PAC learner for a set of distributions \mathcal{H} with sample complexity $n_{\mathcal{H}}(\alpha, \beta, \varepsilon, \delta)$ if it is a PAC learner that satisfies (ε, δ) -differential privacy.

Definition 23 ((ε, δ) -DP (ξ, C) -robust PAC learner) *An algorithm is said to be an (ε, δ) -DP (ξ, C) -robust PAC learner for a set of distributions \mathcal{H} with sample complexity $\tilde{n}_{\mathcal{H}}^C(\alpha, \beta, \varepsilon, \delta)$ if it is a (ξ, C) -robust PAC learner that satisfies (ε, δ) -differential privacy.*

Definition 24 ((ε, δ) -DP C -agnostic PAC learner) *An algorithm is said to be an (ε, δ) -DP C -agnostic PAC learner for a set of distributions \mathcal{H} with sample complexity $n_{\mathcal{H}}^C(\alpha, \beta, \varepsilon, \delta)$ if it is a C -agnostic PAC learner that satisfies (ε, δ) -differential privacy.*

2.4.1. PRIVATE HYPOTHESIS SELECTION AND THE GAP-MAX ALGORITHM

The problem of *hypothesis selection* (sometimes called density estimation, the Le Cam-Birgé method, or the Scheffé estimator) is a classical approach for reducing estimation problems to pairwise comparisons. It provides a generic approach for converting a cover for a set of probability distributions into a learning algorithm, see [Devroye and Lugosi \(2001\)](#) for a reference.

[Bun et al. \(2019\)](#) translated these powerful tools to the differentially private setting, giving an ε -DP algorithm for hypothesis selection using the exponential mechanism with a carefully constructed score function. The following is a modified version where we decouple the accuracy parameter α from the robustness parameter ξ , and boost the success probability to be arbitrarily high. The proof follows immediately from the proof in [Bun et al. \(2019\)](#). For a set of distributions \mathcal{H} , we will denote H^* as the distribution in \mathcal{H} that is closest to the unknown distribution P .

Theorem 25 *Let $\mathcal{H} = \{H_1, \dots, H_m\}$ be a set of probability distributions, $\xi, \alpha, \beta, \varepsilon, \delta \in (0, 1)$ and $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{H}) \leq \xi$. $\text{PHS}(\xi, \alpha, \beta, \varepsilon, \mathcal{H}, D)$ is an ε -DP $(\xi, 3)$ -robust PAC learner with sample complexity*

$$\tilde{n}_{\mathcal{H}}^3(\alpha, \beta, \varepsilon, 0) = O\left(\frac{\log(m/\beta)}{\alpha^2} + \frac{\log(m/\beta)}{\alpha\varepsilon}\right).$$

Furthermore, when the algorithm succeeds it guarantees that $\text{TV}(\hat{H}, H^) \leq 2\xi + \alpha$.*

We note the guarantee that the algorithm gives with respect to H^* in the theorem statement for technical reasons that will become apparent in the proofs of Section 4. Unfortunately the result above requires the number of hypotheses to be finite. Using a uniform convergence argument together with a GAP-MAX algorithm [Bun et al. \(2018\)](#), [Bun, Kamath, Steinke and Wu \(2019\)](#) showed that it may also be possible to get a similar guarantee when the number of hypotheses is infinite, provided that we relax the notion of privacy to approximate differential privacy. The following is an alternate version of ([Bun et al., 2019](#), Theorem 4.1). Again, in this version we decouple the accuracy parameter α from the robustness parameter ξ . The proof follows directly from the proof of Theorem 4.1 in [Bun et al. \(2019\)](#).

Theorem 26 ((alternate) Theorem 4.1 [Bun et al. \(2019\)](#)) *Let \mathcal{H} be a set of probability distributions, $\xi, \alpha, \beta, \varepsilon, \delta \in (0, 1)$ and $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{H}) \leq \xi$. Furthermore, let d be the VC dimension of $\mathcal{F}(\mathcal{H})$ and assume $|\mathcal{B}(3\xi + \alpha, H^*, \mathcal{H})| \leq k$. $\text{GAP-MAX}(\alpha, \xi, \beta, \varepsilon, \delta, k, \mathcal{H}, D)$ is an (ε, δ) -DP $(\xi, 4)$ -robust PAC learner for \mathcal{H} with sample complexity*

$$\tilde{n}_{\mathcal{H}}^4(\alpha, \beta, \varepsilon, \delta) = O\left(\frac{d + \log(1/\beta)}{\alpha^2} + \frac{\log(k/\beta) + \min\{\log(\mathcal{H}), \log(1/\delta)\}}{\alpha\varepsilon}\right).$$

Furthermore, when the algorithm succeeds it guarantees that $\text{TV}(\hat{H}, H^*) \leq 3\xi + \alpha$.

Note that Theorem 26 requires knowledge of $|\mathcal{B}(3\xi + \alpha, H^*, \mathcal{H})|$, which we likely do not know a priori. We can bound this by finding an upper bound on the size of the largest total variation ball centered at any $H' \in \mathcal{H}$, i.e. $\max_{H' \in \mathcal{H}} |\mathcal{B}(3\xi + \alpha, H', \mathcal{H})| \leq k'$. This directly translates to showing \mathcal{H} is $(k', 3\xi + \alpha)$ -locally small.

This lays the foundation for the strategy used in [Bun et al. \(2019\)](#) to construct a private distribution learner for an infinite set of distributions \mathcal{H} : by using a $(k', 6\xi + \alpha)$ -locally small³ ξ -cover for \mathcal{H} as the input to the GAP-MAX algorithm, given the right amount of samples (which depends on k'), with high probability the algorithm outputs a distribution that is $(8\xi + \alpha)$ -close to P .

Unfortunately, the above algorithms are not semi-agnostic and require an upper bound on OPT via ξ . As a first attempt, [Bun et al. \(2019\)](#) give an ε -DP 9-agnostic PAC learner based on the Laplace mechanism. This algorithm – which we will refer to as Naïve-PHS – is similar to the PHS algorithm of Theorem 25. Unfortunately, the sample complexity of the Naïve-PHS algorithm is

$$n_{\mathcal{H}}^9(\alpha, \beta, \varepsilon, 0) = O\left(\frac{\log(|\mathcal{H}|/\beta)}{\alpha^2} + \frac{|\mathcal{H}|^2 \log(|\mathcal{H}|/\beta)}{\alpha\varepsilon}\right),$$

which is exponentially worse than the PHS algorithm. As we will discuss shortly, [Bun et al. \(2019\)](#) also show how to use this algorithm together with the PHS algorithm to get a ε -DP 18-agnostic PAC learner, at the cost of some poly-logarithmic factors. This leads to the natural question of whether there exists an ε -DP semi-agnostic learner which achieves the *same* sample complexity as the PHS algorithm with a comparable agnostic constant. We answer this question in the affirmative and prove the following result:

Theorem 27 *Let $\mathcal{H} = \{H_1, \dots, H_m\}$ be a set of probability distributions, $\alpha, \beta, \varepsilon \in (0, 1)$ and $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{H}) = \text{OPT}$. There exists an ε -DP 3-agnostic PAC learner with sample complexity*

$$n_{\mathcal{H}}^3(\alpha, \beta, \varepsilon, 0) = O\left(\frac{\log(m/\beta)}{\alpha^2} + \frac{\log(m/\beta)}{\alpha\varepsilon}\right).$$

3. Note that the guarantee we can get from any ξ -cover \mathcal{C}_{ξ} is $\text{TV}(H^*, \mathcal{C}_{\xi}) \leq \xi \implies \text{TV}(P, \mathcal{C}_{\xi}) \leq 2\xi$.

We defer discussing the details of this result and its proof to Section 5, however we note that the above result can only handle finite sets of distributions. Recall that while the GAP-MAX algorithm can handle infinite sets of distributions, it is not a semi-agnostic learner. Thus, a natural question is whether we can learn from a set of infinite distributions using some (ε, δ) -DP semi-agnostic PAC learner.

Fortunately, Bun et al. (2019) gave a simple procedure that takes an (ε, δ) -DP robust PAC learner and constructs an (ε, δ) -DP semi-agnostic PAC learner, at the cost of some low order poly-logarithmic factors in the sample complexity bounds, and an increase in the agnostic constant. We can thus use the GAP-MAX algorithm together with this procedure to get an (ε, δ) -DP semi-agnostic PAC learner given an infinite set of distributions. The procedure Bun et al. (2019) came up with works in the following way: run the (ε, δ) -DP robust PAC learner with (a small number of) different values for ξ to get a shortlist of candidates. Use the semi-agnostic NaïvePHS algorithm to select a good hypothesis from the short list. As we mentioned earlier, the guarantee of this approach (Theorem 3.4 in Bun et al. (2019)) is stated specifically in terms of converting the PHS algorithm from Theorem 25 into an ε -DP semi-agnostic PAC learner, however it can be immediately generalized to construct (ε, δ) -DP semi-agnostic PAC learners given any (ε, δ) -DP robust PAC learner. Furthermore, we can replace the Naïve-PHS algorithm with the sample efficient algorithm from Theorem 27 to reduce the agnostic constant, and also remove some logarithmic factors in the sample complexity bound. This yields the following result:

Lemma 28 *Let $\alpha, \beta, \varepsilon, \delta \in (0, 1)$ and $T = \lceil \log_2(1/\alpha) \rceil$. Given a set of distributions \mathcal{H} , an unknown distribution P satisfying $\text{TV}(P, \mathcal{H}) = \text{OPT}$ and an (ε, δ) -DP (ξ, C) -robust PAC learner for \mathcal{H} with sample complexity $\tilde{n}_{\mathcal{H}}^C(\alpha, \beta, \varepsilon, \delta)$, there exists an (ε, δ) -DP $6C$ -agnostic PAC learner for \mathcal{H} with sample complexity*

$$n_{\mathcal{H}}^{6C}(\alpha, \beta, \varepsilon, \delta) = \tilde{n}_{\mathcal{H}}^C \left(\frac{\alpha}{12}, \frac{\beta}{2(T+4)}, \frac{\varepsilon}{2(T+4)}, \frac{\delta}{T+4} \right) + O \left(\frac{\log(T/\beta)}{\alpha^2} + \frac{\log(T/\beta)}{\alpha\varepsilon} \right).$$

We defer the proof to Section B.2.

