

Differentially Private Multivariate Medians

Kelly Ramsay

KRAMSAY2@YORKU.CA

*Department of Mathematics and Statistics
York University
North York, ON M3J1P3, Canada*

Aukosh Jagannath

A.JAGANNATH@UWATERLOO.CA

*Department of Statistics and Actuarial Science
University of Waterloo
Waterloo, ON N2L3G1, Canada*

Shoja'eddin Chenouri

SCHENOURI@UWATERLOO.CA

*Department of Statistics and Actuarial Science
University of Waterloo
Waterloo, ON N2L3G1, Canada*

Editor: Raef Bassily

Abstract

Statistical tools which satisfy rigorous privacy guarantees are necessary for modern data analysis. It is well-known that robustness against contamination is linked to differential privacy. Despite this fact, using multivariate medians for differentially private and robust multivariate location estimation has not been systematically studied. We develop novel finite-sample performance guarantees for differentially private multivariate depth-based medians, which are essentially sharp. Our results cover commonly used depth functions, such as the halfspace (or Tukey) depth, spatial depth, and the integrated dual depth. We show that under Cauchy marginals, the cost of heavy-tailed location estimation outweighs the cost of privacy. We demonstrate our results numerically using a Gaussian contamination model in dimensions up to $d = 100$, and compare them to a state-of-the-art private mean estimation algorithm. As a by-product of our investigation, we prove concentration inequalities for the output of the exponential mechanism about the maximizer of the population objective function. This bound applies to objective functions that satisfy a mild regularity condition.

Keywords: differential privacy, location estimation, robust statistics, depth function, sample complexity

1. Introduction

Protecting user privacy is necessary for a safe and fair society. In order to protect user privacy, many large institutions, including the United States Census Bureau (Abowd et al., 2022), Apple (Differential Privacy Team, 2017), and Google (Guevara, 2021), utilize a strong notion of privacy known as differential privacy. Differential privacy is favored because it protects against many adversarial attacks while requiring minimal assumptions on both the data and the adversary (Dwork et al., 2017). It is now well-known that differential privacy is linked to robustness against contamination, see (Dwork and Lei, 2009; Avella-

Medina, 2020; Liu et al., 2021; Georgiev and Hopkins, 2022; Hopkins et al., 2023; Asi et al., 2023) and the references therein. Surprisingly, however, multivariate medians have not been systematically studied in the differential privacy literature. We study this problem here. Note that there are several variants of differential privacy. Here, we focus on the setting of pure differential privacy, which is the strictest variant of differential privacy.

To date, the literature on differentially private location estimation has largely focused on means. Foygel Barber and Duchi (2014) proved a lower bound on the minimax risk of differentially private mean estimation in terms of the number of moments possessed by the population measure. Other early works focused on the univariate setting (Karwa and Vadhan, 2017; Bun and Steinke, 2019). Kamath et al. (2019) and Bun et al. (2019) studied differentially private mean estimation for several parametric models, including the multivariate Gaussian model. Later, several authors studied differentially private mean estimation for sub-Gaussian measures (Cai et al., 2021; Biswas et al., 2020; Brown et al., 2021; Liu et al., 2022). After which, Kamath et al. (2020); Liu et al. (2021) and Hopkins et al. (2022) studied differentially private mean estimation for measures with a finite second moment. Within this framework, Kamath et al. (2020) considered heavy tails and Liu et al. (2021) and Hopkins et al. (2022) considered robustness in terms of contamination models. Kamath et al. (2020) quantified the sample complexity of their (purely) differentially private mean estimator in terms of the number of moments possessed by the population measure. Liu et al. (2021) quantified the relationship between the accuracy of an (approximately) differentially private mean estimator and the level of contamination in the observed sample. On a similar note, Hopkins et al. (2022) showed that there exists a level of contamination under which the sample complexity of their (purely) differentially private mean estimator remains unchanged.

On the other hand, if one wishes to estimate location robustly, the canonical estimators are medians. The differential privacy literature on medians has focused on the univariate setting: Dwork and Lei (2009) introduced an approximately differentially private median estimator with asymptotic consistency guarantees, Avella-Medina and Brunel (2019); Brunel and Avella-Medina (2020) then introduced several median estimators which achieve sub-Gaussian error rates, and Tzamos et al. (2020) introduced median estimators with minimax optimal sample complexity. Private estimation of multivariate medians, however, has not been carefully studied.

In this paper, we study differentially private, multivariate median estimation through the framework of depth functions. Depth-based median estimation is the standard approach to multivariate median estimation. Here, one has a function, called the depth function, which provides a measure of centrality. Its maximizer is then the median, which is generally robust. Popular depth-based medians include the halfspace (or Tukey) median (Tukey, 1974), the spatial median (Vardi and Zhang, 2000), and the simplicial median (Liu, 1990). For more on the motivation of the depth-based medians approach, see Section 2 below and Small (1990); Vardi and Zhang (2000); Serfling (2006). Our results apply to a broad class of depth functions, including those listed above, and we make minimal distributional assumptions. In particular, we do not require concavity of the depth function or that the population measure has moments of any order.

Before turning to our main results, it is important to note that several authors have recently used depth functions in the context of differential privacy, likely due to their ro-

bustness properties, with a heavy focus on the halfspace depth. Depth functions have been used in settings such as finding a point in a convex hull (Beimel et al., 2019; Gao and Sheffet, 2020), comparing k -norm mechanisms (Awan and Slavković, 2021), and improving robustness for mean estimation for Gaussian models (Brown et al., 2021; Liu et al., 2021). Another related work is that of Ben-Eliezer et al. (2022), who focus on various problems concerning high-dimensional quantile estimation. In particular, one problem Ben-Eliezer et al. (2022) consider is that of privately obtaining one or more points which have halfspace depth above a given threshold. Though this problem is closely related, it is distinct. For instance, their algorithm can be used to estimate a point with high halfspace depth. However, it could not be used to directly estimate the halfspace median, since that would require setting the threshold to be close to the depth of the empirical halfspace median, which depends on the sample. To our knowledge, none of the aforementioned works have directly addressed depth-based median estimation. After completion of this work, Ramsay and Spicker (2025) consider approximately differentially private projection-depth-based medians. Though many popular depth-based medians are covered by our framework, projection-depth-based medians are not covered by it. As a result, their work can be thought of as complementary to this one.

1.1 Our contributions

Our contributions are as follows:

- Our main result is a general finite-sample deviations bound for private multivariate medians based on the exponential mechanism (Theorem 5). This deviations bound applies to private versions of many popular medians arising from depth functions, such as the halfspace median (Tukey, 1974), the simplicial median (Liu, 1988) and the spatial median (Vardi and Zhang, 2000; Serfling, 2002).
- We use Theorem 5 to give upper bounds on the deviations (and sample complexities, see Corollary 10) of several, purely differentially private multivariate medians arising from depth functions. These bounds are essentially sharp given recent results of Kamath et al. (2019); Cai et al. (2021) and do not require any moment assumptions on the population measure.
- We provide a fast implementation of a smoothed version of the integrated dual depth-based median (Cuevas and Fraiman, 2009); we can compute the (non-private) median of ten thousand 100-dimensional samples in less than one second on a personal computer. We show that this smoothed version of the integrated dual depth satisfies desirable properties for a depth function and can be approximated in polynomial time with finite-sample performance guarantees, see Proposition 33.
- As a by-product of our analysis, we uncover a general but elementary concentration bound (Theorem 12) for the exponential mechanism. We give a novel regularity condition “ (K, \mathcal{F}) -regularity” (Condition 3) on the objective function. This regularity condition implies both an upper bound on the sample complexity of a draw from the exponential mechanism and an upper bound on the global sensitivity of the objective function. This bound is applicable to a broader class of problems in statistical learning beyond differential privacy, see the follow-up paper (Ramsay et al., 2025).

- As a second by-product of our analysis, we uncover uniform non-asymptotic error bounds for several depth functions and a differentially private estimator of data depth values.

2. Main results

In this section, we provide a general finite-sample deviation bound for differentially private multivariate median estimation. As the tools we use here combine ideas from two distinct literatures, namely the differential privacy literature and the robust statistics literature, we briefly recall some essential notions from each of these fields before stating our main results. For a more in-depth introduction to differential privacy see Dwork and Roth (2014), and for a more in-depth introduction to robust multivariate median estimation and the depth function framework see Mosler (2002); Small (1990); Liu et al. (2006). We define all mathematical objects throughout the paper as they are used, but it is also helpful to have them collected in one place. Therefore, we also provide a summary of the notation we will use throughout the paper in Appendix A.

2.1 Differential privacy

Let us start first by recalling some essential notions from differential privacy. A *data set* of size $n \times d$ is a collection of n points in \mathbb{R}^d , $\mathbb{X}_n = (X_\ell)_{\ell=1}^n$, with repetitions allowed. Let $\mathbf{D}_{n \times d}$ be the set of all data sets of size $n \times d$. For a data set \mathbb{Y}_n , let $\mathcal{D}(\mathbb{Y}_n, m) = \{\mathbb{Z}_n \in \mathbf{D}_{n \times d} : |\mathbb{Z}_n \triangle \mathbb{Y}_n| = 2m\}$, that is, $\mathcal{D}(\mathbb{Y}_n, m)$ is the collection of data sets of size $n \times d$ which differ from \mathbb{Y}_n in exactly m points. Two data sets \mathbb{Y}_n and \mathbb{Z}_n are said to be *adjacent* if $\mathbb{Z}_n \in \mathcal{D}(\mathbb{Y}_n, 1)$. Finally, for a data set \mathbb{Y}_n with empirical measure $\hat{\mu}_{\mathbb{Y}_n}$ define

$$\widetilde{\mathcal{M}}(\mathbb{Y}_n) = \{\tilde{\mu} \in \mathcal{M}_1(\mathbb{R}^d) : \tilde{\mu} = \hat{\mu}_{\mathbb{Z}_n} \text{ for some } \mathbb{Z}_n \in \mathcal{D}(\mathbb{Y}_n, 1)\},$$

where, for a space S , $\mathcal{M}_1(S)$ denotes the space of probability measures on S .

Let us now recall the notion of differential privacy. Suppose that we observe a data set \mathbb{X}_n comprised of n i.i.d. samples from some unknown measure $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ and produce a statistic, $\hat{\theta}_n$, whose law conditionally on the samples depends on their empirical measure, $\hat{\mu}_n$. Let us denote this law by $P_{\hat{\mu}_n}$. The goal of differential privacy is to produce a statistic such that $P_{\hat{\mu}_n}$ is non-degenerate and, more precisely, obeys a certain “privacy guarantee” which is defined as follows. Following the privacy literature, we call the algorithm that takes $\hat{\mu}_n$ and returns the (random) statistic, $\hat{\theta}_n$, a *mechanism*. We also follow the standard abuse of terminology and call $\hat{\theta}_n$ the mechanism when it is clear from context.

Definition 1 A mechanism $\hat{\theta}_n$ is ϵ -differentially private if for any data set \mathbb{Y}_n , any $\tilde{\mu} \in \widetilde{\mathcal{M}}(\mathbb{Y}_n)$, and any measurable set B , it holds that

$$P_{\hat{\mu}_{\mathbb{Y}_n}}(B) \leq e^\epsilon P_{\tilde{\mu}}(B). \quad (2.1)$$

Here, $\epsilon > 0$ is the *privacy parameter*, for which smaller values enforce stricter levels of privacy. Many general purpose differentially private mechanisms rely on the concept of *global sensitivity*. The global sensitivity GS_n of a function $\phi : \mathbb{R}^d \times \mathcal{M}_1(\mathbb{R}^d) \rightarrow \mathbb{R}$ is given

by

$$\text{GS}_n(\phi) = \sup_{\mathbb{X}_n \in \mathbf{D}_{n \times d}} \sup_{x \in \mathbb{R}^d} |\phi(x, \hat{\mu}_n) - \phi(x, \tilde{\mu}_n)|.$$

$$\tilde{\mu}_n \in \widetilde{\mathcal{M}}(\mathbb{X}_n)$$

The simplest differentially private mechanisms are the additive noise mechanisms, e.g., the Laplace and Gaussian mechanisms. These mechanisms require the non-private estimator to have finite global sensitivity. Unfortunately, virtually all standard location estimators, including both univariate and multivariate medians, have infinite global sensitivity, see e.g., (Avella-Medina and Brunel, 2019).

However, *depth functions*, the objective functions for multivariate medians which we will discuss shortly, all have finite global sensitivity. This fact makes them a good candidate for use with the *exponential mechanism*, which is a general mechanism used to produce differentially private estimators from non-private estimators that are maximizers of a given objective function. For a given $\beta > 0$, base measure (i.e., prior) $\pi \in \mathcal{M}_1(\mathbb{R}^d)$ and objective function $\phi : \mathbb{R}^d \times \mathcal{M}_1(\mathbb{R}^d) \rightarrow \mathbb{R}$, suppose that $\int \exp(\beta\phi(\theta, \nu)) d\pi(\theta) < \infty$ for all $\nu \in \mathcal{M}_1(\mathbb{R}^d)$, then the exponential mechanism with base measure π is given by

$$Q_{\hat{\mu}_n, \beta} \propto \exp(\beta\phi(\theta, \hat{\mu}_n)) d\pi. \quad (2.2)$$

A sample from the exponential mechanism $\tilde{\theta}_n \sim Q_{\hat{\mu}_n, \beta}$ then provides an estimate of the population value $\theta_0 = \arg\max_{\mathbb{R}^d} \phi(\theta, \mu)$. In regard to the choice of β , it was shown by McSherry and Talwar (2007) that if $\tilde{\theta}_n$ is produced by the exponential mechanism, with $\beta \leq \epsilon/2\text{GS}_n(\phi)$ then $\tilde{\theta}_n$ is ϵ -differentially private.

2.2 The depth function framework

When first studying multivariate median estimation, one might try to naively extend the notion of a univariate median to the multivariate setting, i.e., the coordinate-wise median. This, however, is well-known to be an unsatisfactory as a measure of center for many reasons (Small, 1990; Serfling, 2006; Liu et al., 2006). For example, the coordinate-wise median can reside outside the convex hull of the data (Serfling, 2006). The general framework of *depth functions*, or statistical depth functions, was introduced to resolve these issues and is now the standard approach to defining and studying multivariate medians. See (Small, 1990; Zuo and Serfling, 2000; Mosler, 2002; Serfling, 2006; Liu et al., 2006) for a necessarily small introduction to this broad research field.

Let us now recall the basic notions of (statistical) depth functions. Roughly speaking, depth functions are functions of the form $D : \mathbb{R}^d \times \mathcal{M}_1(\mathbb{R}^d) \rightarrow \mathbb{R}^+$ which, given a probability measure and a point in the domain, assign a number. This number represents how central that point is with respect to the measure; it is called the depth of that point. Given a depth function, the corresponding median is then defined to be a maximizer of this depth, that is, the “deepest” point:

$$\text{MED}(\mu; D) = \arg\max_{x \in \mathbb{R}^d} D(x, \mu).$$

Depth-based medians are generally robust, in the sense that they are not affected by outliers. For instance, depth-based medians have a high breakdown point and favorable properties related to the influence function (Chen and Tyler, 2002; Zuo, 2004). Let us now be more

precise. For two random vectors X, Y , we write $X \stackrel{d}{=} Y$ when X is equal in distribution to Y . Recall that a measure $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is centrally symmetric about a point $x \in \mathbb{R}^d$ if $X - x \stackrel{d}{=} x - X$ for $X \sim \mu$. In the following, let $\mathcal{C}_x \subset \mathcal{M}_1(\mathbb{R}^d)$ denote the set of centrally symmetric measures about x and for $A \in \mathbb{R}^{d \times d}$ and $b \in \mathbb{R}^d$, let $A\mu + b$ be the law of $AX + b$ if $X \sim \mu$.

Definition 2 *A function $D: \mathbb{R}^d \times \mathcal{M}_1(\mathbb{R}^d) \rightarrow \mathbb{R}^+$ is a depth function with admissible set $\mathcal{A} \subseteq \mathcal{M}_1(\mathbb{R}^d)$ if, for all $\mu \in \mathcal{A}$, the following four properties hold:*

1. *Similarity invariance: For all orthogonal matrices $A \in \mathbb{R}^{d \times d}$ and $b \in \mathbb{R}^d$ it holds that $D(x, \mu) = D(Ax + b, A\mu + b)$.*
2. *Maximality at center: If $\mu \in \mathcal{C}_{\theta_0}$ then $D(\theta_0, \mu) = \sup_x D(x, \mu)$.*
3. *Decreasing along rays: Suppose D is maximized at θ_0 . For all $p \in (0, 1)$, it holds that $D(x, \mu) \leq D(p\theta_0 + (1 - p)x, \mu) \leq D(\theta_0, \mu)$.*
4. *Vanishing at infinity: $\lim_{c \rightarrow \infty} D(cu, \mu) = 0$ for any unit vector u .*

The first property says that the depth function has limited dependence on the chosen coordinate system. The remaining properties ensure that D measures depth. The admissible set \mathcal{A} is the set of measures or distributions for which the function measures depth.

Remark 3 *There are a few distinct definitions of the notion of a depth function in the statistics literature. See, e.g., (Zuo and Serfling, 2000; Liu et al., 2006; Ramsay et al., 2019) for some alternatives. We choose Definition 2 as it includes as many existing depth functions as possible while retaining the spirit of a depth function. We use here a weaker definition of a depth function than the influential work of Zuo and Serfling (2000), which replaces similarity invariance with affine invariance. For a discussion of desirable properties for a depth function to possess, see (Zuo and Serfling, 2000; Liu et al., 2006; Serfling, 2019).*

Examples of commonly used depth functions include the halfspace depth (or Tukey depth) (Tukey, 1974), the simplicial depth (Liu, 1988), the spatial depth (Vardi and Zhang, 2000; Serfling, 2002), the integrated dual depth (IDD) (Cuevas and Fraiman, 2009), and the integrated-rank-weighted (IRW) depth (Ramsay et al., 2019). For the reader's convenience, we provide a brief review of these depth functions and their basic properties in Section 6. We emphasize that this is a small selection of the large number of depth functions used in the literature, and reviewing them all is outside of the scope of this work.

There are a number of desirable properties of a depth function, relating to transformation invariance, robustness, computability, ability to measure depth for a broad range of distributions and more, from which no depth function arises as the “gold standard” (Zuo and Serfling, 2000; Liu et al., 2006). Each depth function yields a distinct notion of a median, each with its own desirable properties. As our goal in this paper is to unify these two frameworks, we will not investigate the question of which depth is optimal, statistically or computationally, for a given problem. Instead, our main contribution is to give general finite-sample deviations bounds which applies to many of the existing depth functions, under mild conditions on the population measure.

2.3 A finite-sample deviations bound

Our main result relies on a regularity condition, for which we need to introduce the following notation. Let \mathcal{B} denote the space of Borel functions from \mathbb{R}^d to $[0, 1]$. For a family of functions, $\mathcal{F} \subseteq \mathcal{B}$, define a pseudometric on $\mathcal{M}_1(\mathbb{R}^d)$, $d_{\mathcal{F}}(\mu, \nu) = \sup_{g \in \mathcal{F}} |\int g d(\mu - \nu)|$, where $\mu, \nu \in \mathcal{M}_1(\mathbb{R}^d)$. Recall that many standard metrics on probability measures can be written in this fashion. For example, taking \mathcal{F} to be the set of indicator functions of Borel sets yields the Total Variation distance and taking \mathcal{F} to be the set of indicator functions of semi-infinite rectangles yields the Kolmogorov—Smirnov distance. We then define the following important regularity condition.

Definition 4 *We say that D is (K, \mathcal{F}) -regular if there exists a class of functions $\mathcal{F} \subset \mathcal{B}$ such that $D(x, \cdot)$ is K -Lipschitz with respect to the \mathcal{F} -pseudometric uniformly in x , i.e., for all $\mu, \nu \in \mathcal{M}_1(\mathbb{R}^d)$*

$$\sup_x |D(x, \mu) - D(x, \nu)| \leq K d_{\mathcal{F}}(\mu, \nu).$$

Intuitively, (K, \mathcal{F}) -regularity can be thought of as a robustness condition as it says that (K, \mathcal{F}) -regular functions are affected by extreme observations in a very limited manner. For example, it immediately implies the bound $\text{GS}_n(D) \leq K/n$ on the global sensitivity through the boundedness of $g \in \mathcal{F}$, see Lemma 37. This condition is very convenient for two reasons: Firstly, it is elementary to check this condition for many popular depth functions (see, Table 2.4). Secondly, it immediately yields an ϵ -differentially private estimator by choosing $\beta = n\epsilon/2K$, see Lemma 36. While this assumption is convenient for differential privacy, we show in a companion paper (Ramsay et al., 2025) that our main deviations bound holds under a weaker condition and can apply to other statistical learning settings.

We now state the necessary regularity conditions on the pair (μ, D) that are required for our main result. In the following, let $\text{VC}(\mathcal{F})$ denote the Vapnik—Chervonenkis dimension of \mathcal{F} .

Condition 1 *The function $D(\cdot, \mu)$ has a maximizer.*

Condition 2 *The map $x \mapsto D(x, \mu)$ is L -Lipschitz a.e. for some $L > 0$.*

Condition 3 *There is some $K > 0$ and some family of functions $\mathcal{F} \subseteq \mathcal{B}$ with $\text{VC}(\mathcal{F}) < \infty$ such that D is (K, \mathcal{F}) -regular.*

To state our main result, we must also introduce the *discrepancy function*. Let $E_{D, \mu} = \{\theta : D(\theta, \mu) = \max_{x \in \mathbb{R}^d} D(x, \mu)\}$ denote the set of maximizers of $D(\cdot, \mu)$ for fixed D and μ . For a set A and $r > 0$, let $B_r(A) = \{x : \min_{y \in A} \|x - y\| \leq r\}$. The discrepancy function of the pair (D, μ) is

$$\alpha(t) = D(\theta_0, \mu) - \sup_{x \in B_t^c(E_{D, \mu})} D(x, \mu), \quad (2.3)$$

where $\theta_0 \in E_{D, \mu}$. Given $t > 0$, the discrepancy function is the minimum distance between the maximal population depth value, and the maximal population depth value that is at least t -far from the maximizing set $E_{D, \mu}$. For example, often there is one maximizer, the

median, and the discrepancy function measures the minimum difference between the depth of the median, and the depth of any point at least t -far from the median. The faster the discrepancy function grows in t , the faster the depth function decreases as we move away from the median. Note that α is a non-decreasing function because the optimization problem involved is over decreasing sets. In Table 2.3 below, we explicitly calculate α for several depth functions.

We are now ready to state our main result, which provides bounds on the finite-sample deviations of the proposed differentially private multivariate medians via depths. For concreteness, we present our results for two of the most popular choices of priors, namely Gaussian measures and the uniform measure on $A_R(y)$, the d -dimensional cube of side-length R centered at y . In the following, let $d_{R,y}(x)$ denote the minimum distance from a point $x = (x_1, \dots, x_d)$ to a face of the cube $A_R(y)$: $d_{R,y}(x) = \min_{1 \leq i \leq d} |x_i - y_i \pm R/2|$. Similarly, denote the minimum distance from a set $B \subset \mathbb{R}^d$ to a face of the cube $A_R(y)$ by $d_{R,y}(B) = \inf_{x \in B} d_{R,y}(x)$ and let $d(x, B) = \inf_{y \in B} \|x - y\|$ denote the usual point-to-set distance, where $\|\cdot\|$ denotes the Euclidean norm. Define $\alpha^{-1}(t) = \sup\{r \geq 0: \alpha(r) = t\}$, where one notes that α^{-1} exists because α is a monotone function. Lastly, we write $a \lesssim b$ if $a \leq Cb$ for some universal constant $C > 0$.

Theorem 5 *If Conditions 1–3 hold, then the following holds:*

- (i) *Suppose that $L \geq 1$. If $\pi = \mathcal{N}(\theta_\pi, \sigma_\pi^2 I)$, then there exists a universal constant $c > 0$ such that for all $n \geq 8K/\epsilon$, all $d > 2$ and all $0 < \gamma < 1$, with probability at least $1 - \gamma$, we have that*

$$d(E_{D,\mu}, \tilde{\theta}_n) \lesssim \alpha^{-1} \left(cK \left[\sqrt{\frac{\log(1/\gamma) \vee \text{VC}(\mathcal{F}) \log n}{n}} \vee \frac{\log(1/\gamma) \vee \left(\frac{d(E_{D,\mu}, \theta_\pi)^2}{\sigma_\pi^2} + d \log \left(\frac{\sigma_\pi L n \epsilon}{K} \vee d \right) \right)}{n \epsilon} \right] \right). \quad (2.4)$$

- (ii) *If instead $\pi \propto \mathbb{1}\{x \in A_R(\theta_\pi)\}$ where $E_{D,\mu} \subset A_R(\theta_\pi)$, then there exists a universal constant $c > 0$ such that for all $n, d \geq 1$ and all $0 < \gamma < 1$, with probability at least $1 - \gamma$, we have that*

$$d(E_{D,\mu}, \tilde{\theta}_n) \lesssim \alpha^{-1} \left(cK \left[\sqrt{\frac{\log(1/\gamma) \vee \text{VC}(\mathcal{F}) \log n}{n}} \vee \frac{\log(1/\gamma) \vee d \log \left(\frac{Rn}{nd_{R,\theta_\pi}(E_{D,\mu}) \wedge K/L} \right)}{n \epsilon} \right] \right). \quad (2.5)$$

The proof of Theorem 5 can be found in Section 4, and relies on an elementary concentration bound (Theorem 12), which we introduce in Section 3. We also provide a version of Theorem 5 in terms of sample complexity, see Corollary 10. Theorem 5 is a deviations bounds for private multivariate medians. Put simply, it says that $\alpha(d(E_{D,\mu}, \tilde{\theta}_n))$ is bounded

above, with high probability, by two major terms: the sampling error term, which appears first in the upper bound, and the cost of privacy term, which appears second in the upper bound.