3. Covering Unbounded Distributions

In this section, we demonstrate a simple method to prove that a set of distributions has a locally small cover. As an application, we use this result to show that the set of *unbounded* location Gaussians and scale Gaussians have locally small covers. We use these two results to give the first sample complexity result for privately learning unbounded high dimensional Gaussians in Section 4.

3.1. From Covering TV balls to Locally Small Covers

The biggest roadblock to using Theorem 26 is demonstrating the existence of a locally small cover for the set of distributions \mathcal{H} . Unfortunately, explicitly constructing a *global* cover (which is locally small) can be complicated, and may require cumbersome calculations even for “simple” distributions (see, e.g., Lemma 6.13 of Bun et al. (2019)). We offer a conceptually simpler alternative to prove a set of distributions \mathcal{H} has a locally small cover: we demonstrate that if for every $H \in \mathcal{H}$ the total variation ball $\mathcal{B}(\gamma, H, \mathcal{H})$ has an $\frac{\xi}{2}$ -cover of size no more than k , then there exists an ξ -cover for \mathcal{H} that is (k, γ) -locally small.

Lemma 29 *Given a set of distributions \mathcal{H} and $\xi \in (0, 1)$, if for every distribution $H \in \mathcal{H}$ the total variation ball $\mathcal{B}(\gamma, H, \mathcal{H}) \subseteq \mathcal{H}$ has an $\frac{\xi}{2}$ -cover of size no more than k , then there exists a (k, γ) -locally small ξ -cover for \mathcal{H} .*

Proof Fix some $H \in \mathcal{H}$. By assumption, we have that the set of distributions $\mathcal{B}(\gamma, H, \mathcal{H})$ has an $\frac{\xi}{2}$ -cover of size no more than k , which by definition implies that the $\frac{\xi}{2}$ -covering number of $\mathcal{B}(\gamma, H, \mathcal{H})$ is no more than k . By Lemma 12, the ξ -packing number of $\mathcal{B}(\gamma, H, \mathcal{H})$ is also at most k .

Now consider an ξ -packing \mathcal{P}_ξ for the set of distributions \mathcal{H} . We claim any such \mathcal{P}_ξ must be (k, γ) -locally small, and we prove this by contradiction. Suppose to the contrary that there were a distribution $H' \in \mathcal{P}_\xi$ such that $|\mathcal{B}(\gamma, H', \mathcal{P}_\xi)| > k$. This would imply that there is an ξ -packing for $\mathcal{B}(\gamma, H', \mathcal{H})$ with size larger than k , which contradicts the above observation that the packing number of *any* $\mathcal{B}(\gamma, H', \mathcal{H})$ is at most k .

A ξ -packing for \mathcal{H} is called maximal if it is impossible to add a new element of \mathcal{H} to it without violating the ξ -packing property. We claim that any maximal packing \mathcal{P}'_ξ of \mathcal{H} is also a ξ -cover of \mathcal{H} . We can prove this by contradiction. Suppose to the contrary that there were a distribution $P \in \mathcal{H}$ with $\text{TV}(P, \mathcal{P}'_\xi) > \xi$. Then we could add P to \mathcal{P}'_ξ to produce a strictly larger packing, contradicting the maximality of \mathcal{P}'_ξ . Thus taking \mathcal{P}_ξ to be a maximal packing gives us a (k, γ) -locally small ξ -cover. Therefore, it only remains to show that a maximal packing actually exists, which follows from a simple application of Zorn's Lemma.⁴ ■

3.2. Locally Small Gaussian Covers

We now prove that both the set of d -dimensional location Gaussians and scale Gaussians can be covered in a locally small fashion. Our first result shows that the set of d -dimensional location Gaussians \mathcal{G}_d^L has a locally small cover. Our second result is proving the existence of a locally small cover for the set of d -dimensional scale Gaussians \mathcal{G}_d^S .

3.2.1. COVERING LOCATION GAUSSIANS

It is not too difficult to come up with an explicit locally small cover for the set of location Gaussians without using Lemma 29 as is demonstrated in (Bun et al., 2019, Lemma 6.12). Nonetheless we choose to do so as a warmup before attempting to solve the (harder) problem for scale Gaussians. In the case of location Gaussians, our proof is very similar to Bun et al. (2019). Constructing an explicit cover is not too difficult in this case because the geometry of \mathcal{G}_d^L is “simple,” given that the TV distance between any two distributions is determined by the ℓ_2 distance of their means. Unfortunately the situation is not that simple in the scale Gaussian case as we will see shortly. We begin by showing that the total variation ball centered at any Gaussian $\mathcal{N}(\mu, I)$ with respect to \mathcal{G}_d^L can be covered, as long as the radius is not too large.

4. Let M be the set of all γ -packings of \mathcal{H} . Define a partial order on M by the relation $\mathcal{P}_1 \leq \mathcal{P}_2 \iff \mathcal{P}_1 \subseteq \mathcal{P}_2$ where $\mathcal{P}_1, \mathcal{P}_2 \in M$. We claim that every chain in this partially ordered set has an upper bound in M ; by Zorn's lemma, this would imply that M has a maximal element which concludes the proof. To see why every (possibly infinite) chain $\mathcal{P}_1 \leq \mathcal{P}_2 \leq \dots$ has an upper bound in M , we consider the following upper bound $U = \cup_i \mathcal{P}_i$. Note that $U \in M$ since otherwise there would be an index i such that $\mathcal{P}_i \notin M$.

Lemma 30 For any $d \in \mathbb{N}$, $\mu \in \mathbb{R}^d$, $\xi \in (0, 1)$ and $\gamma \in (\xi, c_1)$ where c_1 is a universal constant, there exists an ξ -cover for the set of distributions $\mathcal{B}(\gamma, \mathcal{N}(\mu, I), \mathcal{G}_d^L)$ of size

$$\left(\frac{\gamma}{\xi}\right)^{O(d)}.$$

Proof Fix some $\mathcal{N}(\mu, I) \in \mathcal{G}_d^L$. From (Devroye et al., 2018, Theorem 1.2) we have

$$\frac{1}{200} \min\{1, \|\mu_1 - \mu_2\|_2\} \leq \text{TV}(\mathcal{N}(\mu_1, I), \mathcal{N}(\mu_2, I)) \leq \frac{9}{2} \min\{1, \|\mu_1 - \mu_2\|_2\}. \quad (2)$$

For any γ smaller than the universal constant c_1 , the lower bound in (2) implies that any $\mathcal{N}(\tilde{\mu}, I) \in \mathcal{B}(\gamma, \mathcal{N}(\mu, I), \mathcal{G}_d^L)$ must satisfy $\|\mu - \tilde{\mu}\|_2 \leq 200\gamma$. We thus propose the following cover:

$$\mathcal{C}_\xi = \left\{ \mathcal{N}(\mu + \hat{z}, I) : \hat{z} \in \left(\frac{2\xi}{9\sqrt{d}}\right) \mathbb{Z}^d, \|\hat{z}\|_2 \leq 200\gamma \right\}.$$

We now prove that \mathcal{C}_ξ is a valid ξ -cover. Fix some $\mathcal{N}(\tilde{\mu}, I) \in \mathcal{B}(\gamma, \mathcal{N}(\mu, I), \mathcal{G}_d^L)$ and define $z = \tilde{\mu} - \mu$. We know $\|z\|_2 \leq 200\gamma$. Let $\hat{z} = \left(\frac{2\xi}{9\sqrt{d}}\right) \lfloor \left(\frac{9\sqrt{d}}{2\xi}\right) z \rfloor$ and $\hat{\mu} = \mu + \hat{z}$. Note that we have $\mathcal{N}(\hat{\mu}, I) \in \mathcal{C}_\xi$. Furthermore, z and \hat{z} are element-wise close ($\|z - \hat{z}\|_\infty \leq 2\xi/9\sqrt{d}$) therefore we have

$$\begin{aligned} \text{TV}(\mathcal{N}(\tilde{\mu}, I), \mathcal{N}(\hat{\mu}, I)) &\leq \frac{9}{2} \|\tilde{\mu} - \hat{\mu}\|_2 \\ &= \frac{9}{2} \|z - \hat{z}\|_2 \\ &\leq \frac{9\sqrt{d}}{2} \|z - \hat{z}\|_\infty \\ &\leq \frac{9\sqrt{d}}{2} \cdot \frac{2\xi}{9\sqrt{d}} \\ &= \xi, \end{aligned}$$

where the first inequality follows from (2). We now bound the size of this cover.

$$\begin{aligned}
 |\mathcal{C}_\xi| &= \left| \left\{ \mathcal{N}(\mu + \hat{z}, I) : \hat{z} \in \left(\frac{2\xi}{9\sqrt{d}} \right) \mathbb{Z}^d, \|\hat{z}\|_2 \leq 200\gamma \right\} \right| \\
 &\leq \left| \left\{ \hat{z} : \hat{z} \in \mathbb{Z}^d, \|\hat{z}\|_2 \leq \frac{900\sqrt{d}\gamma}{\xi} \right\} \right| \leq \left| \left\{ \hat{z} : \hat{z} \in \mathbb{Z}^d, \|\hat{z}\|_1 \leq \frac{900d\gamma}{\xi} \right\} \right| \\
 &\leq \left| \left\{ z_1 - z_2 : z_1, z_2 \in \mathbb{Z}_+^d, \|z_1\|_1 \leq \left\lceil \frac{900d\gamma}{\xi} \right\rceil, \|z_2\|_1 \leq \left\lceil \frac{900d\gamma}{\xi} \right\rceil \right\} \right| \\
 &\leq \left| \left\{ z : z \in \mathbb{Z}_+^d, \|z\|_1 \leq \left\lceil \frac{900d\gamma}{\xi} \right\rceil \right\} \right|^2 \\
 &\leq \left(\sum_{i=1}^{\lceil 900d\gamma/\xi \rceil} \binom{i+d-1}{d-1} \right)^2 \\
 &\leq \left(\left\lceil \frac{900d\gamma}{\xi} \right\rceil \binom{\lceil 900d\gamma/\xi \rceil + d - 1}{d - 1} \right)^2 \leq \left(\frac{\gamma}{\xi} \right)^{O(d)},
 \end{aligned}$$

where the third last inequality follows from the standard solution to the stars and bars problem. \blacksquare