To further help interpret Theorem 5, note that we are mainly concerned about small deviations. For many choices of μ and D , one can show that α^{-1} is at most linear for small $t > 0$ under various regularity conditions that capture many popular depths, see Section 5 where we present such conditions. In this case one can remove the α^{-1} from the right-hand sides of (2.4)-(2.5). In addition, this fact coupled with the fact that many of the depth functions satisfy Condition 3 with $K = O(1)$ and $VC(\mathcal{F}) = O(d)$, and, in many contexts, the population median is unique, recovers a bound similar to those obtained in the case of private mean estimation (Kamath et al., 2019, 2020; Cai et al., 2021). That is, the bound given in Theorem 5 then reduces to the simpler bound

$$\left\| \text{MED}(\mu; D) - \tilde{\theta}_n \right\| \lesssim \sqrt{\frac{\log(1/\gamma) \vee d \log n}{n}} \bigvee \frac{\log(1/\gamma) \vee d \log \left(\frac{Rn}{nd_{R,\theta_\pi}(E_{D,\mu}) \wedge 1/L} \right)}{n\epsilon}.$$

Under such μ and D , assuming a reasonable prior and ignoring logarithmic terms, the cost of privacy term is greater than the sampling error term only when $\sqrt{d/n\epsilon} > 1$. For concrete examples of Theorem 5 applied to specific depth functions and population measures ν , see Examples 1 and 2.

While we state our result for general depth functions satisfying the above assumptions, our results apply to many of the most popular depths, such as halfspace depth and, more broadly, Type D depths in the sense of Zuo and Serfling (2000), spatial depth, simplicial depth and both the integrated rank-weighted and integrated dual depths. Depth functions to which these results apply and the conditions under which they apply are discussed at length in Section 6. The conditions for specific popular depths are summarized in Table 2.4.

We note that our bound is essentially sharp, see Lemma 18. Given that if μ is Gaussian, then the population mean equals the population median, the bound on the sample complexity implied by Theorem 5 (see Corollary 10) matches the lower bound on the sample complexity for private Gaussian mean estimation given by Kamath et al. (2019), Theorem 6.5 and Hopkins et al. (2022), Theorem 7.2 up to logarithmic factors.

Alternatively, Cai et al. (2021) give a minimax lower bound for approximately differentially private sub-Gaussian mean estimation. Unsurprisingly, given that it is designed for the approximately differentially private setting, this lower bound becomes $-\infty$ if we apply it to the setting of pure differential privacy. However, it is still interesting to compare Theorem 5 to the bound of Cai et al. (2021) in some way. A trivial manipulation of the conclusion of Lemma 18 shows that the bound given in Theorem 5 matches the lower bound of Cai et al. (2021), up to logarithmic terms, when the “approximate differential privacy parameter”, i.e., δ is taken to be at least n^{-k} for some $k > 0$. Therefore, as stated above, the bound in Theorem 5 is essentially sharp. Furthermore, at least for the setting of Gaussian mean estimation, relaxing to the setting of approximate differential privacy will not yield faster rates than those given by Theorem 5.¹

Recall that estimating a parameter from an unbounded parameter space, i.e., unbounded estimation, is non-trivial in the setting of differential privacy, e.g., see Avella-Medina and

1. It should not be difficult to extend this conclusion to the more general case of symmetric sub-Gaussian median estimation.

Brunel (2019). Theorem 5 shows that choosing σ_π relatively large allows one to perform unbounded estimation at the cost of an extra $\log d$ in the second term of (2.4), or, equivalently, an extra $\log d$ to the sample complexity. To elaborate, using a depth function in the exponential mechanism with a Gaussian prior constitutes unbounded estimation of the population median. Theorem 5 then shows that $d(E_{D,\mu}, \theta_\pi)/\sigma_\pi = O(\sqrt{d \log d})$ is sufficient for the deviations under a shifted location Gaussian prior to match the deviations under a correctly specified uniform prior, up to a logarithmic factor in d . Thus, we only require $d(E_{D,\mu}, \theta_\pi)$ to be polynomial in d , i.e., choosing σ_π to be some large polynomial in d will ensure that $\sigma_\pi \geq d(E_{D,\mu}, \theta_\pi)/\sqrt{d \log d}$.

Remark 6 (Projection depth) *One popular notion of depth not covered by our analysis is the projection depth (Zuo, 2003), see also (Stahel, 1981; Donoho, 1982). This is because a private median drawn from the exponential mechanism based on the empirical version of this depth function, is inconsistent μ -a.s, see Lemma 38. For a private version of the projection-depth-based median, see (Ramsay and Spicker, 2025).*

Remark 7 (Comparison to non-private bounds) *The first term in the maximum in Theorem 5 is a deviations bound for non-private depth-based medians. If we assume that $\theta_0 \in A_R(0)$ and choose the depth function to be halfspace depth (Definition 11), then this upper bound recovers the upper bound on the non-private estimation error of the halfspace median for the Gaussian measure, see Theorem 2.1 of Chen et al. (2018), see also Theorem 3 of Zhu et al. (2020).*

Remark 8 (Relaxing the Lipschitz requirement) *The Lipschitz requirement for D can be relaxed to obtain more general deviation bounds, at the cost of complicating the presentation. Furthermore, all the depth functions we consider are Lipschitz functions a.e. under mild assumptions, see Theorem 30. It is also not necessary to require that $E_{D,\mu} \subset A_R(\theta_\pi)$. It is sufficient to assume that $E_{D,\mu} \cap A_R(\theta_\pi) \neq \emptyset$. In this case, one can replace $d_{R,\theta_\pi}(E_{D,\mu})$ with $d_{R,\theta_\pi}(E_{D,\mu} \cap A_R(\theta_\pi))$.*

Remark 9 (The univariate setting) *Theorem 12 can be applied to obtain an optimal univariate private median estimator. Tzamos et al. (2020) gave a lower bound on the min-max sample complexity for estimating the one-dimensional median of a population measure whose density is bounded below by a positive constant. Applying Theorem 5 when $d = 1$ and D being the halfspace depth to such measures yields an upper bound with matches that lower bound. This upper bound recovers the one given by Karwa and Vadhan (2017) for one-dimensional Gaussian mean estimation.*

Note that Theorem 5 has a straightforward restatement in terms of sample complexity.

Corollary 10 *Suppose that Conditions 1–3 hold. Then the following holds*

- (i) *Suppose that $D(x, \mu) \leq 1$ and $L \geq 1$. If $\pi = \mathcal{N}(\theta_\pi, \sigma_\pi^2 I)$, then there exists universal constants $C, c > 0$ such that for all $t \geq 0$, all $d > 2$ and all $0 < \gamma < 1$, we have that*

D/μ :	$N_d(\theta_0, \Sigma)$	d -version symmetric
HD	$\Phi\left(\frac{t}{\sqrt{\lambda_1}}\right) - \frac{1}{2}$	$\frac{1}{2} - \sup_{\ v\ =1} \inf_{\ u\ =1} F_0\left(\frac{-tv^\top u}{a(u)}\right)$
IRW	$\int \left \frac{1}{2} - \Phi\left(\frac{tv_1^\top u}{\sqrt{u^\top \Sigma u}}\right) \right d\nu$	$\inf_{\ v\ =1} \int \left \frac{1}{2} - F_0\left(\frac{tv^\top u}{a(u)}\right) \right d\nu$
IDD	$\frac{1}{4} - \int \Phi\left(\frac{-tv_1^\top u}{\sqrt{u^\top \Sigma u}}\right) \Phi\left(\frac{tv_1^\top u}{\sqrt{u^\top \Sigma u}}\right) d\nu$	$\frac{1}{4} - \sup_{\ v\ =1} \int F_0\left(\frac{tv^\top u}{a(u)}\right) F_0\left(\frac{-tv^\top u}{a(u)}\right) d\nu$

Table 1: Table of the discrepancy function $\alpha(t)$ for different depth functions and underlying population measures. Recall that μ is d -version symmetric about zero if for any unit vector u , $X^\top u \stackrel{d}{=} a(u)Z$ where $X \sim \mu$ and Z is a random variable such that $Z \stackrel{d}{=} -Z$ and $a(u) = a(-u)$. Here, v_1 is the eigenvector associated with the largest eigenvalue λ_1 of Σ and $F_0(x) = F(x, u, \mu)$, see (2.9), and ν is the uniform measure on the $(d-1)$ -dimensional unit sphere.

$d(E_{D,\mu}, \tilde{\theta}_n) < t$ with probability at least $1 - \gamma$ provided

$$n > CK^2 \left(\frac{\log(1/\gamma) \vee [\text{VC}(\mathcal{F}) \log(\frac{K}{c\alpha(t)}) \vee 1]}{\alpha(t)^2} \vee \frac{\log(1/\gamma) \vee \left(\frac{d(E_{D,\mu}, \theta_\pi)^2}{\sigma_\pi^2} \right) \vee d \log\left(\frac{\sigma_\pi L}{\alpha(t)} \vee d\right)}{\epsilon \alpha(t)} \right). \quad (2.6)$$

(ii) If instead $\pi \propto \mathbb{1}\{x \in A_R(\theta_\pi)\}$ where $E_{D,\mu} \subset A_R(\theta_\pi)$, then there exists universal constants $C, c > 0$ such that for all $t \geq 0$, all $d \geq 1$ and all $0 < \gamma < 1$, we have that $d(E_{D,\mu}, \tilde{\theta}_n) \leq t$ with probability at least $1 - \gamma$ provided

$$n > CK^2 \left(\frac{\log(1/\gamma) \vee [\text{VC}(\mathcal{F}) \log(\frac{K}{c\alpha(t)}) \vee 1]}{\alpha(t)^2} \vee \frac{\log(1/\gamma) \vee d \log\left(\frac{R}{d_{R,\theta_\pi}(E_{D,\mu}) \wedge \alpha(t)/L}\right)}{\alpha(t)\epsilon} \right). \quad (2.7)$$

The proof is deferred to Appendix B.

2.4 Examples

Let us now turn to some concrete examples that illustrate our result. We first introduce the halfspace depth, which will be used in our examples, where we apply our main results to the halfspace median. For $X \sim \mu$, define

$$F(x, u, \mu) = \mu\left(X^\top u \leq x^\top u\right). \quad (2.8)$$

Halfspace depth, otherwise known as Tukey depth (Tukey, 1974), is defined as follows:

Definition 11 (Halfspace depth) *The halfspace depth HD of a point $x \in \mathbb{R}^d$ with respect to $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is*

$$\text{HD}(x, \mu) = \inf_{\|u\|=1} F(x, u, \mu). \quad (2.9)$$

The halfspace depth of a point $x \in \mathbb{R}^d$ is the minimum probability mass contained in a closed halfspace containing x . The halfspace median is the maximizer of the halfspace depth, when it is unique. We first consider the canonical problem of multivariate Gaussian location estimation under a Gaussian prior with shifted location, where here, the center of the prior θ_π is far from the population halfspace median. We quantify the relationship between finite-sample deviations, prior location misspecification, privacy budget, dimension and statistical accuracy. In this case, omitting logarithmic terms, for a location misspecification level of $\|\theta_0 - \theta_\pi\|/\sigma_\pi \leq \sqrt{d}$, the finite-sample deviations of the private halfspace median reduce to $O(\sqrt{\lambda_1 d/n} \vee \sqrt{\lambda_1 d/n\epsilon})$. Here, λ_1 is the largest eigenvalue of the covariance matrix of μ .

Next, we consider estimating the location parameter of a population measure μ made up of Cauchy marginals, which is known as the poly-Cauchy distribution. This setup is more difficult in the sense that the mean of the probability measure μ does not exist. For instance, theory concerning differentially private mean estimators cannot be applied in this setting.

Example 1 (Multivariate Gaussians with Gaussian prior with shifted location)

Consider the problem of estimating the location parameter θ_0 of a multivariate Gaussian measure $\mu = \mathcal{N}(\theta_0, \Sigma)$ and we wish to estimate θ_0 . For this example, let $\lambda_1 > \dots > \lambda_d > 0$ be the ordered eigenvalues of the covariance matrix Σ . Suppose that in order to estimate θ_0 we use the exponential mechanism paired with the halfspace depth $D = \text{HD}$ (see Definition 11) in conjunction with a misspecified Gaussian prior: $\pi = \mathcal{N}(\theta_\pi, \sigma_\pi^2 I)$.