Combining Lemma 29 with Lemma 30 immediately gives us the following corollary:

Corollary 31 *For any $d \in \mathbb{N}$, $\xi \in (0, 1)$ and $\gamma \in (\xi, c_1)$ where c_1 is a universal constant, there exists an ξ -cover \mathcal{C}_ξ for the set of d -dimensional location Gaussians \mathcal{G}_d^L that is $((2\gamma/\xi)^{O(d)}, \gamma)$ -locally small.*

3.2.2. COVERING SCALE GAUSSIANS

It is not a trivial exercise to come up with an explicit cover for the class of scale Gaussians due to the complicated nature of the geometry of \mathcal{G}_d^S . Fortunately for us, Lemma 29 simplifies things significantly. It turns out that if we want to cover the TV ball centered at any $\mathcal{N}(0, \Sigma)$, we can use a cover for $\mathcal{N}(0, I)$ and “stretch” the covariance matrices of every distribution in the cover (using Σ) so that the modified cover becomes a valid cover for the TV ball centered at $\mathcal{N}(0, \Sigma)$. The following lemma tells us that we can cover the total variation ball centered at $\mathcal{N}(0, I)$ with respect to \mathcal{G}_d^S as long as the radius is not too large.

Lemma 32 *For any $d \in \mathbb{N}$, $\xi \in (0, 1)$ and $\gamma \in (\xi, c_2)$ where c_2 is a universal constant, there exists an ξ -cover for the set of distributions $\mathcal{B}(\gamma, \mathcal{N}(0, I), \mathcal{G}_d^S)$ of size*

$$\left(\frac{\gamma}{\xi} \right)^{O(d^2)}.$$

Proof

From (Devroye et al., 2018, Theorem 1.1) we have,

$$\text{TV}(\mathcal{N}(0, I), \mathcal{N}(0, \Sigma)) \geq \frac{1}{100} \min \left\{ 1, \sqrt{\sum_{i=1}^d \lambda_i^2} \right\}, \quad (3)$$

where $\lambda_1 \dots \lambda_d$ are the eigenvalues of $\Sigma - I$, and it holds that $\sqrt{\sum_{i=1}^d \lambda_i^2} = \|\Sigma - I\|_F$.

For any γ smaller than the universal constant c'_2 , the lower bound in (3) implies two things: 1) for any $\mathcal{N}(0, \Sigma) \in \mathcal{B}(\gamma, \mathcal{N}(0, I), \mathcal{G}^S)$, $\|\Sigma - I\|_F \leq 100\gamma$ and 2) the minimum eigenvalue of Σ , λ_{\min} , satisfies $\lambda_{\min} \geq 1 - 100\gamma$. We thus propose the following cover:

$$\mathcal{C}_\xi = \left\{ \mathcal{N}(0, I + \hat{\Delta}) : \hat{\Delta} \in \rho \mathbb{Z}^{d \times d} \cap \mathbb{S}_d, \|\hat{\Delta}\|_F \leq 100\gamma \right\},$$

where $\rho = \frac{\xi\sqrt{2\pi e}(1-100\gamma)}{d+\xi\sqrt{2\pi e}}$. First we will show that this is a valid cover. Consider an arbitrary $\mathcal{N}(0, \Sigma) \in \mathcal{B}(\gamma, \mathcal{N}(0, I), \mathcal{G}_d^S)$. We want to show that there is a distribution $\mathcal{N}(0, \hat{\Sigma}) \in \mathcal{C}_\xi$ that is ξ -close to $\mathcal{N}(0, \Sigma)$. Let $\Delta = \Sigma - I$, let $\hat{\Delta} = \rho \lfloor \Delta / \rho \rfloor$ and let $\hat{\Sigma} = I + \hat{\Delta}$. Since $\|\Delta\|_F = \|\Sigma - I\|_F \leq 100\gamma$, $\mathcal{N}(0, \hat{\Sigma})$ is indeed in the cover.

Next we show that $\text{TV}(\mathcal{N}(0, \Sigma), \mathcal{N}(0, \hat{\Sigma})) \leq \xi$. We use Proposition 32 in Valiant and Valiant (2010), which states for any two positive definite matrices Σ and $\hat{\Sigma}$, if $\|\Sigma - \hat{\Sigma}\|_{\infty, \infty} \leq \rho'$ and the smallest eigenvalue of Σ satisfies $\lambda_{\min} > \eta$, then we have

$$\text{TV}(\mathcal{N}(0, \Sigma), \mathcal{N}(0, \hat{\Sigma})) \leq \frac{d\rho'}{\sqrt{2\pi e}(\eta - \rho')}. \quad (4)$$

By the definition of \mathcal{C}_ξ , $\|\hat{\Delta} - \Delta\|_{\infty, \infty} \leq \rho$. Since any valid Σ must satisfy $\lambda_{\min} \geq 1 - 100\gamma$, our choice of setting $\rho = \frac{\xi\sqrt{2\pi e}(1-100\gamma)}{d+\xi\sqrt{2\pi e}}$ implies that

$$\text{TV}(\mathcal{N}(0, \Sigma), \mathcal{N}(0, \hat{\Sigma})) \leq \xi,$$

for any γ smaller than the universal constant c'_2 . We now bound the size of the cover in a similar manner to the case of location Gaussians.

$$\begin{aligned} |\mathcal{C}_\xi| &= \left| \left\{ \hat{\Delta} \in \rho \mathbb{Z}^{d \times d} \cap \mathbb{S}_d : \|\hat{\Delta}\|_F \leq 100\gamma \right\} \right| \\ &\leq \left| \left\{ \hat{\Delta} \in \mathbb{Z}^{d \times d} : \|\hat{\Delta}\|_F \leq 100\gamma/\rho \right\} \right| \\ &\leq \left| \left\{ \hat{\Delta} \in \mathbb{Z}^{d \times d} : \|\hat{\Delta}\|_{1,1} \leq 100\gamma d/\rho \right\} \right| \\ &\leq \left| \left\{ \hat{\Delta} \in \mathbb{Z}_+^{d \times d} : \|\hat{\Delta}\|_{1,1} \leq \lceil 100\gamma d/\rho \rceil \right\} \right|^2 \\ &\leq \left(\left\lceil \frac{100\gamma d(d + \xi\sqrt{2\pi e})}{\xi\sqrt{2\pi e}(1 - 100\gamma)} \right\rceil \cdot \left(\left\lceil \frac{100\gamma d(d + \xi\sqrt{2\pi e})}{\xi\sqrt{2\pi e}(1 - 100\gamma)} \right\rceil + d^2 - 1 \right) \right)^2, \end{aligned}$$

for any γ smaller than the universal constant c''_2 we have,

$$|\mathcal{C}_\xi| \leq \left(\frac{\gamma}{\xi} \right)^{O(d^2)}.$$

Setting $c_2 = \max\{c'_2, c''_2\}$ completes the proof. ■

The following corollary is a direct consequence of Lemma 32 and Proposition 44.

Corollary 33 For any $d \in \mathbb{N}$, $\xi \in (0, 1)$, $\gamma \in (0, c_2)$ where c_2 is a universal constant and $\Sigma \in \mathbb{S}_d$, there exists an ξ -cover for the set of distributions $\mathcal{B}(\gamma, \mathcal{N}(0, \Sigma), \mathcal{G}_d^S)$ of size

$$\left(\frac{\gamma}{\xi}\right)^{O(d^2)}.$$

Proof Fix some $\Sigma \in \mathbb{S}_d$ and define $\Sigma^{1/2}$ as one of its matrix square-roots. By Proposition 44 and Equation (1) we have:

$$\text{TV}\left(\mathcal{N}\left(0, \Sigma^{1/2}\Sigma_1\Sigma^{1/2}\right), \mathcal{N}\left(0, \Sigma^{1/2}\Sigma_2\Sigma^{1/2}\right)\right) \leq \text{TV}\left(\mathcal{N}\left(0, \Sigma_1\right), \mathcal{N}\left(0, \Sigma_2\right)\right). \quad (5)$$

We can thus take the cover \mathcal{C}_ξ in Lemma 32, and replace every distribution $\mathcal{N}(0, \Sigma_1) \in \mathcal{C}_\xi$ with $\mathcal{N}(0, \Sigma^{1/2}\Sigma_1\Sigma^{1/2})$. Note that our modified cover will have the same size. From (5), our new cover will be a valid ξ -cover for $\mathcal{B}(\gamma, \mathcal{N}(0, \Sigma), \mathcal{G}_d^S)$ since the TV distance can not increase between any two distributions $\mathcal{N}(0, \Sigma_1), \mathcal{N}(0, \Sigma_2) \in \mathcal{B}(\gamma, \mathcal{N}(0, I), \mathcal{G}_d^S)$ after applying the transformation above. \blacksquare

We can now combine Lemma 29 with Corollary 33 to get the following:

Corollary 34 For any $d \in \mathbb{N}$, $\xi \in (0, 1)$, and $\gamma \in (\xi, c_2)$ where c_2 is a universal constant, there exists an ξ -cover \mathcal{C}_ξ for the set of scale Gaussians \mathcal{G}_d^S that is $\left((2\gamma/\xi)^{O(d^2)}, \gamma\right)$ -locally small.

4. Beyond GAP-MAX: Boosting Weak Hypotheses

As we mentioned before, by using a $(k, 6\xi + \alpha)$ -locally small ξ -cover for an infinite set of distributions \mathcal{H} , one can utilize Theorem 26 to privately learn a distribution to low error. Unfortunately, this approach will yield a sample complexity bound that has a term of order $O(\log(1/\delta)/\alpha\varepsilon)$. In the case of learning an unbounded univariate Gaussian in the realizable setting, it is known that the sample complexity is $O(1/\alpha^2 + \log(1/\delta)/\varepsilon)$ Karwa and Vadhan (2018), however the upper bound on the sample complexity achieved by Theorem 26 (together with an appropriate locally smaller cover) is $O(1/\alpha^2 + \log(1/\delta)/\alpha\varepsilon)$ (Bun et al., 2019, Corollary 6.15). In order to overcome the poor dependence on $\log(1/\delta)$, we can instead aim for a two step approach:

1. Use the GAP-MAX algorithm in Theorem 26 but with constant accuracy C to learn a distribution H' that is roughly C -close to the true Gaussian for some appropriately selected constant $C < 1$.
2. Build a *finite* cover for $\mathcal{B}(C, H', \mathcal{H})$ and use the private hypothesis selection algorithm (Theorem 25) to learn a distribution \hat{H} that is α -close to the true Gaussian.

Running the GAP-MAX algorithm with constant accuracy C thus removes the dependence on α in the $O(\log(1/\delta)/\varepsilon)$ term. Intuitively, this approach learns a “rough” estimate of the right distribution using approximate differential privacy. Since we know that we are roughly C -close to the true Gaussian, we can cover $\mathcal{B}(C, H', \mathcal{H})$ with a finite cover, and use the ε -differentially private hypothesis selection algorithm. This two step approach which we dub *boosting* gets us a much better dependence on the privacy parameter δ in our sample complexity bounds, and as we will see it holds more generally in the robust learning setting.

Remark 35 We note that the first step in the above approach may only need to produce an exceptionally coarse estimate to the true distribution – one to which it bears very little resemblance at all! We illustrate this with the simple problem of privately estimating a univariate Gaussian $\mathcal{N}(\mu, 1)$ (in the realizable case).

We work backwards: our overall target is an algorithm with sample complexity $\tilde{O}(1/\alpha^2 + 1/\alpha\varepsilon + \log(1/\delta)/\varepsilon)$. Since using the pure DP hypothesis selection algorithm of Theorem 25 takes $O(\log |\mathcal{C}_\alpha|(1/\alpha^2 + 1/\alpha\varepsilon))$ samples, we require only that $|\mathcal{C}_\alpha|$ is less than some quasi-polynomial in $1/\alpha$. For the sake of exposition, suppose we restrict further and require $|\mathcal{C}_\alpha| \leq 1/\alpha^{101}$. This can be achieved by starting at any point which is at most $1/\alpha^{100}$ from the true mean μ and taking an α -additive grid over this space. But if we only require a starting point $\hat{\mu}$ which is $1/\alpha^{100}$ -close to the true mean μ , this corresponds (by Gaussian tail bounds) to a distribution whose total variation distance is roughly $1 - \exp(-1/\alpha^{200})$ with respect to the true distribution.

We can see that the first step in the procedure truly requires an exceptionally coarse estimate of the distribution. The estimate of the mean described is significantly further from the true mean than any individual point will be. Interestingly, note that if one requires a more accurate final distribution, the distribution output in the first step is allowed to be less accurate.

4.1. Warmup: Learning Location Gaussians

As a first step, we can show that Algorithm 1 can achieve a slightly more general guarantee than a robust PAC learning sample complexity bound. We make Algorithm 1 more general than it needs to be to give a robust learning guarantee for \mathcal{G}_d^L in order to make use of it as a subroutine in Algorithm 2 which robustly learns \mathcal{G}_d .

Lemma 36 Let b be a positive constant. For any $\beta, \varepsilon, \delta \in (0, 1)$, $\alpha \in (0, c_3)$ and $\xi \in (0, c_4)$ where c_3 and c_4 are constants that depend only on b , given a dataset $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{G}_d^L) \leq b\xi + \alpha/4$, $\text{BOOST}_1(b, \xi, \alpha, \beta, \varepsilon, \delta, \mathcal{C}_\xi, D)$ is an (ε, δ) -differentially private algorithm which outputs some $\hat{H} \in \mathcal{G}_d^L$ such that $\text{TV}(\hat{H}, P) \leq 3(b+1)\xi + \alpha$ with probability no less than $1 - \beta$, so long as

$$n = O\left(\frac{d + \log(1/\beta)}{\alpha^2} + \frac{d + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

Algorithm 1 Boosting for learning \mathcal{G}_d^L : $\text{BOOST}_1(b, \xi, \alpha, \beta, \varepsilon, \delta, \mathcal{C}_\xi, D)$.

Input : Positive constant b , robustness parameter $\xi \in (0, 1)$, accuracy parameters $\alpha, \beta \in (0, 1)$, privacy parameters $\varepsilon, \delta \in (0, 1)$, locally small ξ -cover \mathcal{C}_ξ for \mathcal{G}_d^L and dataset D of size n .

Output: Distribution $\hat{H} \in \mathcal{G}_d^L$.

1. $H' = \text{GAP-MAX}\left(b\xi + \frac{\alpha}{4}, \left(\frac{1}{200(600b+1)} - \frac{3\alpha}{4}\right), \beta/2, \varepsilon/2, \delta, k, \mathcal{C}_\xi, D\right)$ // $H' = \mathcal{N}(\mu', I)$
 2. Build an ξ -cover $\tilde{\mathcal{C}}_\xi$ for $\mathcal{B}\left(\left(3b\xi + \frac{1}{200(600b+1)}\right), H', \mathcal{G}_d^L\right)$
 3. **Return** $\hat{H} = \text{PHS}((b+1)\xi + \alpha, \alpha, \beta/2, \varepsilon/2, \tilde{\mathcal{C}}_\xi, D)$ // $\hat{H} = \mathcal{N}(\hat{\mu}, I)$
-

Proof [Proof of Lemma 36] We first show Algorithm 1 satisfies (ε, δ) -differential privacy. Line 1 of the algorithm is $(\varepsilon/2, \delta)$ -differentially private by the guarantee of Theorem 26. Line 2 maintains

$(\varepsilon/2, \delta)$ -privacy by post-processing (Lemma 21). Finally, line 3 is $(\varepsilon/2, 0)$ -differentially private by Theorem 25. By composition (Lemma 20), the entire algorithm is (ε, δ) -differentially private.

We now argue about the accuracy of the algorithm. By Lemma 17, Corollary 31 and Theorem 26, as long as $n = O\left(\frac{d+\log(1/\beta\delta)}{\varepsilon}\right)$, with probability no less than $1 - \beta/2$ the GAP-MAX algorithm in line 1 outputs a distribution H' that is $\left(3b\xi + \frac{1}{200(600b+1)}\right)$ -close to H^* , for any ξ and α smaller than constants c'_3 and c_4 , respectively, that depend only on b . For the remainder of the proof we condition on this event.

In line 2, we build an ξ -cover $\tilde{\mathcal{C}}_\xi$ for $\mathcal{B}\left(\left(3b\xi + \frac{1}{200(600b+1)}\right), H', \mathcal{G}_d^L\right)$ that satisfies $\text{TV}(H^*, \tilde{\mathcal{C}}_\xi) \leq \xi$. By the triangle inequality $\text{TV}(P, \tilde{\mathcal{C}}_\xi) \leq (b+1)\xi + \alpha/4$. By Lemma 30, we can indeed construct $\tilde{\mathcal{C}}_\xi$ such that $|\tilde{\mathcal{C}}_\xi| \leq (C)^{O(d)}$ (for some constant C) as long as ξ is smaller than a constant c''_3 that depends only on b . By the stated accuracy and size of $\tilde{\mathcal{C}}_\xi$, with probability no less than $1 - \beta/2$, Theorem 25 guarantees that line 3 outputs \hat{H} satisfying

$$\text{TV}(P, \hat{H}) \leq 3((b+1)\xi + \alpha/4) + \alpha/4 = 3(b+1)\xi + \alpha,$$

as long as $n = O\left(\frac{d+\log(1/\beta)}{\alpha^2} + \frac{d+\log(1/\beta)}{\alpha\varepsilon}\right)$. Setting $n = O\left(\frac{d+\log(1/\beta)}{\alpha^2} + \frac{d+\log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right)$ together with a union bound completes the proof. \blacksquare

The following result can be derived from the BOOST_1 algorithm by taking the standard assumption in the robust setting of $\text{TV}(P, \mathcal{H}) \leq \xi$. The proof is nearly identical to the proof of Lemma 36.

Lemma 37 *Let $\alpha, \beta, \varepsilon, \delta \in (0, 1)$ and $\xi \in (0, c_5)$ where c_5 is a universal constant, and let $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{G}_d^L) \leq \xi$. Furthermore, let \mathcal{C}_ξ be an appropriately selected locally small ξ -cover for \mathcal{G}_d^L . $\text{BOOST}_1(1, \xi, \alpha, \beta, \varepsilon, \delta, \mathcal{C}_\xi, D)$ is an (ε, δ) -DP $(\xi, 6)$ -robust PAC learner for \mathcal{G}_d^L with sample complexity*

$$\tilde{n}_{\mathcal{G}_d^L}^6(\alpha, \beta, \varepsilon, \delta) = O\left(\frac{d + \log(1/\beta)}{\alpha^2} + \frac{d + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

We can now use Lemma 28 together with Lemma 37 to get a semi-agnostic algorithm.

Theorem 38 *Let $\alpha, \beta, \varepsilon, \delta \in (0, 1)$ and $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{G}_d^L) = \text{OPT}$. For any OPT smaller than a universal constant c_5 , there exists an (ε, δ) -DP 36-agnostic PAC learner for \mathcal{G}_d^L with sample complexity*

$$\tilde{n}_{\mathcal{G}_d^L}^{36}(\alpha, \beta, \varepsilon, \delta) = \tilde{O}\left(\frac{d + \log(1/\beta)}{\alpha^2} + \frac{d + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

4.2. Learning Gaussians

We can now show that Algorithm 2 achieves the following sample complexity bound for robust learning.