In order to apply Theorem 5, we need to check Condition 1, 3 and Lipschitz continuity of HD. First, note that Condition 1 holds for any bounded depth function as defined in Definition 2. Next, one can show that (see Theorem 30) the halfspace depth is $(1, \mathcal{F})$ -regular where \mathcal{F} is the set of closed halfspaces in \mathbb{R}^d , thus, $\text{VC}(\mathcal{F}) = d + 2$. Lastly, the map $x \mapsto \text{HD}(x, \mu)$ is L -Lipschitz with $L = (2\pi\lambda_d)^{-1/2}$, implying Condition 2 is satisfied. We may now apply Theorem 5, for which we must compute $\alpha(t)$. To this end, it holds that $\alpha(t) \gtrsim tC_R/\sqrt{\lambda_1}$, for some large $R > t$, see Lemmas 21 and 40. Therefore, under halfspace depth, with probability at least $1 - \gamma$, it holds that

$$\|\theta_0 - \tilde{\theta}_n\| \lesssim \frac{\sqrt{\lambda_1}}{C_R} \left[\sqrt{\frac{\log(1/\gamma) \vee d \log n}{n}} \vee \frac{\log(1/\gamma) \vee \frac{\|\theta_0 - \theta_\pi\|^2}{\sigma_\pi^2} \vee d \log(n\epsilon\sigma_\pi/\lambda_d \vee d)}{n\epsilon} \right].$$

Suppose we would like to achieve a success rate of $1 - e^{-d}$, i.e., $\gamma = e^{-d}$. This reduces the deviations bound to

$$\|\theta_0 - \tilde{\theta}_n\| \lesssim \frac{\sqrt{\lambda_1}}{C_R} \left[\sqrt{\frac{d \log n}{n}} \vee \frac{\frac{\|\theta_0 - \theta_\pi\|^2}{\sigma_\pi^2} \vee d \log(n\epsilon\sigma_\pi/\lambda_d \vee d)}{n\epsilon} \right].$$

Observe now that in order for the prior effect to be negligible, we require $\|\theta_0 - \theta_\pi\| / \sigma_\pi \leq \sqrt{d \log d}$ and that σ_π is at most polynomial in d . This suggests that, up to a point, a larger prior variance σ_π^2 produces lower deviations. In this case, when the error level, ϵ and Σ are fixed, we only need n to grow slightly faster than the dimension $n \geq Cd \log d$ in order to maintain a consistent level of error. If the prior effect is negligible in the sense just described, then, omitting logarithmic terms for brevity, the deviations bound reduces to

$$\|\theta_0 - \tilde{\theta}_n\| \lesssim \frac{\sqrt{\lambda_1}}{C_R} \cdot \left(\sqrt{\frac{d}{n}} \vee \frac{d}{n\epsilon} \right), \quad (2.10)$$

We then have that the cost of privacy is eclipsed by the sampling error whenever $n \gtrsim d/\epsilon^2$. Therefore, the cost of differential privacy is asymptotically “free”, provided that d is not too large. That is, the error incurred by differential privacy is much smaller than that of the sampling error. Furthermore, (2.10), makes it is easy to see that in this case, the private medians achieve the minimax lower bound for approximately differentially private sub-Gaussian mean estimation (Cai et al., 2021), if $\delta \propto n^{-k}$. There exists other private estimators that are minimax optimal for private Gaussian mean estimation, e.g., (Kamath et al., 2019; Biswas et al., 2020). These estimators may be computable in polynomial time, whereas the private halfspace median would take exponential time. However, existing estimators may not be robust. For example, they may not perform well if the population measure is not multivariate Gaussian but instead is multivariate Cauchy. Robustness against violations of assumptions is especially important in private data analysis, where we may not have enough privacy budget to check assumptions. In addition, the parameter of interest may be the population median. In that case, mean estimators would not be suitable if the population measure is not symmetric. Note that a bound in terms of Mahalanobis distance (with respect to Σ), denoted $\|\cdot\|_\Sigma$, is implied by (2.10), through the inequality $\|\theta_0 - \tilde{\theta}_n\|_\Sigma \leq \|\theta_0 - \tilde{\theta}_n\| / \sqrt{\lambda_d}$, viz.

$$\|\theta_0 - \tilde{\theta}_n\|_\Sigma \lesssim \frac{\sqrt{\lambda_1 \lambda_d}}{C_R} \cdot \left(\sqrt{\frac{d}{n}} \vee \frac{d}{n\epsilon} \right).$$

However, this resulting bound is not optimal for Gaussian mean estimation, see, e.g., (Brown et al., 2021). This does not imply that this estimator is not optimal for this problem, it is just that the optimal Mahalanobis error is not implied by our bound. It is possible the techniques used to prove Theorem 12 could be used to prove a general deviations bound in terms of the Mahalanobis distance.

Example 2 (Cauchy marginals) Suppose that μ is the product probability measure constructed from independent Cauchy marginals with location parameter 0 and scale parameters $\sigma_1, \dots, \sigma_d$ and we would like to use the exponential mechanism to privately estimate the half-space median of μ . In this case, μ is d -version symmetric², with $a(u) = \sum_{j=1}^d \sigma_j |u_j|$ where u_j is the j^{th} coordinate of u . A similar analysis as to that of Example 1, gives that halfspace depth (under this μ) satisfies the conditions of Theorem 5. Note that in this setting,

2. Recall that μ is d -version symmetric about zero if for any unit vector u , $X^\top u \stackrel{d}{=} a(u)Z$ where $X \sim \mu$ and $Z \sim \nu$ such that $Z \stackrel{d}{=} -Z$ and $a(u) = a(-u)$.

	HD	IDD	IRW	SMD	SD	MSD
$\text{VC}(\mathcal{F})$	$O(d)$	$O(d)$	$O(d)$	$O(d^2 \log d)$	$O(d)$	$O(d)$
K	1	3	4	$d + 1$	d	1
L	$\sup_u \ f_u\ $	$3 \sup_u \ f_u\ $	$2 \sup_u \ f_u\ $	$\sup_u \ f_u\ $	$2L'$	$2\sqrt{d}L'$
\mathcal{A}	$\mathcal{M}_1(\mathbb{R}^d)$	\mathcal{C}	\mathcal{C}	\mathcal{N}	\emptyset (see Remark 28)	\emptyset

Table 2: Table of values for which different depth functions satisfy Theorem 5. Letting μ_u be the law of $X^\top u$ if $X \sim \mu$, f_u is the density of μ_u with respect to the Lebesgue measure and L' denotes $\sup_y \mathbb{E} \|y - X\|^{-1}$. Note \mathcal{A} is the set of admissible measures for which D is a depth function, in the sense of Definition 2. Here, $\mathcal{C} \subset \mathcal{M}_1(\mathbb{R}^d)$ denotes the set of centrally symmetric measures over \mathbb{R}^d and $\mathcal{N} \subset \mathcal{M}_1(\mathbb{R}^d)$ denotes the set of angularly symmetric measures over \mathbb{R}^d that are absolutely continuous with respect to the Lebesgue measure.

the Lipschitz constant scales with the dimension: $L = O(\sqrt{d})$. The next step is to compute $\alpha(t)$. Let $\bar{\sigma} = d^{-1} \sum_{j=1}^d \sigma_j$. Observe that when the scale parameters are constant in d , $\sup_u a(u) = \sqrt{d} \bar{\sigma}$. It follows that under the halfspace depth, $\alpha(t) = \frac{1}{\pi} \arctan\left(\frac{t}{\sqrt{d}\bar{\sigma}}\right)$. If $t < R$ for some large R , then we can write $\alpha(t) \gtrsim t(\bar{\sigma}\sqrt{d}(1 + R^2/d\bar{\sigma}^2))^{-1}$. If we use a Gaussian prior with unit scale,³ then, omitting logarithmic terms, the deviations bound becomes

$$\|\tilde{\theta}_n\| \lesssim \frac{d\bar{\sigma}}{\sqrt{n}} \vee \frac{d^{1/2} \|\theta_\pi\|^2 \vee d^{3/2}}{\epsilon n}.$$

Here, observe that the deviations are a factor of \sqrt{d} larger than in the Gaussian setting. The increase in the sampling error term is due to the fact that if X follows the poly-Cauchy distribution, then there is a direction u where $X^\top u$ has scale proportional to \sqrt{d} . This causes the discrepancy function to be dependent on the dimension. For the poly-Cauchy measure, considering all directions is actually a drawback of the halfspace median. On the other hand, the coordinate-wise median does not consider the direction with scale proportional to \sqrt{d} , and does not have this issue. Indeed, when we apply Theorem 5 to the univariate private median, the resulting deviations bound is $O(\sqrt{d \max_{i \in [d]} \sigma_i^2/n} \vee [d^{3/2}(\|\theta_\pi\|_\infty^2 \vee 1)/\epsilon n])$. Observe that the sampling error term is $O(\sqrt{d/n})$. (Note that any other optimal, univariate private median estimator would achieve the same deviations bound, e.g., (Tzamos et al., 2020).) On the other hand, the cost of privacy term in the deviations bound is $O(d^{3/2}/n\epsilon)$ for both the halfspace median and coordinate-wise median. This represents the difficulty of estimating location when μ is high-dimensional and heavy-tailed.

In the Gaussian setting, the discrepancy function is relatively similar between the halfspace depth to the integrated depths. However, in the Cauchy setting, the discrepancy functions differ between the halfspace depth, the integrated-rank weighted depth and the integrated

3. We don't lose generality here; we can arbitrarily shrink the scale of the μ .




Figure 1: Logarithm of the inverse of the value of the discrepancy function at $t = 1$, $\log(1/\alpha(1))$, as the dimension increases. Here, μ is made up of d independent standard Cauchy marginals. Notice that halfspace depth has the smallest values of $1/\alpha(1)$ which implies it has a smaller bound on the sample complexity. This demonstrates the trade-off between accuracy and computability when choosing a depth function. For the integrated depths, the value of the integrand is approximated via Monte Carlo simulation with 10,000 unit vectors.

dual depth. To see this, we must evaluate the integrals in Table 2.3. For the IRW this integral is given by

$$\alpha(t) = \frac{1}{\pi} \inf_{\|v\|=1} \int \left| \arctan \left(\frac{tv^\top u}{\sum_{j=1}^d \sigma_j |u_j|} \right) \right| d\nu(u).$$

We evaluate these integrals via Monte Carlo simulation with 10,000 unit vectors. Figure 1 plots the logarithm of the inverted discrepancy function computed at $t = 1$, i.e., $\log(1/\alpha(1))$, for the three depth functions as the dimension increases. Observe that the halfspace depth's discrepancy function decreases the slowest. Thus, the halfspace depth gives the smallest sample complexity for heavy tailed estimation when the dimension is large. The difference in accuracy between the halfspace depth and the integrated depths demonstrates the trade-off between computability and statistical accuracy.

2.5 Simulation

To complement our theoretical results, we demonstrate the cost of privacy through simulation. In order to simulate a private median, we use (non-private) discretized Langevin dynamics. Given that the discretized Langevin dynamics approximates a draw from the exponential mechanism, we don't expect the use of discretized Langevin dynamics to have a meaningful impact on our results. We recognize that privacy guarantees may be affected by the use of Markov chain Monte Carlo (MCMC) methods; the current implementation may not guarantee ϵ -differential privacy. This appears to be an open problem. The purpose of the current implementation is to investigate the statistical accuracy of the methods. The computation of specific private medians via private MCMC is an interesting direction of future research. For a discussion on the level of privacy guaranteed by MCMC, see (Seeman

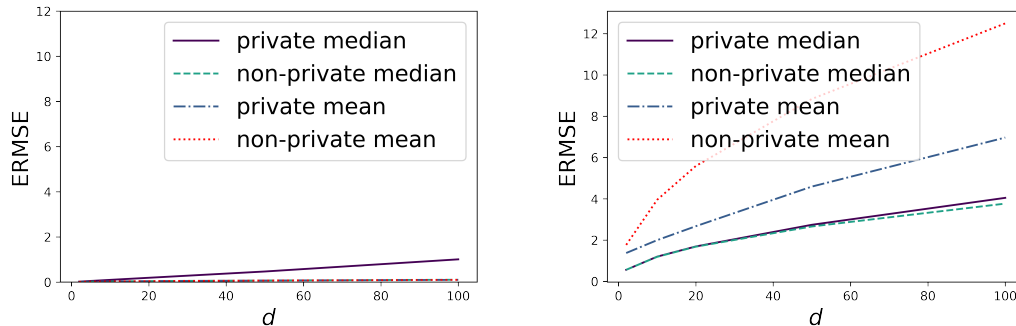


Figure 2: Empirical root mean squared error (ERMSE) of the location estimators under (left) Gaussian data and (right) contaminated Gaussian data when $\epsilon = 10$. In the uncontaminated setting, the ERMSE of the private median grows at a slightly faster rate than its non-private counterpart and the mean estimators. This is due to the cost of privacy. On the other hand, in the contaminated setting, the ERMSE of median estimators is much lower than that of the mean estimators. In particular, the private median outperforms the coin-press mean. This demonstrates the usual trade-off between accuracy and robustness between the mean and the median. For a zoomed in version of the left plot, see Figure 4 in the appendix.

et al., 2021). As a by-product of our work, we also provide a fast implementation of a modified version of the non-private integrated dual depth-based median. This implementation shows that this multivariate median can be computed quickly, even when the dimension is large. For example, we can compute the median of ten thousand 100-dimensional samples in less than one second on a personal computer.⁴ We call this modified version of the integrated dual depth the smoothed integrated dual depth, see Section 6.1 for more details. Our code is available on GitHub.

Aside from the private smoothed integrated dual depth median, we also computed the non-private smoothed integrated dual depth median and the private coin-press mean of Biswas et al. (2020). When computing the coin-press mean, we used the existing GitHub implementation provided by the authors. The bounding ball for the coin-press algorithm had radius $10\sqrt{d}$ and was run for four iterations.

We simulated fifty instances with $n = 10,000$ over a range of dimensions from two to one hundred. The data were either generated from a standard Gaussian measure or contaminated standard Gaussian measure. In the contaminated model, 25% of the observations had mean $(5, \dots, 5)$. We tested privacy parameters $\epsilon = 2, 5$ and 10, the smoothing parameter s was fixed at 10 and $\pi = \mathcal{N}(0, 25dI)$.

Figure 2 shows the empirical root mean squared error (ERMSE) for the different location estimators as the dimension increases when $\epsilon = 10$. Results for the remaining values of ϵ can be seen in Appendix D. Notice that the private median ERMSE grows at a similar

4. The code was run on a desktop computer with an Intel i7-8700K 3.70GHz chipset and an Nvidia GTX 1660 super graphics card.

rate as that of the non-private ERMSE. The mean estimators perform very well in the uncontaminated setting, but poorly in the contaminated setting. In particular, the private median outperforms the coin-press mean. The median estimators do not perform as well as the mean estimators in the standard Gaussian setting, but do not become corrupted in the contaminated setting. This demonstrates the usual trade-off between accuracy and robustness between the mean and the median.