Lemma 39 *Let $\beta, \varepsilon, \delta \in (0, 1)$, $\alpha \in (0, c_6)$ and $\xi \in (0, c_7)$ where c_6 and c_7 are universal constants, and let $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{G}_d^L) \leq \xi$. Furthermore, let \mathcal{C}_ξ^1 and \mathcal{C}_ξ^2 be appropriately selected locally small covers for \mathcal{G}_d and \mathcal{G}_d^L respectively. $\text{BOOST}_2(\xi, \alpha, \beta, \varepsilon, \delta, \mathcal{C}_\xi^1, \mathcal{C}_\xi^2, D)$ is an (ε, δ) -DP $(\xi, 33)$ -robust PAC learner for \mathcal{G}_d with sample complexity*

$$\tilde{n}_{\mathcal{G}_d}^{33}(\alpha, \beta, \varepsilon, \delta) = O\left(\frac{d^2 + \log(1/\beta)}{\alpha^2} + \frac{d^2 + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

Algorithm 2 Boosting for learning \mathcal{G}_d : $\text{BOOST}_2(\xi, \alpha, \beta, \varepsilon, \delta, \mathcal{C}_\xi^1, \mathcal{C}_\xi^2, D)$.

Input : Robustness parameter $\xi \in (0, 1)$, accuracy parameters $\alpha, \beta \in (0, 1)$, privacy parameters $\varepsilon, \delta \in (0, 1)$, locally small ξ -cover \mathcal{C}_ξ^1 for \mathcal{G}_d^S , locally small ξ -cover \mathcal{C}_ξ^2 for \mathcal{G}_d^L and dataset D of size $2n$.

Output: Distribution $\hat{H} \in \mathcal{G}_d$.

1. Set $D' = \{Y_1, \dots, Y_n\}$ where $Y_i = (X_{2i} - X_{2i-1})/\sqrt{2}$
 2. $H'_1 = \text{GAP-MAX}\left(4\xi, \frac{1}{100(1201)}, \beta/4, \varepsilon/4, \delta/2, k_1, \mathcal{C}_\xi^1, D'\right)$ // $H'_1 = \mathcal{N}(0, \Sigma')$
 3. Build an ξ -cover $\tilde{\mathcal{C}}_\xi^1$ for $\mathcal{B}\left(\left(12\xi + \frac{1}{100(1201)}\right), H'_1, \mathcal{G}_d^S\right)$
 4. $\hat{H}_1 = \text{PHS}(4\xi, \alpha, \beta/4, \varepsilon/4, \tilde{\mathcal{C}}_\xi^1, D')$ // $\hat{H}_1 = \mathcal{N}(0, \hat{\Sigma})$
 5. Set $D'' = \{W_1, \dots, W_{2n}\}$ where $W_i = \hat{\Sigma}^{-1/2}X_i$
 6. $\hat{H}_2 = \text{BOOST}_1\left(10, \xi, \alpha, \beta/2, \varepsilon/2, \delta/2, \mathcal{C}_\xi^2, D''\right)$ // $\hat{H}_2 = \mathcal{N}(\hat{\mu}, I)$
 7. **Return**: $\hat{H} = \mathcal{N}\left(\hat{\Sigma}^{1/2}\hat{\mu}, \hat{\Sigma}\right)$
-

Proof [Proof of Lemma 39] We first show Algorithm 2 satisfies (ε, δ) -differential privacy. Line 2 of the algorithm is $(\varepsilon/4, \delta/2)$ -differentially private by the guarantee of Theorem 26. Line 3 maintains $(\varepsilon/4, \delta/2)$ -privacy by post-processing (Lemma 21). Line 4 is $(\varepsilon/4, 0)$ -differentially private by Theorem 25. Line 5 maintains privacy by post processing (Lemma 21). Finally, line 6 is $(\varepsilon/2, \delta/2)$ -differentially private by the privacy of Algorithm 1 proved in Lemma 36. By composition (Lemma 20) the entire algorithm is (ε, δ) -differentially private.

We now argue about the accuracy of the algorithm. Let $H^* = \mathcal{N}(\mu^*, \Sigma^*)$ be the hypothesis in \mathcal{G}_d that satisfies $\text{TV}(H^*, P) \leq \xi$. By Lemma 46 $Y_i \sim Q$ where $\text{TV}(Q, \mathcal{N}(0, \Sigma^*)) \leq 3\xi$, which implies that $\text{TV}(\mathcal{N}(0, \Sigma^*), \mathcal{C}_\xi^1) \leq 4\xi$. From Lemma 17, Corollary 34 and Theorem 26, as long as $n = O\left(\frac{d^2 + \log(1/\beta\delta)}{\varepsilon}\right)$, with probability no less than $1 - \beta/4$ the GAP-MAX algorithm in line 2 outputs a distribution $H'_1 = \mathcal{N}(0, \Sigma')$ that is $\left(12\xi + \frac{1}{100(1201)}\right)$ -close to $\mathcal{N}(0, \Sigma^*)$, for any ξ smaller than a universal constant c'_6 . We condition on this event.

In line 3, we build a ξ -cover $\tilde{\mathcal{C}}_\xi^1$ for $\mathcal{B}\left(12\xi + \frac{1}{100(1201)}, H'_1, \mathcal{G}_d^S\right)$ that satisfies $\text{TV}(\mathcal{N}(0, \Sigma^*), \tilde{\mathcal{C}}_\xi^1) \leq \xi$. It follows immediately that $\text{TV}(Q, \tilde{\mathcal{C}}_\xi^1) \leq 4\xi$ from the triangle inequality. By Corollary 33, we can indeed construct $\tilde{\mathcal{C}}_\xi^1$ such that $|\tilde{\mathcal{C}}_\xi^1| \leq (C)^{O(d^2)}$ for any ξ smaller than a universal constant c''_6 . By the stated accuracy and size of $\tilde{\mathcal{C}}_\xi^1$ Theorem 25 guarantees, with probability at least

$1 - \beta/4$, that line 4 outputs $\widehat{H}_1 = \mathcal{N}(0, \widehat{\Sigma})$ such that $\text{TV}(\widehat{H}_1, \mathcal{N}(0, \Sigma^*)) \leq 8\xi + \alpha/4$ as long as $n = O\left(\frac{d^2 + \log(1/\beta)}{\alpha^2} + \frac{d^2 + \log(1/\beta)}{\alpha\varepsilon}\right)$. We further condition on this event.

Let R be the distribution satisfying $W_i \sim R$.⁵ By Corollary 45, $\text{TV}(R, \mathcal{N}(\mu^*, \widehat{\Sigma}^{-1/2}\Sigma^*\widehat{\Sigma}^{-1/2})) \leq \xi$. By application of the triangle inequality, Corollary 45 and Equation (1),

$$\begin{aligned} \text{TV}(R, \mathcal{N}(\mu^*, I)) &\leq \text{TV}(R, \mathcal{N}(\mu^*, \widehat{\Sigma}^{-1/2}\Sigma^*\widehat{\Sigma}^{-1/2})) + \text{TV}(\mathcal{N}(\mu^*, \widehat{\Sigma}^{-1/2}\Sigma^*\widehat{\Sigma}^{-1/2}), \mathcal{N}(\mu^*, I)) \\ &\leq \xi + \text{TV}(\mathcal{N}(\mu^*, \Sigma^*), \mathcal{N}(\mu^*, \widehat{\Sigma})) \\ &= \xi + \text{TV}(\mathcal{N}(0, \Sigma^*), \widehat{H}_1) \\ &\leq 9\xi + \alpha/4. \end{aligned}$$

It follows from the triangle inequality that $\text{TV}(R, \mathcal{C}_\xi^2) \leq 10\xi + \alpha/4$. This together with Lemma 36 implies that, with probability greater than $1 - \beta/2$, $\text{BOOST}_1(10, \xi, \alpha, \beta/2, \varepsilon/2, \delta/2, \mathcal{C}_\xi^2, D'')$ outputs $\widehat{H}_2 = \mathcal{N}(\widehat{\mu}, I)$ satisfying $\text{TV}(\widehat{H}_2, \mathcal{N}(\mu^*, I)) \leq 33\xi + \alpha$ for any ξ and α smaller than universal constants c_6''' and c_7 respectively, as long as $n = O\left(\frac{d + \log(1/\beta)}{\alpha^2} + \frac{d + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right)$. Using the triangle inequality and Corollary 45 it follows that $\text{TV}(\mathcal{N}(\widehat{\Sigma}^{1/2}\widehat{\mu}, \widehat{\Sigma}), P) \leq 33\xi + \alpha$. Setting $n = O\left(\frac{d^2 + \log(1/\beta)}{\alpha^2} + \frac{d^2 + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right)$ and $c_6 = \max\{c_6', c_6'', c_6'''\}$ together with a union bound completes the proof. \blacksquare

Finally, we can combine the above result with Lemma 28 to get a semi-agnostic sample complexity bound for modest levels of model misspecification.

Theorem 40 *Let $\beta, \varepsilon, \delta \in (0, 1)$, $\alpha \in (0, c_6)$ for some universal constant c_6 , and let $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{G}_d^L) = \text{OPT}$. For any OPT smaller than a universal constant c_7 , there exists an (ε, δ) -DP 198-agnostic PAC learner for \mathcal{G}_d with sample complexity*

$$n_{\mathcal{G}_d}^{198}(\alpha, \beta, \varepsilon, \delta) = \widetilde{O}\left(\frac{d^2 + \log(1/\beta)}{\alpha^2} + \frac{d^2 + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

4.3. Bounds for the Realizable Setting

The following sample complexity bounds hold for (ε, δ) -DP (realizable) PAC learning. The proofs are very similar to the proofs in Section 4.1 and 4.2 for (ε, δ) -DP robust PAC learning, where the slight difference is that we can build the covers directly with accuracy α (instead of ξ) since we assume realizability. The first bound is tight and the second one is conjectured to be tight.

Lemma 41 *For any $\beta, \varepsilon, \delta \in (0, 1)$ and $\alpha \in (0, c_8)$ where c_8 is a universal constant, there exists an (ε, δ) -DP PAC learner for \mathcal{G}_d^L with sample complexity*

$$n_{\mathcal{G}_d^L}(\alpha, \beta, \varepsilon, \delta) = O\left(\frac{d + \log(1/\beta)}{\alpha^2} + \frac{d + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

5. If $\widehat{\Sigma}$ is not invertible, the range of $\widehat{\Sigma}$ is an r -dimensional linear subspace of \mathbb{R}^d , for some $r < d$. Let Π be a $d \times r$ matrix whose columns form a basis for the range of $\widehat{\Sigma}$. It follows that $\widetilde{\Sigma} = \Pi^T \widehat{\Sigma} \Pi$ is a valid r -dimensional full-rank covariance matrix. Moreover, by construction, $\widetilde{\Sigma}$ is identical to $\widehat{\Sigma}$ after projection on to the subspace defined by the range of $\widehat{\Sigma}$. We can thus project our data onto the range of $\widehat{\Sigma}$ and continue the algorithm using $\widetilde{\Sigma}$.

Lemma 42 For any $\beta, \varepsilon, \delta \in (0, 1)$ and $\alpha \in (0, c_9)$ where c_9 is a universal constant, there exists an (ε, δ) -DP PAC learner for \mathcal{G}_d with sample complexity

$$n_{\mathcal{G}_d}(\alpha, \beta, \varepsilon, \delta) = O\left(\frac{d^2 + \log(1/\beta)}{\alpha^2} + \frac{d^2 + \log(1/\beta)}{\alpha\varepsilon} + \frac{\log(1/\beta\delta)}{\varepsilon}\right).$$

5. Agnostic Private Hypothesis Selection

In this section we present an ε -DP semi-agnostic PAC learner that achieves the same sample complexity as the PHS algorithm of Theorem 25. The PHS algorithm is based on the celebrated *Scheffé tournament* (see, e.g., Chapter 6 of Devroye and Lugosi (2001)), where the distributions in \mathcal{H} play a round robin tournament against one another. The winner of this tournament is then chosen as the output. One of the technical difficulties in constructing privatized versions of the Scheffé tournament via the exponential mechanism is that a single sample can quite drastically change the outcome of the tournament, which makes choosing score functions based on tournaments challenging. We sidestep this issue completely by considering another approach to hypothesis selection called the *minimum distance estimate* (MDE). The MDE approach is based on *maximizing* a particular function of the data and \mathcal{H} as we will see shortly. Fortunately, this estimator is already in the form of a maximization problem and the function we aim to maximize has low sensitivity. Thus, using the exponential mechanism together with the MDE is a very natural way to privatize semi-agnostic hypothesis selection. The MDE requires $O(m^3)$ computations, where m is the number of hypotheses in \mathcal{H} . Mahalanabis and Stefankovic (2008) presented a modified MDE that is very similar to the original MDE, but only requires $O(m^2)$ computations. Fortunately, this modified algorithm maintains the guarantee of the original algorithm, so we will privatize the modified MDE instead of the original MDE. We formally restate Theorem 27 below.

Theorem 43 Let $\mathcal{H} = \{H_1, \dots, H_m\}$ be a set of probability distributions, $\alpha, \beta, \varepsilon \in (0, 1)$ and $D \sim P^n$ where P satisfies $\text{TV}(P, \mathcal{H}) = \text{OPT}$. There exists an ε -DP 3-agnostic PAC learner with sample complexity

$$n_{\mathcal{H}}^3(\alpha, \beta, \varepsilon, 0) = O\left(\frac{\log(m/\beta)}{\alpha^2} + \frac{\log(m/\beta)}{\alpha\varepsilon}\right).$$

Before we prove the result, we define a few things. For an ordered pair of distributions (H_i, H_j) over a common domain \mathcal{X} , we define their *Scheffé set* as $A_{ij} = \{x \in \mathcal{X} : H_i(x) > H_j(x)\}$. A useful version of the TV distance between two distributions we will make use of is

$$2 \text{TV}(H_i, H_j) = (H_i(A_{ij}) - H_j(A_{ij})) + (H_j(A_{ji}) - H_i(A_{ji})).$$

We note that most of the analysis below is standard in proving the correctness of the MDE (e.g., see the proof of Theorem 6.3 in Devroye and Lugosi (2001)), and is slightly adapted using the analysis of the modified MDE algorithm in Theorem 4 of Mahalanabis and Stefankovic (2008). The only difference here is the use of the exponential mechanism.

Proof [Proof of Theorem 27] Let \mathcal{X} be the common domain of the distributions in \mathcal{H} . For a dataset D and set $A \subseteq \mathcal{X}$, we define $\hat{P}(A, D) = \frac{1}{n} \cdot |\{x \in D : x \in A\}|$. For a distribution H and a set

$A \subseteq \mathcal{X}$, let $R(H, A) = H(A) - \widehat{P}(A)$. For any $H_i \in \mathcal{H}$, we define the score function

$$\begin{aligned} S(D, H_i) &= - \sup_{j \in [m] \setminus \{i\}} \left| (H_i(A_{ij}) - \widehat{P}(A_{ij}, D)) - (H_i(A_{ji}) - \widehat{P}(A_{ji}, D)) \right| \\ &= - \sup_{j \in [m] \setminus \{i\}} |R(H_i, A_{ij}) - R(H_i, A_{ji})|. \end{aligned}$$

With this in place, the algorithm is simple: run the exponential mechanism [McSherry and Talwar \(2007\)](#) with this score function, on the set of candidates \mathcal{H} , with dataset D , and return whichever distribution it outputs.

It is not hard to see that the score function has sensitivity $2/n$. Let $H_k \in \mathcal{H}$ be any distribution that maximizes the score function. From [Theorem 19](#), it follows that running the exponential mechanism with our dataset D , the set of distributions \mathcal{H} , privacy parameter ε and the the score function above outputs a distribution $H_{k'} \in \mathcal{H}$ that guarantees, with probability no less than $1 - \beta/2$,

$$\begin{aligned} S(D, H_{k'}) &\geq S(D, H_k) - \frac{4 \log(2m/\beta)}{n\varepsilon} \\ &\geq S(D, H_k) - \alpha, \end{aligned}$$

where the last line holds so long as $n = O\left(\frac{\log(m/\beta)}{\alpha\varepsilon}\right)$. We condition on this event, which can equivalently be stated as

$$\sup_{j \in [m] \setminus \{k'\}} |R(H_{k'}, A_{k'j}) - R(H_{k'}, A_{jk'})| \leq \sup_{j \in [m] \setminus \{k\}} |R(H_k, A_{kj}) - R(H_k, A_{jk})| + \alpha. \quad (6)$$

We can now bound the total variation distance between the unknown distribution P and the output $H_{k'}$. Let H_l be any distribution in \mathcal{H} that satisfies $\text{TV}(H_l, P) = \text{OPT}$.⁶ Using the triangle inequality we have,

$$\begin{aligned} 2 \text{TV}(H_{k'}, P) &= \|H_{k'} - P\|_1 \\ &\leq \|H_l - P\|_1 + \|H_{k'} - H_l\|_1. \end{aligned} \quad (7)$$

We now look at the right most term in (7). By the definition of the total variation distance and an application of the triangle inequality we have

$$\begin{aligned} \|H_{k'} - H_l\|_1 &= (H_{k'}(A_{k'l}) - H_l(A_{k'l})) + (H_l(A_{lk'}) - H_{k'}(A_{lk'})) \\ &= |(H_{k'}(A_{k'l}) - H_{k'}(A_{lk'})) + (H_l(A_{lk'}) - H_l(A_{k'l}))| \\ &\leq |R(H_{k'}, A_{k'l}) - R(H_{k'}, A_{lk'})| + |R(H_l, A_{lk'}) - R(H_l, A_{k'l})| \\ &\leq \sup_{j \in [m] \setminus \{k'\}} |R(H_{k'}, A_{k'j}) - R(H_{k'}, A_{jk'})| + \sup_{j \in [m] \setminus \{l\}} |R(H_l, A_{lj}) - R(H_l, A_{jl})|. \end{aligned}$$

6. Note that this implies that $\|H_l - P\|_1 = 2\text{OPT}$.

Using (6), the fact that H_k maximizes the score function, and the triangle inequality all together yields,

$$\begin{aligned}
 \|H_{k'} - H_l\|_1 &\leq \sup_{j \in [m] \setminus \{k\}} |R(H_k, A_{kj}) - R(H_k, A_{jk})| + \sup_{j \in [m] \setminus \{l\}} |R(H_l, A_{lj}) - R(H_l, A_{jl})| + \alpha \\
 &\leq 2 \sup_{j \in [m] \setminus \{l\}} |R(H_l, A_{lj}) - R(H_l, A_{jl})| + \alpha \\
 &\leq 2 \sup_{j \in [m] \setminus \{l\}} \left| (H_l(A_{lj}) - P(A_{lj})) + (P(A_{jl}) - H_l(A_{jl})) \right| \\
 &\quad + 2 \sup_{j \in [m] \setminus \{l\}} \left| (P(A_{lj}) - \widehat{P}(A_{lj}, D)) + (\widehat{P}(A_{jl}, D) - P(A_{jl})) \right| + \alpha.
 \end{aligned}$$

Notice that the first term on the right hand side of the final inequality is at most twice the ℓ_1 distance between H_l and P . Let $\Delta(P) = 2 \sup_{j \in [m] \setminus \{l\}} \left| (P(A_{lj}) - \widehat{P}(A_{lj}, D)) + (\widehat{P}(A_{jl}, D) - P(A_{jl})) \right|$. This gives us

$$\|H_{k'} - H_l\|_1 \leq 2\|H_l - P\|_1 + \Delta(P) + \alpha.$$

Furthermore, notice that the term $\Delta(P)$ is small when the difference between the empirical and the true probability measures assigned by P to the Scheffè sets is small. We can thus upper bound this term by α by using $2\binom{m}{2}$ standard Chernoff bounds together with a union bound to get,

$$\|H_{k'} - H_l\|_1 \leq 2\|H_l - P\|_1 + 2\alpha, \tag{8}$$

with probability no less than $1 - \beta/2$ so long as $n = O\left(\frac{\log(m/\beta)}{\alpha^2}\right)$. Putting (7) and (8) together gives us,

$$\begin{aligned}
 \text{TV}(H_{k'}, P) &\leq \frac{3}{2}\|H_l - P\|_2 + \alpha \\
 &= 3\text{OPT} + \alpha.
 \end{aligned}$$

A union bound together with setting $n = O\left(\frac{\log(m/\beta)}{\alpha^2} + \frac{\log(m/\beta)}{\alpha\varepsilon}\right)$ completes the proof. \blacksquare

6. Conclusion

We provide the first finite sample complexity bounds for privately learning Gaussians with unbounded parameters. We do this via a method for converting small local covers, to global covers which are locally small. In this paper, we only prove sample complexity upper bounds, and our methods are not computational in nature. One natural direction is to design polynomial time algorithms for learning unbounded Gaussians. Another direction is to explore applications of our method to other classes of distributions. The most immediate class that comes to mind is Gaussians mixture models (GMMs) – given upper and lower bounds on the total variation distance between GMMs based on their parameter distance, it should not be difficult to derive corresponding sample complexity bounds.