3. An elementary concentration bound

In this section, we introduce the underlying concentration bound used to produce Theorem 5, which we imagine will be useful more broadly. Our result shows that for a general objective function ϕ , $\tilde{\theta}_n$ concentrates around $\theta_0 = \operatorname{argmax}_{\theta \in \mathbb{R}^d} \phi(\theta, \mu)$ with high probability. In particular, Theorem 5 is specifically for privately estimating multivariate medians, given normal or uniform priors, meanwhile the concentration bound presented in this section applies more broadly to $\tilde{\theta}_n$ drawn from the more general (2.2) as an estimate of $\theta_0 = \operatorname{argmax}_{\theta \in \mathbb{R}^d} \phi(\theta, \mu)$, under a general prior and a relaxed version of Condition 2. The relaxed version of Condition 2 is stated below.

Condition 4 *The map $x \mapsto \phi(x, \mu)$ is π -a.s. uniformly continuous with modulus of continuity ω .*

We now introduce some new important functions. Replacing the depth function D with ϕ , we can define the generalized discrepancy function of the pair (ϕ, μ) as $\alpha(t) = \phi(\theta_0, \mu) - \sup_{x \in B_t^c(E_{\phi, \mu})} \phi(x, \mu)$, where $\theta_0 \in E_{\phi, \mu}$. The *calibration function* of the triple (ϕ, μ, π) is the function

$$\psi(\lambda) = \min_{r > 0} [\lambda \cdot \omega(r) - \log \pi(B_r(E_{\phi, \mu}))] \quad (3.1)$$

Note that $\psi(\lambda)$ is increasing and tends to ∞ as $\lambda \rightarrow \infty$. The *rate function* of the prior is given by

$$I(t) = -\log \pi(B_t^c(E_{\phi, \mu})). \quad (3.2)$$

The concentration/entropy function of the prior is the rate at which the prior concentrates around the maximizers of the population risk. For large t , it is the rate of decay of the tails of π .

The concentration bound is given in the following theorem, which quantifies the accuracy of $\tilde{\theta}_n$ drawn from (2.2) as an estimate of θ_0 .

Theorem 12 *Suppose that for the pair (ϕ, μ) , Conditions 1, 3 and 4 hold for some $K > 0$. Then, there are universal constants $c_1, c_2 > 0$ such that the following holds. For any $t, \beta > 0$, we have that*

$$\Pr \left(\max_{\theta_0 \in E_{\phi, \mu}} \|\tilde{\theta}_n - \theta_0\| \geq t \right) \leq \left(\frac{c_1 n}{\operatorname{VC}(\mathcal{F})} \right)^{\operatorname{VC}(\mathcal{F})} e^{-c_2 n \left[\frac{1}{\beta} I(t) \vee \alpha(t) \right]^2 / K^2} + e^{-\beta \alpha(t) / 2 - I(t) / 2 + \psi(\beta)}.$$

Remark 13 (Comparison to some well-known bounds) *It is helpful to compare our result to the widely used accuracy bound of McSherry and Talwar (2007). The latter controls the accuracy of the estimator $\tilde{\theta}_n$ as measured by the empirical objective value: it bounds the difference between the empirical objective value achieved by the exponential mechanism and*

the maximal empirical objective value. More precisely, if we let $S_t = \{x \in \mathbb{R}^d: \phi(x, \hat{\mu}_n) - \sup_{\theta} \phi(\theta, \hat{\mu}_n) \geq 2t\}$, the result of McSherry and Talwar (2007) says that

$$\Pr \left(\sup_{\theta} \phi(\theta, \hat{\mu}_n) - \phi(\tilde{\theta}_n, \hat{\mu}_n) > 2t \right) \leq e^{-t\epsilon} / \pi(S_t).$$

As well, note that there is no analogous sample complexity bound because this approximation is not improving with the sample size. By contrast, (3.10) concerns the consistency of $\tilde{\theta}_n$ as an estimator, namely the distance between the estimator and any maximizer of the population objective, that is $\|\tilde{\theta}_n - \theta_0\|$. Furthermore, the quality of this concentration bound depends on n and thus a sample complexity bound can be read off readily. In addition to McSherry and Talwar (2007), Chaudhuri and Hsu (2012) also provide an accuracy bound for M -estimators drawn from the exponential mechanism in the univariate setting. They give the number of samples required for the error of the private estimator to be within $t > 0$ of the error of the non-private estimator. Our bound can be applied to the same problem, and we also recover the requirement that $n \gtrsim 1/t\epsilon$.

Proof We first prove that the regularity conditions imply that the empirical objective function concentrates around the population objective function. By Condition 3, there are some $K > 0$ and \mathcal{F} with $\text{VC}(\mathcal{F}) < \infty$ such that ϕ is (K, \mathcal{F}) -regular, i.e.,

$$\sup_x |\phi(x, \mu) - \phi(x, \hat{\mu}_n)| \leq K \sup_{g \in \mathcal{F}} |\mathbb{E}_{\hat{\mu}_n} g(X) - \mathbb{E}_{\mu} g(X)|. \quad (3.3)$$

Now, for a probability measure Q over \mathcal{B} , let $N(\tau, \mathcal{F}, L_2(Q))$ be the τ -covering number of \mathcal{F} with respect to the $L^2(Q)$ -norm. By assumption, for any $g \in \mathcal{F}$, $\|g\|_{\infty} \leq 1$ and $\text{VC}(\mathcal{F}) < \infty$. This fact implies that there is some universal $c > 0$ such that for any $0 < \tau < 1$, it holds that,

$$\sup_Q N(\tau, \mathcal{F}, L^2(Q)) \lesssim \text{VC}(\mathcal{F}) \left(\frac{c}{\tau} \right)^{2\text{VC}(\mathcal{F})}.$$

See (Sen, 2022, Theorem 7.12), or (Kosorok, 2008, Theorem 9.3). A straightforward manipulation of Talagrand's inequality (Talagrand, 1994, Theorem 1.1) implies that there is some universal $c > 0$ such that for all $n \geq 1$ and for all $t > 0$, it holds that

$$\Pr \left(\sqrt{n} \sup_{g \in \mathcal{F}} \left| \frac{1}{n} \sum_{i=1}^n g(X_i) - \mathbb{E} g(X) \right| \geq t \right) \leq \left(\frac{cn}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-2t^2}.$$

For all $t > 0$, it holds that

$$\Pr \left(\sup_{x \in \mathbb{R}^d} |\phi(x, \hat{\mu}_n) - \phi(x, \mu)| > t \right) \leq \left(\frac{cn}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-2nt^2/K^2}. \quad (3.4)$$

where to go from the second inequality to the third we multiplied both sides of the inequality by \sqrt{n} .

We can now use the concentration result (3.4) to prove the rest of the theorem. For brevity, let $D_{n,t} = \{\max_{\theta_0 \in E_{\phi, \mu}} \|\tilde{\theta}_n - \theta_0\| > t\}$, and for any $y > 0$, define the event $A_{n,y} =$

$\{\|\phi(\cdot, \hat{\mu}_n) - \phi(\cdot, \mu)\|_\infty < y\}$. It follows from the law of total probability and (3.4) that for any $y > 0$ it holds that

$$\begin{aligned} \Pr(D_{n,t}) &= \Pr(D_{n,t} \cap A_{n,y}^c) + \Pr(D_{n,t} \cap A_{n,y}) \leq \Pr(A_{n,y}^c) + \Pr(D_{n,t} \cap A_{n,y}) \\ &\leq \left(\frac{cn}{K \text{VC}(\mathcal{F})}\right)^{\text{VC}(\mathcal{F})} e^{-2ny^2/K^2} + \Pr(D_{n,t} \cap A_{n,y}). \end{aligned} \quad (3.5)$$

We now focus on the right-hand term above. By definition, we have that

$$\frac{1}{\beta} \log \Pr(D_{n,t} \cap A_{n,y}) = \frac{1}{\beta} \log \int_{A_{n,y}} \frac{\int_{B_t^c(E_{\phi,\mu})} \exp(\beta\phi(x, \hat{\mu}_n)) d\pi}{\int_{\mathbb{R}^d} \exp(\beta\phi(x, \hat{\mu}_n)) d\pi} d\mu. \quad (3.6)$$

On $A_{n,y}$ it holds that $\exp(\beta\phi(x, \mu)) \exp(-\beta y) \leq \exp(\beta\phi(x, \hat{\mu}_n)) \leq \exp(\beta\phi(x, \mu)) \exp(\beta y)$. Applying this to the right-hand side of (3.6) results in

$$\frac{1}{\beta} \log \Pr(D_{n,t} \cap A_{n,y}) \leq \frac{1}{\beta} \log \frac{\int_{B_t^c(E_{\phi,\mu})} \exp(\beta\phi(x, \mu)) d\pi}{\int_{\mathbb{R}^d} \exp(\beta\phi(x, \mu)) d\pi} + 2y. \quad (3.7)$$

We now bound $\mathbb{E}_\pi \exp(\beta\phi(x, \mu))$ below. For any $r > 0$, Condition 4 implies that

$$\begin{aligned} \frac{1}{\beta} \log \int_{\mathbb{R}^d} \exp(\beta\phi(x, \mu)) d\pi &\geq \frac{1}{\beta} \log \int_{B_r(E_{\phi,\mu})} \exp(\beta\phi(x, \mu)) d\pi \\ &= \phi(\theta_0, \mu) + \frac{1}{\beta} \log \int_{B_r(E_{\phi,\mu})} \exp(\beta(\phi(x, \mu) - \phi(\theta_0, \mu))) d\pi \\ &\geq \phi(\theta_0, \mu) - \omega(r) + \beta^{-1} \log \pi(B_r(E_{\phi,\mu})). \end{aligned}$$

Maximizing the right-hand side of the above, and recalling the definition of the calibration function, ψ , from (3.1), yields

$$\frac{1}{\beta} \log \int_{\mathbb{R}^d} \exp(\beta\phi(x, \mu)) d\pi \geq \phi(\theta_0, \mu) - \psi(\beta)/\beta.$$

Plugging this lower bound into (3.7) yields

$$\begin{aligned} \frac{1}{\beta} \log \Pr(D_{n,t} \cap A_{n,y}) &\leq \frac{1}{\beta} \log \int_{B_t^c(E_{\phi,\mu})} \exp(\beta\phi(x, \mu)) d\pi - \phi(\theta_0, \mu) + \psi(\beta)/\beta + 2y \\ &\leq \sup_{x \in B_t^c(E_{\phi,\mu})} \phi(x) - \phi(\theta_0) + \frac{1}{\beta} \log \pi(B_t^c(E_{\phi,\mu})) + \psi(\beta)/\beta + 2y \\ &= -\alpha(t) - I(t)/\beta + \psi(\beta)/\beta + 2y, \end{aligned}$$

where the last line follows from the definitions of $\alpha(t)$ and $I(t)$. Rewriting the last inequality gives that

$$\Pr(D_{n,t} \cap A_{n,y}) \leq \exp(-\beta\alpha(t) - I(t) + \psi(\beta) + 2\beta y). \quad (3.8)$$

Now, setting $y = f(t, \beta) = \alpha(t)/2 \vee I(t)/2\beta$ and plugging (3.8) into (3.5) results in

$$\Pr \left(\max_{\theta_0 \in E_{\phi, \mu}} \left\| \tilde{\theta}_n - \theta_0 \right\| > t \right) \leq \left(\frac{cn}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-nf(t, \beta)^2/8K^2} + e^{-\beta\alpha(t) - I(t) + \psi(\beta) + 2\beta f(t, \beta)}. \quad (3.9)$$

Next, note that $\alpha(t) + I(t)/\beta - f(t, \beta) \geq \alpha(t)/2 + I(t)/2\beta$. Plugging this inequality into (3.9) yields the desired result. \blacksquare

It is helpful to view the two terms in the above bounds separately. The first term measures the sample complexity of estimation and, in the absence of sampling considerations, recovers the concentration properties of classical (non-private) estimators. (In particular, the non-private bound can be obtained from the above by formally setting $\beta = \infty$.) The second term is novel and is a measure of the cost of sampling and encodes the trade-off between the prior, π , the parameter β , and the objective function α .

Through taking $\beta = n\epsilon/2K$, Theorem 12 immediately provides a consistency bound (and thereby a sample complexity bound) for the exponential mechanism:

$$\Pr \left(\max_{\theta_0 \in E_{\phi, \mu}} \left\| \tilde{\theta}_n - \theta_0 \right\| \geq t \right) \leq \left(\frac{c_1 n}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-c_2 n \left[\frac{1}{n\epsilon} I(t) \vee \alpha(t)/2K \right]^2} + e^{-n\epsilon\alpha(t)/4K - I(t)/2 + \psi(n\epsilon/2K)}. \quad (3.10)$$

Let us pause here to interpret this result. The cost of consistent estimation is encoded in the first terms of (3.10) and, in particular, can be related to the non-private setting as well. The cost of privacy is encoded in the second term in (3.10). Finally, we note that the second term in (3.10) encodes a trade-off between an effect of the prior and the privacy parameter.

Equation (3.10) readily yields the sample complexity of an estimator drawn from the exponential mechanism.⁵ This yields Corollary 10, and, for instance, Lemma 15 below.

4. Proof of Theorem 5

We now prove Theorem 5 with a series of lemmas. Let us first note the following fact, whose proof is deferred to Appendix B.