Acknowledgments

GK would like to thank Mark Bun, Adam Smith, Thomas Steinke, and Zhiwei Steven Wu for helpful conversations and suggestions which led to the results in Section 5. IA was supported by an Ontario Graduate Scholarship. HA was supported by an NSERC Discovery grant and a McMaster University startup grant. GK was supported by an NSERC Discovery grant a Compute Canada RRG grant, and a University of Waterloo startup grant.

References

- Jayadev Acharya, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Sorting with adversarial comparators and application to density estimation. In *Proceedings of the 2014 IEEE International Symposium on Information Theory, ISIT '14*, pages 1682–1686, Washington, DC, USA, 2014. IEEE Computer Society.
- Jayadev Acharya, Moein Falahatgar, Ashkan Jafarpour, Alon Orlitsky, and Ananda Theertha Suresh. Maximum selection and sorting with adversarial comparators. *Journal of Machine Learning Research*, 19(1):2427–2457, 2018.
- Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private assouad, fano, and le cam. *arXiv preprint arXiv:2004.06830*, 2020.
- Ishaq Aden-Ali and Hassan Ashtiani. On the sample complexity of learning sum-product networks. In *International Conference on Artificial Intelligence and Statistics*, pages 4508–4518. PMLR, 2020.
- Martin Anthony. Classification by polynomial surfaces. *Discrete Applied Mathematics*, 61(2):91–103, 1995.
- 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 31, NeurIPS '18*, pages 3412–3421. Curran Associates, Inc., 2018a.
- Hassan Ashtiani, Shai Ben-David, and Abbas Mehrabian. Sample-efficient learning of mixtures. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, AAAI'18*, pages 2679–2686. AAAI Publications, 2018b. URL <https://arxiv.org/abs/1706.01596>.
- Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. *Machine Learning*, 94(3):401–437, 2014.
- Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan Ullman. Coinpress: Practical private mean and covariance estimation. *arXiv preprint arXiv:2006.06618*, 2020.
- Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th ACM Symposium on Operating Systems Principles, SOSP '17*, pages 441–459, New York, NY, USA, 2017. ACM.

- Olivier Bousquet, Daniel M. Kane, and Shay Moran. The optimal approximation factor in density estimation. In *Proceedings of the 32nd Annual Conference on Learning Theory, COLT '19*, pages 318–341, 2019.
- Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Proceedings of the 14th Conference on Theory of Cryptography, TCC '16-B*, pages 635–658, Berlin, Heidelberg, 2016. Springer.
- Mark Bun and Thomas Steinke. Average-case averages: Private algorithms for smooth sensitivity and mean estimation. In *Advances in Neural Information Processing Systems 32, NeurIPS '19*, pages 181–191. Curran Associates, Inc., 2019.
- Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Proceedings of the 46th Annual ACM Symposium on the Theory of Computing, STOC '14*, pages 1–10, New York, NY, USA, 2014. ACM.
- Mark Bun, Kobbi Nissim, and Uri Stemmer. Simultaneous private learning of multiple concepts. In *Proceedings of the 7th Conference on Innovations in Theoretical Computer Science, ITCS '16*, pages 369–380, New York, NY, USA, 2016. ACM.
- Mark Bun, Thomas Steinke, and Jonathan Ullman. Make up your mind: The price of online queries in differential privacy. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '17*, pages 1306–1325, Philadelphia, PA, USA, 2017. SIAM.
- Mark Bun, Cynthia Dwork, Guy N. Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM Symposium on the Theory of Computing, STOC '18*, pages 74–86, New York, NY, USA, 2018. ACM.
- Mark Bun, Gautam Kamath, Thomas Steinke, and Zhiwei Steven Wu. Private hypothesis selection. In *Advances in Neural Information Processing Systems 32, NeurIPS '19*, pages 156–167. Curran Associates, Inc., 2019.
- 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.
- Aref N. Dajani, Amy D. Lauger, Phyllis E. Singer, Daniel Kifer, Jerome P. Reiter, Ashwin Machanavajjhala, Simson L. Garfinkel, Scot A. Dahl, Matthew Graham, Vishesh Karwa, Hang Kim, Philip Lelerc, Ian M. Schmutte, William N. Sexton, Lars Vilhuber, and John M. Abowd. The modernization of statistical disclosure limitation at the U.S. census bureau, 2017. Presented at the September 2017 meeting of the Census Scientific Advisory Committee.
- 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.
- Constantinos Daskalakis, Ilias Diakonikolas, and Rocco A. Servedio. Learning Poisson binomial distributions. In *Proceedings of the 44th Annual ACM Symposium on the Theory of Computing, STOC '12*, pages 709–728, New York, NY, USA, 2012. ACM.

- Luc Devroye and Gábor Lugosi. A universally acceptable smoothing factor for kernel density estimation. *The Annals of Statistics*, 24(6):2499–2512, 1996.
- Luc Devroye and Gábor Lugosi. Nonasymptotic universal smoothing factors, kernel complexity and Yatracos classes. *The Annals of Statistics*, 25(6):2626–2637, 1997.
- Luc Devroye and Gábor Lugosi. *Combinatorial methods in density estimation*. Springer, 2001.
- Luc Devroye, Abbas Mehrabian, and Tommy Reddad. The total variation distance between high-dimensional Gaussians. *arXiv preprint arXiv:1810.08693*, 2018.
- 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.
- Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '16, pages 655–664, Washington, DC, USA, 2016. IEEE Computer Society.
- Differential Privacy Team, Apple. Learning with privacy at scale. <https://machinelearning.apple.com/docs/learning-with-privacy-at-scale/appledifferentialprivacysystem.pdf>, December 2017.
- Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems 30*, NIPS '17, pages 3571–3580. Curran Associates, Inc., 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.
- Wenxin Du, Canyon Foot, Monica Moniot, Andrew Bray, and Adam Groce. Differentially private confidence intervals. *arXiv preprint arXiv:2001.02285*, 2020.
- Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the 41st Annual ACM Symposium on the Theory of Computing*, STOC '09, pages 371–380, New York, NY, USA, 2009. ACM.
- Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, EUROCRYPT '06, pages 486–503, Berlin, Heidelberg, 2006a. Springer.
- 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, 2006b. Springer.

- 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.
- Cynthia Dwork, Adam Smith, Thomas Steinke, and Jonathan Ullman. Exposed! a survey of attacks on private data. *Annual Review of Statistics and Its Application*, 4(1):61–84, 2017.
- 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.
- Sivakanth Gopi, Gautam Kamath, Janardhan Kulkarni, Aleksandar Nikolov, Zhiwei Steven Wu, and Huanyu Zhang. Locally private hypothesis selection. In *Proceedings of the 33rd Annual Conference on Learning Theory, COLT '20*, 2020.
- Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd Annual ACM Symposium on the Theory of Computing, STOC '10*, pages 705–714, New York, NY, USA, 2010. ACM.
- Nils Homer, Szabolcs Szelinger, Margot Redman, David Duggan, Waibhav Tembe, Jill Muehling, John V. Pearson, Dietrich A. Stephan, Stanley F. Nelson, and David W. Craig. Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genetics*, 4(8):1–9, 2008.
- Gautam Kamath and Jonathan Ullman. A primer on private statistics. *arXiv preprint arXiv:2005.00010*, 2020.
- Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Proceedings of the 32nd Annual Conference on Learning Theory, COLT '19*, pages 1853–1902, 2019a.
- Gautam Kamath, Or Sheffet, Vikrant Singhal, and Jonathan Ullman. Differentially private algorithms for learning mixtures of separated Gaussians. In *Advances in Neural Information Processing Systems 32, NeurIPS '19*, pages 168–180. Curran Associates, Inc., 2019b.
- Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Proceedings of the 33rd Annual Conference on Learning Theory, COLT '20*, 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, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. Releasing search queries and clicks privately. In *Proceedings of the 18th International World Wide Web Conference, WWW '09*, pages 171–180, New York, NY, USA, 2009. ACM.
- Yuhan Liu, Ananda Theertha Suresh, Felix Yu, Sanjiv Kumar, and Michael Riley. Learning discrete distributions: User vs item-level privacy. *arXiv preprint arXiv:2007.13660*, 2020.

- Satyaki Mahalanabis and Daniel Stefankovic. Density estimation in linear time. In *Proceedings of the 21st Annual Conference on Learning Theory, COLT '08*, pages 503–512, 2008.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS '07*, pages 94–103, Washington, DC, USA, 2007. IEEE Computer Society.
- Kobbi Nissim and Uri Stemmer. Clustering algorithms for the centralized and local models. In *Algorithmic Learning Theory, ALT '18*, pages 619–653. JMLR, Inc., 2018.
- Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39th Annual ACM Symposium on the Theory of Computing, STOC '07*, pages 75–84, New York, NY, USA, 2007. ACM.
- Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Locating a small cluster privately. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS '16*, pages 413–427, New York, NY, USA, 2016. ACM.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *Proceedings of the 38th IEEE Symposium on Security and Privacy, SP '17*, pages 3–18, Washington, DC, USA, 2017. IEEE Computer Society.
- Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the 43rd Annual ACM Symposium on the Theory of Computing, STOC '11*, pages 813–822, New York, NY, USA, 2011. ACM.
- 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, Washington, DC, USA, 2017b. IEEE Computer Society.
- 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.
- Michel Talagrand. Sharper bounds for gaussian and empirical processes. *The Annals of Probability*, pages 28–76, 1994.
- Gregory Valiant and Paul Valiant. A CLT and tight lower bounds for estimating entropy. *Electronic Colloquium on Computational Complexity (ECCC)*, 17(179), 2010.
- Vladimir Naumovich Vapnik and Alexey Yakovlevich Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.