Lemma 14 *Let $a, b, c \in \mathbb{R}^+$. The function $h(x) = ax - b \log(cx/b)$ is increasing and positive for*

$$x \geq \frac{2b}{a} [\log(c/a) \vee 1].$$

5. To see this briefly, assume for simplicity that π is the uniform measure on a hypercube whose side lengths are at most polynomial in d . In addition, assume that θ_0 is well captured by this hypercube, i.e., θ_0 is at least e^{-d} far from any boundary of the cube. Suppose we want to estimate a maximizer of $\phi(\cdot, \mu)$ within some error level t with probability at least $1 - e^{-d}$. Lemma 35 and Corollary 10 give that it is sufficient for $\text{VC}(\mathcal{F})$ to be polynomial in d to ensure that the sample complexity is polynomial in d .

Lemma 15 *Suppose that Conditions 1 and 3 hold with $\phi = D$ for some $K > 0$. In addition, suppose that the map $x \mapsto D(x, \mu)$ is L -Lipschitz for some $L > 0$. Then, there exists a universal constant $c > 0$, such that for all $n, d \geq 1$ and all $0 < \gamma < 1$, with probability at least $1 - \gamma$, it holds that*

$$d(E_{D, \mu}, \tilde{\theta}_n) \lesssim \alpha^{-1} \left(cK \left[\sqrt{\frac{\log(1/\gamma) \vee \text{VC}(\mathcal{F}) \log n}{n}} \vee \frac{\log(1/\gamma) \vee \left(1 - \log \pi(B_{\frac{2K}{Ln\epsilon}}(E_{\phi, \mu}))\right)}{n\epsilon} \right] \right). \quad (4.1)$$

Proof Note that all of the conditions of Theorem 12 hold, and so we can apply Theorem 12 with $\beta = n\epsilon/2K$. Proving the deviation bound then amounts to finding a lower bound on t which ensures that:

$$\left(\frac{c_1 n}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-c_2 n [2KI(t)/n\epsilon \vee \alpha(t)]^2 / K^2} + e^{-n\epsilon \alpha(t)/2K - I(t)/2 + \psi(n\epsilon/2K)} := I + II \leq \gamma. \quad (4.2)$$

We will do this by finding t such that both I and II are less than $\gamma/2$. We start with finding t such that $I \leq \gamma/2$. First, note that

$$\left(\frac{c_1 n}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-c_2 n [2KI(t)/n\epsilon \vee \alpha(t)]^2 / K^2} \leq \left(\frac{c_1 n}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} e^{-c_2 n \alpha(t)^2 / K^2}.$$

If we can find a lower bound on t which implies that

$$\left(\frac{c_1 n}{\text{VC}(\mathcal{F})} \right)^{\text{VC}(\mathcal{F})} \leq e^{c_2 n \alpha(t)^2 / 2K^2}, \quad (4.3)$$

then such a lower bound, together with the bound $t \gtrsim \alpha^{-1} \left(\sqrt{K^2 \log(2/\gamma)/n} \right)$ yields that $I \leq \gamma/2$. It remains to find t such that (4.3) holds.

To this end, define the function $h(n) = c_2 n \alpha(t)^2 / 2K^2 - \text{VC}(\mathcal{F}) \log(c_1 n / \text{VC}(\mathcal{F}))$. The condition (4.3) is then equivalent to the condition $h(n) \geq 0$. Note that $h(n)$ is in the form of Lemma 14 with $a = c_2 \alpha(t)^2 / 2K^2$, $b = \text{VC}(\mathcal{F})$ and $c = c_1$. Applying Lemma 14 yields that $h(n) \geq 0$ provided that

$$n \gtrsim K^2 \frac{\text{VC}(\mathcal{F}) \log \left(\frac{K^2}{c \alpha(t)^2} \right) \vee 1}{\alpha(t)^2}, \quad (4.4)$$

for some universal $c > 0$. Now, provided that $\alpha^2(t) \gtrsim K^2/n$, (4.4) is satisfied when

$$\alpha(t)^2 \gtrsim K^2 \text{VC}(\mathcal{F}) \log n / n.$$

It follows that there exists a universal constant $c > 0$ such that (4.3) holds provided that

$$t \gtrsim \alpha^{-1} \left(cK \sqrt{\frac{\log(1/\gamma) \vee \text{VC}(\mathcal{F}) \log n}{n}} \right). \quad (4.5)$$

We now turn to finding t such that $II \leq \gamma/2$. The $II \leq \gamma/2$ is satisfied when

$$\alpha(t) > 4K \frac{\log(2/\gamma) \vee \psi(n\epsilon/2K)}{n\epsilon} > 2K \frac{\log(2/\gamma) - I(t)/2 + \psi(n\epsilon/2K)}{n\epsilon}. \quad (4.6)$$

It remains to simplify $\psi(n\epsilon/2K)$. To this end, given that by assumption D is L -Lipschitz, we have that $\omega(r) \leq Lr$. Plugging this in to the definition of ψ yields:

$$\psi(n\epsilon/2K) \leq \min_{r>0} \left[\frac{n\epsilon}{2K} \cdot L \cdot r - \log \pi(B_r(E_{\phi,\mu})) \right].$$

Next, taking the specific choice $r = 2K/Ln\epsilon$ yields $\psi(n\epsilon/2K) \leq 1 - \log \pi(B_{2K/Ln\epsilon}(E_{\phi,\mu}))$. Plugging this inequality into (4.6), and combining the result with (4.5) yields the desired result. \blacksquare

We now note the following lemmas whose proofs are deferred to Appendix B.

Lemma 16 *If $\pi = \mathcal{N}(\theta_\pi, \sigma_\pi^2 I)$ with $\sigma_\pi \geq 1/4$, then for all $E \subset \mathbb{R}^d$, all $d > 2$ and all $R \leq \sigma_\pi$ it holds that*

$$-\log \pi(B_R(E)) \lesssim \frac{d(E, \theta_\pi)^2}{\sigma_\pi^2} + d \log \left(\frac{\sigma_\pi}{R} \vee d \right). \quad (4.7)$$

Lemma 17 *Suppose that $\pi \propto \mathbb{1}\{x \in A_R(\theta_\pi)\}$ where $E \subset A_R(\theta_\pi)$, then for all $E \subset \mathbb{R}^d$, $d > 0$ and all $r > 0$ it holds that*

$$-\log \pi(B_r(E)) \lesssim d \log \left(\frac{R}{d_{R,\theta_\pi}(E) \wedge r} \right).$$

We can now prove Theorem 5.

Proof [Proof of Theorem 5] First note that by assumption, all the conditions of Lemma 15 are satisfied. To prove (2.4) and (2.5) we apply Lemma 16 and Lemma 17, respectively, in conjunction with Lemma 15. To apply Lemma 16, first, observe that together, the assumptions that $L > 1$, $\sigma_\pi^2 \geq 1/16$ and $n \geq 8K/\epsilon$ imply that $2K/Ln\epsilon \leq \sigma_\pi$. Applying Lemma 16 with $R = 2K/Ln\epsilon$ and $E = E_{D,\mu}$ in conjunction with (4.1), yields (2.4). Applying Lemma 17 with $r = 2K/Ln\epsilon$ and $E = E_{D,\mu}$ in conjunction with (4.1), yields (2.5) as desired. \blacksquare

4.1 Optimality of Theorem 5

Lemma 18 *The upper bound on the deviation of $\tilde{\theta}_n$ about $E_{D,\mu}$ given in Theorem 5 is sharp, up to logarithmic factors. The upper bound on the sample complexity of $\tilde{\theta}_n$ given in Corollary 10 is also sharp, up to logarithmic factors.*

Proof We show that for a specific set of parameters, the upper bound given in Corollary 10 matches the lower bound of Theorem 6.5 of (Kamath et al., 2019) up to logarithmic factors in d and t^{-1} . Specifically, set $D = \text{HD}$, $\theta_0 \in A_R(0)$, $\pi = A_R(0)$, that $\mu = \mathcal{N}(\theta_0, \sigma^2 I)$. For $t \geq R$, the sample complexity is 1. We then focus on the case where $t \leq R$.

We now check the conditions of Corollary 10 under these parameters. To this end, $\sup_x D(x, \mu) = D(\theta_0, \mu)$ and so Condition 1 is satisfied. Next, HD satisfies Condition 3 with

$K = 1$ and $\text{VC}(\mathcal{F}) = d + 2$, as shown in Lemma 44. Furthermore, we have that $\text{HD}(\cdot, \mu)$ is σ^{-1} -Lipschitz. Therefore, the conditions of Corollary 10 are satisfied. Plugging these values into Corollary 10 and setting $\gamma = 1/3$, as in Theorem 6.5 of (Kamath et al., 2019), yields

$$n \gtrsim \frac{d \log \left(\frac{K}{\alpha(t)} \vee e \right)}{\alpha(t)^2} \vee \frac{d \log \left(\frac{R/\sigma}{d_{R,\theta_\pi}(\theta_0) \wedge \alpha(t)} \vee d \right)}{\alpha(t)\epsilon}. \quad (4.8)$$

Note that in this case, we have that $\alpha(t) = \Phi(t/\sigma) - 1/2$. Let $a(R) = (\Phi(R/\sigma) - 1/2)/R$. Now, since $t \leq R$, concavity of Φ on \mathbb{R}^+ implies that $\alpha(t) \geq a(R)t$. Applying this fact in conjunction with (4.8) results in

$$n \gtrsim \frac{d \log \left(\frac{1}{a(R)t} \vee e \right)}{a(R)^2 \cdot t^2} \vee \frac{d \log \left(\frac{R/\sigma}{d_{R,\theta_\pi}(\theta_0) \wedge a(R)t} \vee d \right)}{a(R) \cdot t \cdot \epsilon}, \quad (4.9)$$

which matches the bound given in Theorem 6.5 of (Kamath et al., 2019) up to logarithmic factors in d and t^{-1} . A simple manipulation of the terms in (4.9) gives an equivalent deviations bound, which matches the minimax lower bounds given by Cai et al. (2021), up to logarithmic factors, if δ is taken to be n^{-k} for some $k > 0$. ■

5. Inverse discrepancy function for small t

As stated earlier, when applying Theorem 5, if α^{-1} is often at most linear, at least for small t , then this is very convenient technically, in that it simplifies the bounds given in Theorem 5. This is because, in that case, we need not be concerned about inverting α directly. Indeed, if there exists $C_\mu > 0$ such that $\alpha(t) \geq C_\mu t$ for small t , then this implies that the $\alpha^{-1}(t)$ in Theorem 5 can be replaced by t/C_μ , thus, simplifying the bound greatly. The purpose of this section is to give conditions for there to exist $C_\mu > 0$ such that $\alpha(t) \geq C_\mu t$ for small t . We first give general conditions under which $\alpha(t) \geq C_\mu t$ for an arbitrary depth function.

Lemma 19 *Suppose that $\sup_{x \in \mathbb{R}^d} D(x, \mu) = D(\theta_0, \mu)$, $D(\cdot, \mu)$ is translation invariant, decreasing along rays, continuous and continuously differentiable on $B_r(\theta_0)$ for some $r > 0$. In addition, suppose that there exists $0 < a < R < r$ such that $u_* = \arg\max_{u \in \mathbb{S}^{d-1}} D(u \cdot t, \mu)$ is constant in t for $a \leq t \leq R$. Then, if for $a \leq t \leq R$, it holds that*

$$g^*(t) = u_*^\top \left[\frac{d}{dx} D(x, \mu) \Big|_{x=t \cdot u_*} \right]$$

is non-decreasing on $[a, R]$, we have that $\alpha(t) \geq t\alpha(R)/R$.

Note that the proofs of results in this section are deferred to Appendix B. The conditions of Lemma 19 will hold when μ and D are such that the contours of the depth function are not changing in shape as we move away from the center. For example, the conditions of Lemma 19 hold for the class of elliptical distributions and the integrated depths. We give a more general version of this result below, using Lemma 19 to show a general bound for integrated depth functions. In addition, recall that ν denotes the uniform measure on the $(d-1)$ -dimensional unit sphere.

Lemma 20 Suppose μ is centrally symmetric about θ_0 and $D(x, \mu) = \int_{\mathbb{S}^{d-1}} g(x^\top u, \mu_u) d\nu$ where for all $u \in \mathbb{S}^{d-1}$, we have that

1. $g(ay + b, a\mu_u + b) = g(y, \mu_u)$,
2. $g(\cdot, \mu_u)$ are bounded, continuous and continuously differentiable,
3. $g(\theta_0^\top u, \mu_u) = 0$ and $g(\theta_0^\top u - y, \mu_u) = g(\theta_0^\top u + y, \mu_u)$,
4. there exists $R > 0$ such that $g(\theta_0^\top u + y, \mu_u)$ is non-increasing for $0 \leq y \leq R$ and
5. there exists $v \in \mathbb{S}^{d-1}$ such that $g(c(\theta_0 + v)^\top u, \mu_u) \geq g(c(\theta_0 + w)^\top u, \mu_u)$ for all $c > 0$ and $w \in \mathbb{S}^{d-1}$,

then for $t \in [0, R]$ it holds that

$$\alpha(t) \geq t\alpha(R)/R.$$

Observe that Items 1-4 are satisfied for the smoothed integrated dual depth (Definition 31) and the integrated-rank-weighted depth (Definition 26) provided the projected cumulative distribution functions $h(y, u) = \mu(X^\top u \leq y)$ are continuous and continuously differentiable in y for all $u \in \mathbb{S}^{d-1}$. Item 5 depends on the distribution, but is fairly general. For instance, it is easy to see that item 5 is satisfied by the class of elliptical distributions, or anytime μ is such that the depth function has nested, convex contours.

In the univariate setting, deviations bounds for the univariate median often only require a lower bound on the density of the population measure in a neighborhood of the median. We next show a result of this flavor, for the halfspace depth. Recall that a measure $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is halfspace symmetric about $\theta_0 \in \mathbb{R}^d$ if for every closed halfspace H which contains θ_0 , it holds that $\mu(X \in H) \geq 1/2$ (Zuo and Serfling, 2000).

Lemma 21 Suppose that $D = HD$, $\arg\max_{x \in \mathbb{R}^d} HD(x, \mu) = \theta_0$ and suppose that there is some $R > 0$ such that for all $u \in \mathbb{S}^{d-1}$, μ_u are halfspace symmetric and absolutely continuous. If there exists $C_\mu > 0$ such that for all $u \in \mathbb{S}^{d-1}$, $\inf_{\theta_0^\top u \pm R} f_{\mu_u}(x) \geq C_\mu$, then for all $\theta_0^\top u - R \leq t \leq \theta_0^\top u + R$, we have that $\alpha(t) \geq tC_\mu$.

6. Depth functions and their properties

For the convenience of the reader, we collect here some basic facts about depth functions. We first review several common depth functions in the literature and their properties. We then introduce a new depth function, the smoothed integrated dual depth, which has desirable properties from a computational perspective and is used in our simulations. For proofs of the results from this section, see Appendix F.

One classical depth function is the simplicial depth (Liu, 1988, 1990). Let $\Delta(x_1, \dots, x_{d+1})$ be the simplex in \mathbb{R}^d with vertices x_1, \dots, x_{d+1} .

Definition 22 (Simplicial Depth) Suppose that Y_1, \dots, Y_{d+1} are i.i.d. from $\mu \in \mathcal{M}_1(\mathbb{R}^d)$. The simplicial depth of a point $x \in \mathbb{R}^d$ with respect to μ is

$$\text{SMD}(x, \mu) = \Pr(\Delta(Y_1, \dots, Y_{d+1}) \ni x).$$

The simplicial depth of a point $x \in \mathbb{R}^d$ is the probability that a random simplex defined by $d+1$ draws from μ contains x . The next depth function we introduce is spatial depth (Vardi and Zhang, 2000; Serfling, 2002), which is maximized at the spatial median (Vardi and Zhang, 2000). Define the spatial sign function some as $\text{SR}(y) = y/\|y\|$ (with $\text{SR}(0) = 0$). We can then define the spatial rank function as $\mathbb{E}_\mu \text{SR}(x - X)$. The norm of the spatial rank of $x \in \mathbb{R}^d$ is a measure of the outlyingness of x with respect to μ . With this in mind, the spatial depth is defined as follows:

Definition 23 *The spatial depth SD of a point $x \in \mathbb{R}^d$ with respect to $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is $\text{SD}(x, \mu) = 1 - \|\mathbb{E}_\mu \text{SR}(x - X)\|$.*

We will see below that the spatial depth is (K, \mathcal{F}) -regular, however, $K = O(d)$. Replacing the norm with its square eliminates this inconvenience, and still yields a depth function which is maximized at the spatial median. Therefore, we define the modified spatial depth as follows:

Definition 24 *The modified spatial depth MSD of a point $x \in \mathbb{R}^d$ with respect to $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is $\text{MSD}(x, \mu) = 1 - \|\mathbb{E}_\mu \text{SR}(x - X)\|^2$.*

Lastly, we define the integrated depth functions. Here, the idea is to define the depth of a point $x \in \mathbb{R}^d$ as the average, univariate depth of $x^\top u$ over all directions u . The first integrated depth function⁶ is the integrated dual depth of Cuevas and Fraiman (2009), which, letting ν be the uniform measure on \mathbb{S}^{d-1} , is defined as follows:

Definition 25 (Integrated Dual Depth) *The integrated dual depth of a point $x \in \mathbb{R}^d$ with respect to $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is*

$$\text{IDD}(x, \mu) = \int_{\mathbb{S}^{d-1}} F(x, u, \mu) (1 - F(x, u, \mu)) d\nu(u). \quad (6.1)$$

The integrated dual depth of $x \in \mathbb{R}^d$ is the average univariate simplicial depth of $x^\top u$ over all directions $u \in \mathbb{S}^{d-1}$. The integrated rank-weighted depth (Ramsay et al., 2019) replaces the simplicial depth in (6.1) with a slightly modified version of halfspace depth. For $X \sim \mu$, define $F(x, u, \mu) = \mu(X^\top u < x^\top u)$.

Definition 26 (Integrated Rank-Weighted Depth) *The integrated rank-weighted depth of a point $x \in \mathbb{R}^d$ with respect to $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is $\text{IRW}(x, \mu) = 2\mathbb{E}_\nu F(x, U, \mu) \wedge (1 - F(x, -U, \mu))$.*

We now turn to a brief discussion of these functions satisfy Definition 2. To this end, recall the following classes of measures. Recall that a measure μ is *angularly symmetric* about a point $\theta_0 \in \mathbb{R}^d$ if $\frac{X - \theta_0}{\|X - \theta_0\|} \stackrel{d}{=} \frac{\theta_0 - X}{\|\theta_0 - X\|}$ for $X \sim \mu$ (Liu, 1990). Let $\mathcal{N} \subset \mathcal{M}_1(\mathbb{R}^d)$ denote the set of angularly symmetric measures over \mathbb{R}^d that are absolutely continuous with respect to the Lebesgue measure. The following proposition collects existing results in the literature:

6. The first integrated multivariate depth function, to be precise. The first integrated depth was developed for functional data by Fraiman and Muniz (2001).

Proposition 27 *The aforementioned depth functions satisfy the following:*

- *Halfspace depth is a depth function with admissible set $\mathcal{M}_1(\mathbb{R}^d)$.*
- *Simplicial depth is a depth function with admissible set \mathcal{N} .*
- *Spatial depth and the modified spatial depth satisfy properties (1), (2) and (4) for all $\mu \in \mathcal{M}_1(\mathbb{R}^d)$.*
- *Integrated dual depth and integrated rank-weighted depth are depth functions with admissible set \mathcal{C} . Furthermore, they satisfy properties (1), (2) and (4) for any $\mu \in \mathcal{M}_1(\mathbb{R}^d)$.*

Remark 28 *Spatial depth does not satisfy Definition 2, as property (3) fails. However, spatial depth satisfies spatial angle monotonicity and spatial rank monotonicity (Serfling, 2019, see Section 4.3), which together provide a weaker alternative to property (3). Simplicial depth and halfspace depth are affine invariant, which is a stronger version of property (1).*

Remark 29 *Both the halfspace depth and the simplicial depth function have an exact computational running time of $O(n^{d-1})$, see, e.g., (Dyckerhoff and Mozharovskiy, 2016; Afshani et al., 2016) and the references therein. By contrast, the spatial depth and the integrated depth functions are computable in high dimensions (Chaudhuri, 1996; Ramsay et al., 2019). For instance, the integrated depth functions can be approximated easily via sampling, say, m points from ν . One can then approximate the sample depth value of a given point in $O(mnd)$ time (Ramsay et al., 2019). Going further, it is easy to show via Talagrand's inequality (Talagrand, 1994; Sen, 2022) that if one wants to approximate a sample depth value to within error t , with > 0.99 probability then we need $m \gtrsim nd \log(1/t)/t^2$ unit vectors, which results in $O(n^2 d^2 \log(1/t)/t^2)$ time.*

The following theorem demonstrates that the depth functions defined in this section satisfy the conditions of Theorem 12. Let μ_u be the law of $X^\top u$ if $X \sim \mu$.

Theorem 30 *The aforementioned depth functions satisfy the following:*

- *Conditions 1 and 3 hold for each of the depth functions defined in Section 6.*
- *Suppose that π is absolutely continuous and that for all $u \in \mathbb{S}^{d-1}$, the measure μ_u is absolutely continuous with density f_u and $\sup_{\|u\|=1} \|f_u\|_\infty = L < \infty$. Then there exists a universal constant $C > 0$ such that for all $s \in (0, \infty]$, halfspace, simplicial, integrated rank-weighted and the smoothed integrated dual depth⁷ satisfy Condition 4 with $\omega(r) = C \cdot L \cdot r$.*
- *Suppose that μ has a bounded density and that $\sup_y \mathbb{E} \|y - X\|^{-1} = L' \leq \infty$. Then modified spatial depth satisfies Condition 4 with $\omega(r) = 2\sqrt{d} \cdot L' \cdot r$ and spatial depth satisfies Condition 4 with $\omega(r) = 2 \cdot L' \cdot r$.*

The above results are also summarized in Table 2.4.

7. The smoothed integrated dual depth is introduced in Section 6.1 below and equals the integrated dual depth when $s = \infty$.

6.1 The smoothed integrated dual depth

One issue with many of the aforementioned depth functions is that their empirical versions contain indicator functions, which are non-smooth. This fact is inconvenient from an optimization perspective, e.g., we cannot apply gradient descent to estimate the median. In addition, this non-smoothness implies that the empirical depth functions will possess flat contours. This can cause the estimated medians to perform poorly when the population measure has atoms near its center (Lalanne et al., 2023). Therefore, to resolve this issue, recall that $\mathbb{1}\{X_i^\top u \leq x^\top u\} \approx \sigma(s(x - X)^\top u)$, for large s , where σ is the usual sigmoid function: $\sigma(x) = (1 + e^{-x})^{-1}$. This motivates the following definition of depth, which we use in our simulation study from Section 2.5. Recall that ν denotes the uniform measure on the $(d - 1)$ -dimensional unit sphere.

Definition 31 (Smoothed Integrated Dual Depth) *The smoothed integrated dual depth with smoothing parameter $s > 0$ of a point $x \in \mathbb{R}^d$ with respect to $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is*

$$\text{IDD}(x, \mu, s) = \int_{\mathbb{S}^{d-1}} \mathbb{E} \sigma(s(x - X)^\top u) \left(1 - \mathbb{E} \sigma(s(x - X)^\top u)\right) d\nu(u). \quad (6.2)$$

As $s \rightarrow \infty$, the smoothed integrated dual depth converges to the integrated dual depth. However, the smoothed integrated dual depth is convenient in that its empirical version is differentiable. Furthermore, the smoothed integrated dual depth is a depth with admissible set \mathcal{C} .

Proposition 32 *For any $s > 0$, the smoothed integrated dual depth is a depth function with admissible set \mathcal{C} . In addition, properties (1) and (4) hold for any $\mu \in \mathcal{M}_1(\mathbb{R}^d)$.*

In the simulations presented in Section 2.5, we approximate (6.2) with

$$\widehat{\text{IDD}}(x, \mu, s) = \frac{1}{M} \sum_{m=1}^M \mathbb{E} \sigma(s(x - X)^\top u_m) \left(1 - \mathbb{E} \sigma(s(x - X)^\top u_m)\right),$$

where u_1, \dots, u_M are drawn i.i.d. from ν . One can show that the number of unit vectors needed to maintain a fixed error level $t > 0$ with high probability $1 - \gamma$ is polynomial in n and d .

Proposition 33 *Let u_1, \dots, u_M be i.i.d. from ν . Then for all $d, s, t > 0$ and all $0 < \gamma < 1$, there exists a universal constant $c_1 > 0$ such that $\sup_x \left| \widehat{\text{IDD}}(x, \hat{\mu}_n, s) - \text{IDD}(x, \hat{\mu}_n, s) \right| \leq t$, with probability $1 - \gamma$, provided that*

$$M \geq c_1 [\log(1/\gamma) \vee dn \log(1/t \vee e)] t^{-2}.$$

Observe that Proposition 33 is a sample complexity style result, where instead of bounding the number of samples needed for achieve some error tolerance $t > 0$, we bound the number of simulated unit vectors needed to achieve error tolerance $t > 0$. Specifically, the result says that we need, roughly, $O(nd/t^2)$ unit vectors for our approximation to be within t of $\text{IDD}(x, \hat{\mu}_n, s)$, with probability > 0.99 . Lastly, note that the median estimators are relatively insensitive to the choice of s . For $s \geq 10$, the specific value of s appears to have minimal effect on the estimation error and convergence of the gradient descent algorithm, see Appendix C for more details. We recommend choosing $s = 100$.

7. Computing private depth values

The depth values themselves can be of interest. For example, they are used in visualization and a host of inference procedures (Li and Liu, 2004; Chenouri et al., 2011). The theoretical analysis from the previous sections yields simple differentially private methods for estimating depth values. If we want to compute the depth of a known point, then we can use an additive noise mechanism, such as the Laplace mechanism or the Gaussian mechanism (Dwork et al., 2006), i.e., we can just add noise that is calibrated to ensure differential privacy. For example, it follows from privacy of the Laplace mechanism (Dwork et al., 2006) that if D is (K, \mathcal{F}) -regular, then for all $x \in \mathbb{R}^d$ it holds that $\tilde{D}(x, \hat{\mu}_n) = D(x, \hat{\mu}_n) + W_1 K/n\epsilon$, $W_1 \sim \text{Laplace}(1)$, is ϵ -differentially private. Theorem 12 then implies the concentration of private depth values generated from the Laplace mechanism.