Yannis G. Yatracos. Rates of convergence of minimum distance estimators and Kolmogorov's entropy. *The Annals of Statistics*, 13(2):768–774, 1985.

Huanyu Zhang, Gautam Kamath, Janardhan Kulkarni, and Zhiwei Steven Wu. Privately learning Markov random fields. In *Proceedings of the 37th International Conference on Machine Learning*, ICML '20. JMLR, Inc., 2020.

Appendix A. Useful Inequalities

Proposition 44 *Let X and Y be random variables taking values in the same set. For any function f , we have $\text{TV}(f(X), f(Y)) \leq \text{TV}(X, Y)$.*

Proof For any set A we have,

$$\Pr[f(X) \in A] - \Pr[f(Y) \in A] = \Pr[X \in f^{-1}(A)] - \Pr[Y \in f^{-1}(A)] \leq \text{TV}(X, Y).$$

Taking the supremum of the left hand side completes the proof. \blacksquare

Corollary 45 *Let X and Y be random variables taking values in the same set. For any invertible function f , we have $\text{TV}(f(X), f(Y)) = \text{TV}(X, Y)$.*

Proof By Proposition 44, $\text{TV}(f(X), f(Y)) \leq \text{TV}(X, Y)$ and $\text{TV}(f^{-1}(f(X)), f^{-1}(f(Y))) \leq \text{TV}(f(X), f(Y))$. \blacksquare

Proposition 46 *For any $\xi \in (0, 1)$, Gaussian $\mathcal{N}(\mu, \Sigma)$ and distributions P that satisfies $\text{TV}(P, \mathcal{N}(\mu, \Sigma)) \leq \xi$, the following holds. If $X_1, X_2 \sim P^2$, and $Y = (X_1 - X_2)/\sqrt{2} \sim Q$, then $\text{TV}(Q, \mathcal{N}(0, \Sigma)) \leq 3\xi$.*

Proof Given $X \sim P$, the density P is given by $P = \mathcal{N}(\mu, \Sigma) + \Delta$.⁷ It follows from the definition of the TV distance that $\|\Delta\|_1 \leq 2\xi$. Given a sample $X \sim P$, we let P^- be the distribution that satisfies $-X \sim P^-$. The density of P^- is given by $P^- = \mathcal{N}(-\mu, \Sigma) + \Delta^-$ where $\|\Delta^-\|_1 \leq 2\xi$. The sample $Y' = X_1 + (-X_2)$ has density

$$\begin{aligned} P_{\text{conv}} &= P \otimes P^- = (\mathcal{N}(\mu, \Sigma) + \Delta) \otimes (\mathcal{N}(-\mu, \Sigma) + \Delta^-) \\ &= \mathcal{N}(0, 2\Sigma) + \mathcal{N}(\mu, \Sigma) \otimes \Delta^- + \Delta \otimes \mathcal{N}(-\mu, \Sigma) + \Delta \otimes \Delta^-. \end{aligned}$$

We can now bound the TV distance between P_{conv} and $\mathcal{N}(0, 2\Sigma)$.

$$\begin{aligned} \text{TV}(P_{\text{conv}}, \mathcal{N}(0, 2\Sigma)) &= \frac{1}{2} \|P_{\text{conv}} - \mathcal{N}(0, 2\Sigma)\|_1 \\ &= \frac{1}{2} \|\mathcal{N}(0, 2\Sigma) + \mathcal{N}(\mu, \Sigma) \otimes \Delta^- + \Delta \otimes \mathcal{N}(-\mu, \Sigma) + \Delta \otimes \Delta^- - \mathcal{N}(0, 2\Sigma)\|_1 \\ &= \frac{1}{2} \|\mathcal{N}(\mu, \Sigma) \otimes \Delta^- + \Delta \otimes \mathcal{N}(-\mu, \Sigma) + \Delta \otimes \Delta^-\|_1 \\ &\leq \frac{1}{2} (\|\mathcal{N}(\mu, \Sigma)\|_1 \|\Delta^-\|_1 + \|\Delta\|_1 \|\mathcal{N}(-\mu, \Sigma)\|_1 + \|\Delta\|_1 \|\Delta^-\|_1) \\ &\leq \frac{1}{2} (2\xi + 4\xi^2) \leq 3\xi, \end{aligned}$$

7. Recall in this paper we define distributions by their densities so $\mathcal{N}(\mu, \Sigma)$ is the Gaussian density.

where the first inequality follows from the triangle inequality together with Young's inequality. Since $\text{TV}(P_{\text{conv}}, \mathcal{N}(0, 2\Sigma)) \leq 3\xi$, the statement follows immediately from Corollary 45 and equation (1). \blacksquare

Appendix B. Omitted proofs from Section 2

B.1. Proof of Lemma 12

We first prove the inequality on the left hand side. Let \mathcal{C}_γ be a γ -cover for \mathcal{H} that has size $N(\mathcal{H}, \gamma)$. If $N(\mathcal{H}, \gamma) = \infty$, we are done. Otherwise, we claim there is no 2γ -packing $\mathcal{P}'_{2\gamma}$ of \mathcal{H} of size at least $N(\mathcal{H}, \gamma) + 1$. We prove this by contradiction. Assume to the contrary that there exists a 2γ -packing $\mathcal{P}'_{2\gamma}$ of \mathcal{H} such that $|\mathcal{P}'_{2\gamma}| = N(\mathcal{H}, \gamma) + 1$. By the pigeonhole principle, there exists $Q \in \mathcal{C}_\gamma$ and two distributions $P, P' \in \mathcal{P}'_{2\gamma}$ such that $\text{TV}(Q, P) \leq \gamma$ and $\text{TV}(Q, P') \leq \gamma$. By the triangle inequality it follows that $\text{TV}(P, P') \leq 2\gamma$ which is a contradiction, so $\mathcal{P}'_{2\gamma}$ cannot be a 2γ -packing of \mathcal{H} . This shows that $M(\mathcal{H}, 2\gamma) \leq N(\mathcal{H}, \gamma)$.

We now prove the inequality on the right hand side. Let \mathcal{P}_γ be a maximal γ -packing with size $M(\mathcal{H}, \gamma)$. If $M(\mathcal{H}, \gamma) = \infty$, we are done. Otherwise, we claim that \mathcal{P}_γ is also an γ -cover of \mathcal{H} , and hence $N(\mathcal{H}, \gamma) \leq M(\mathcal{H}, \gamma)$. We can prove this by contradiction. Suppose to the contrary that there were a distribution $P \in \mathcal{H}$ with $\text{TV}(P, \mathcal{P}_\gamma) > \gamma$. Then we could add P to \mathcal{P}_γ to produce a strictly larger packing, contradicting the maximality of \mathcal{P}_γ . Therefore it only remains to show that a maximal packing actually exists, which follows from a simple application of Zorn's lemma (See proof of Lemma 29).

B.2. Proof of Lemma 28

Fix accuracy and privacy parameters $\alpha, \beta, \varepsilon, \delta \in (0, 1)$. Let $T = \lceil \log_2(1/\alpha) \rceil$ and define sequences $\xi_1 = \alpha/12C, \xi_2 = 2\alpha/12C, \dots, \xi_{T+4} = 2^{T+3}\alpha/12C$. Furthermore for all $t \in [T+4]$ set $\alpha_t = \alpha/12, \beta_t = \beta/2(T+4), \varepsilon_t = \varepsilon/2(T+4)$ and $\delta_t = \delta/(T+4)$. For each t , let H_t denote the outcome of a run of the (ξ, C) -robust algorithm using robustness parameter ξ_t accuracy parameters α_t, β_t and privacy parameters ε_t, δ_t . We then use the algorithm of Theorem 27 to select a hypothesis from H_1, \dots, H_{T+4} using accuracy parameter $\alpha/2, \beta/2$ and privacy parameter $\varepsilon/2$. By composition of DP (Lemma 20) it follows the output of the two step procedure is (ε, δ) -DP.

We can now argue about the accuracy of the algorithm. Given $\tilde{n}_{\mathcal{H}}^C \left(\frac{\alpha}{12}, \frac{\beta}{2(T+4)}, \frac{\varepsilon}{2(T+4)}, \frac{\delta}{T+4} \right)$ samples, a union bound guarantees that with probability at least $1 - \beta/2$ each H_t is accurate (assuming ξ_t is a correct upper bound for OPT). We condition on this event.

The success of run t of the robust PAC learner implies that $\text{OPT} \in (\xi_{t-1}, \xi_t]$, so $\text{TV}(P, H_t) \leq C\xi_t + \alpha_t \leq 2C \cdot \text{OPT} + \alpha/12$. Similarly, if $\text{OPT} \leq \xi_1$ then H_1 satisfies $\text{TV}(P, H_1) \leq C\xi_1 + \alpha_1 = \alpha/12 + \alpha/12 = \alpha/6$. Finally, if $\text{OPT} > \xi_{T+4}$ this implies $\text{OPT} > 8/12C$, so any distribution H' satisfies $\text{TV}(P, H') \leq 2C \cdot \text{OPT}$ trivially. So regardless of OPT, there is a run t that satisfies $\text{TV}(P, H_t) \leq 2C \cdot \text{OPT} + \alpha/6$. Finally, Theorem 27 guarantees that, with probability greater than $1 - \beta/2$, the second step in our above procedure will output a distribution \hat{H} such that $\text{TV}(P, \hat{H}) \leq 3(2C \cdot \text{OPT} + \alpha/6) + \alpha/2 = 6C \cdot \text{OPT} + \alpha$, as long as $n = O \left(\frac{\log(T/\beta)}{\alpha^2} + \frac{\log(T/\beta)}{\alpha\varepsilon} \right)$. To complete the proof, we use a union bound together with setting

$$n = \tilde{n}_{\mathcal{H}}^C \left(\frac{\alpha}{12}, \frac{\beta}{2(T+4)}, \frac{\varepsilon}{2(T+4)}, \frac{\delta}{T+4} \right) + O \left(\frac{\log(T/\beta)}{\alpha^2} + \frac{\log(T/\beta)}{\alpha\varepsilon} \right).$$