Corollary 34 *Suppose that D satisfies Condition 3. For $x \in \mathbb{R}^d$ chosen independently of the data, $\tilde{D}(x, \hat{\mu}_n) = D(x, \hat{\mu}_n) + W_1 k/n\epsilon$, is ϵ -differentially private. In addition, there exists universal constants $c_1, c_2 > 0$ such that for all $t \geq 0$ it holds that*

$$\Pr(|\tilde{D}(x, \hat{\mu}_n) - D(x, \mu)| > t) \leq (c_1 n / \text{VC}(\mathcal{F}))^{\text{VC}(\mathcal{F})} e^{-c_2 n t^2 / K^2} + e^{-n \epsilon t / 2K}.$$

Corollary 34 suggests taking $\epsilon \propto t/K$ when simulating private depth values. An obvious problem of interest is estimating the vector of depth values at the sample points, i.e., can we estimate $\hat{\mathbf{D}}(\hat{\mu}_n) = (D(X_1, \hat{\mu}_n), \dots, D(X_n, \hat{\mu}_n))$ privately? Define the global 1-sensitivity of $\hat{\mathbf{D}}(\hat{\mu}_n)$ to be

$$\text{GS}_n(\hat{\mathbf{D}}(\hat{\mu}_n)) = \sup_{\substack{\mathbb{X}_n \in \mathbf{D}_{n \times d}, \\ \hat{\mu}_n \in \hat{\mathcal{M}}(\mathbb{X}_n)}} \left\| \hat{\mathbf{D}}(\hat{\mu}_n) - \hat{\mathbf{D}}(\tilde{\mu}_n) \right\|_1,$$

where $\|\cdot\|_1$ denotes the L_1 norm. In general, $\text{GS}_n(\hat{\mathbf{D}}(\hat{\mu}_n)) = O(1)$, see Lemma 39. As a result, if \mathbf{D} is the vector of private depth values generated from the Laplace mechanism then, $\|\tilde{\mathbf{D}}(\hat{\mu}_n) - \hat{\mathbf{D}}(\hat{\mu}_n)\| \gtrsim n/\epsilon$ with high probability; the level of noise is greater than that of the sampling error. Here, $\|\cdot\|$ denotes the L_2 norm. This result is intuitive; these vectors reveal more information about the population as n grows, which differs markedly from the single depth value case, where the amount of information received is fixed in n . In fact, for large n the vector of depth values at the sample points contains a significant amount of information about μ ; the depth function can, under certain conditions, characterize the distribution of the input measure μ (see Struyf and Rousseeuw, 1999; Nagy, 2021, and the references therein). In order to release so much information about the population privately, we must inject non-negligible noise. We cannot, then, simply plug in the n private sample depth values into an inference procedure and proceed. An interesting topic for future research is how to extend depth-based inference procedures to the private setting.

Acknowledgments

The authors acknowledge Gautam Kamath for his helpful comments and discussion, which improved the paper. The authors acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). A.J. acknowledges the support of the Canada

Research Chairs Program. Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), [DGEER-2023-00311, RGPIN-2020-04597, DGEER-2020-00199] et du Programme des chaires de recherche du Canada [CRC-2022-00142]

Appendix A. Notation conventions

We collect here some notation we will use throughout the paper. Note that all notations are introduced throughout the paper, but having the notation collected in one place may be a useful reference. For sets A, B , we let $A \triangle B$ be the symmetric difference between A and B . Given $v \in \mathbb{R}^d$, let $\|v\|$ denote the Euclidean norm, and for a function $f: \mathbb{R} \rightarrow \mathbb{R}$, let $\|f\| = \sup_{x \in \mathbb{R}} |f(x)|$. For a set A and $r > 0$, let $B_r(A) = \{x: \min_{y \in A} \|x - y\| \leq r\}$. Let $A_R(y)$ be the d -dimensional cube of side-length R centered at y and let $d_{R,y}(x)$ denote the minimum distance from a point $x = (x_1, \dots, x_d)$ to a face of the cube $A_R(y)$: $d_{R,y}(x) = \min_{1 \leq i \leq d} |x_i - y_i \pm R/2|$. Similarly, denote the minimum distance from a set $B \subset \mathbb{R}^d$ to a face of the cube $A_R(y)$ by $d_{R,y}(B) = \inf_{x \in B} d_{R,y}(x)$ and let $d(x, B) = \inf_{y \in B} \|x - y\|$ denote the usual point-to-set distance. For two numbers $a, b \in \mathbb{R}$, we have that $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$ and we write $a \lesssim b$ ($a \gtrsim b$) if $a \leq Cb$ ($a \geq Cb$) for some universal constant $C > 0$.

For a space S , $\mathcal{M}_1(S)$ denotes the space of probability measures on S . We say that a data set of size $n \times d$ is a collection of n points in \mathbb{R}^d , $\mathbb{X}_n = (X_\ell)_{\ell=1}^n$, with repetitions allowed. We assume that we have observed a data set of size $n \times d$, which are realizations of n , independent and identically distributed random variables from a population measure $\mu \in \mathcal{M}_1(\mathbb{R}^d)$. The empirical measure of these observations is denoted by $\hat{\mu}_n$. We let \mathcal{B} denote the space of Borel functions from \mathbb{R}^d to $[0, 1]$. For a family of functions, $\mathcal{F} \subseteq \mathcal{B}$ let $\text{VC}(\mathcal{F})$ denote the Vapnik—Chervonenkis dimension of \mathcal{F} . We sometimes use the following pseudometric on $\mathcal{M}_1(\mathbb{R}^d)$, $d_{\mathcal{F}}(\mu, \nu) = \sup_{g \in \mathcal{F}} |\int g d(\mu - \nu)|$, where $\mu, \nu \in \mathcal{M}_1(\mathbb{R}^d)$.

Next, we let $\mathcal{N}(\mu, \Sigma)$ denote the Gaussian measure with mean μ and covariance matrix Σ . We let \mathbb{S}^{d-1} denote the $(d-1)$ -dimensional unit sphere and we let ν be the uniform measure on \mathbb{S}^{d-1} , where the dimension is clear from the context. We write $X \sim \mu$ if the random vector X is drawn from $\mu \in \mathcal{M}_1(\mathbb{R}^d)$. Furthermore, we let μ_u denote the law of $X^\top u$ for $u \in \mathbb{S}^{d-1}$ and f_u be the associated density function (when it exists). For two random vectors X, Y , we write $X \stackrel{d}{=} Y$ if X is equal in distribution to Y . For a random vector $X \sim \mu$, we write its expectation as $\mathbb{E}_\mu X = \int X d\mu$ and omit the measure, μ , when it is clear from the context. Recall that a measure $\mu \in \mathcal{M}_1(\mathbb{R}^d)$ is centrally symmetric about a point $x \in \mathbb{R}^d$ if $X - x \stackrel{d}{=} x - X$ for $X \sim \mu$ and a measure μ is angularly symmetric about a point $\theta_0 \in \mathbb{R}^d$ if $\frac{X - \theta_0}{\|X - \theta_0\|} \stackrel{d}{=} \frac{\theta_0 - X}{\|\theta_0 - X\|}$ for $X \sim \mu$ (Liu, 1990). We let $\mathcal{C} \subset \mathcal{M}_1(\mathbb{R}^d)$ denote the set of centrally symmetric measures over \mathbb{R}^d and we let $\mathcal{A} \subset \mathcal{M}_1(\mathbb{R}^d)$ denote the set of angularly symmetric measures over \mathbb{R}^d that are absolutely continuous with respect to the Lebesgue measure.

Appendix B. Technical proofs

Proof [Proof of Corollary 10] First note that by assumption, all of the conditions of Lemma 35 are satisfied. To prove (2.6) and (2.7) we apply Lemma 16 and Lemma 17, respectively, in conjunction with Lemma 35. To apply Lemma 16, first, observe that together, the assumptions that $L > 1$, $\sigma_\pi^2 \geq 1/16$ and $D(x; \mu) \leq 1$ imply that $\alpha(t)/4L\sigma_\pi \leq 1$. Applying Lemma 16 with $R = \alpha(t)/4L$ and $E = E_{D,\mu}$ in conjunction with (B.11) yields that there exists a universal constant $c > 0$ such that for all $t \geq 0$, all $d > 2$ and all $0 < \gamma < 1$, we have that $d(E_{D,\mu}, \tilde{\theta}_n) < t$ with probability at least $1 - \gamma$ provided that

$$n \gtrsim K^2 \left[\frac{\log(1/\gamma) \vee (\text{VC}(\mathcal{F}) \log(\frac{K}{c\alpha(t)}) \vee 1)}{\alpha(t)^2} \right] \bigvee \left[K \frac{\log(1/\gamma) \vee \frac{d(E_{D,\mu}, \theta_\pi)^2}{\sigma_\pi^2} \vee d \log\left(\frac{L\sigma_\pi}{\alpha(t)} \vee d\right)}{\epsilon \alpha(t)} \right].$$

To show (2.7), we apply Lemma 17 with $r = \alpha(t)/4L$ and $E = E_{D,\mu}$ in conjunction with (B.11), which yields that there exists a universal constant $c_1 > 0$ such that for all $t \geq 0$, all $d \geq 1$ and all $0 < \gamma < 1$, we have that $d(E_{D,\mu}, \tilde{\theta}_n) < t$ with probability at least $1 - \gamma$ provided that

$$n \gtrsim K^2 \left[\frac{\log(1/\gamma) \vee (\text{VC}(\mathcal{F}) \log(\frac{K}{c\alpha(t)}) \vee 1)}{\alpha(t)^2} \right] \bigvee \left[K \frac{\log(1/\gamma) \vee d \log\left(\frac{R}{d_{R,\theta_\pi}(E_{D,\mu}) \wedge \alpha(t)/L}\right)}{\alpha(t)\epsilon} \right],$$

which completes the proof. \blacksquare

Proof [Proof of Lemma 19] Without loss of generality, assume that $\theta_0 = 0$. It suffices to show that $\alpha(t)$ is concave on $[a, R]$. This holds if and only if $g(t) = \sup_{\|x\| \geq t} D(x, \mu)$ is convex on $[0, R]$. By the decreasing along rays property and the fact that D is maximized at 0, it holds that $g(t) = \sup_{\|u\|=1} D(t \cdot u, \mu)$. Next, let $u_{*,t} = \arg\max_{\|u\|=1} D(t \cdot u, \mu)$. Observe that by assumption $u_{*,t} = u_*$, i.e., $u_{*,t}$ is constant in t . Therefore, $\sup_{\|u\|=1} D(t \cdot u, \mu) = D(t \cdot u_*, \mu)$. Next, it holds that

$$\frac{d}{dt} g(t) = \frac{d}{dt} D(t \cdot u_*, \mu) = u_*^\top \left[\frac{d}{dx} D(x, \mu) \Big|_{x=t \cdot u_*} \right],$$

which is non-decreasing by assumption. This implies that $g(t)$ is convex on $[a, R]$. It follows that $\alpha(t) \geq t\alpha(R)/R$. \blacksquare

Proof [Proof of Lemma 20] Again, without loss of generality, take $\theta_0 = 0$. We have that

$$\frac{d}{dx} \int_{\mathbb{S}^{d-1}} g(x^\top u, \mu_u) d\nu = \int_{\mathbb{S}^{d-1}} \frac{d}{dx} g(x^\top u, \mu_u) d\nu = \int_{\mathbb{S}^{d-1}} g^{(1)}(x^\top u, \mu_u) u d\nu.$$

Now,

$$u_*^\top \left[\frac{d}{dx} D(x, \mu) \Big|_{x=t \cdot u_*} \right] = \int_{\mathbb{S}^{d-1}} u_*^\top u g^{(1)}(t \cdot u_*^\top u, \mu_u) d\nu = 2 \int_{\mathbb{S}^{d-1}} |u_*^\top u| g^{(1)}(t \cdot |u_*^\top u|, \mu_u) d\nu,$$

which is non-decreasing by the assumptions on g . Now, assumptions 1–5 and the preceding observation imply the conditions of Lemma 19 are satisfied, and the result follows. \blacksquare

Proof [Proof of Lemma 21] Given that HD is translation invariant, there is no loss of generality in assuming that $\theta_0 = 0$. First, by definition and the decreasing along rays property of HD, we have that

$$\alpha(t) = \text{HD}(0, \mu) - \sup_{\|x\| \geq t} \text{HD}(x, \mu) = 1/2 - \sup_{\|x\|=t} \inf_{\|u\|=1} F(x, u, \mu). \quad (\text{B.1})$$

Next, we bound the term $\sup_{\|x\|=t} \inf_{\|u\|=1} F(x, u, \mu)$ above. To this end, we have that

$$\sup_{\|x\|=t} \inf_{\|u\|=1} F(x, u, \mu) \leq \sup_{\|u\|=1} \left(\frac{1}{2} - \left| \frac{1}{2} - h(t, u) \right| \right).$$

This inequality, in combination with (B.1) and the mean value theorem yields that for all $0 \leq t \leq R$, it holds

$$\alpha(t) \geq \left| \frac{1}{2} - h(t, u) \right| \geq tC_\mu,$$

which completes the proof. ■

Proof [Proof of Lemma 14] First, we use the derivative test to determine when the function $h(x)$ is increasing. Doing so yields that $h(x)$ is increasing for $x \geq b/a$. Now, it remains to find $x \geq b/a$ such that $h(x) \geq 0$. In other words, for $y_r = rb/a$, we want to find $r \geq 1$ such that $h(y_r) \geq 0$. By definition,

$$h(y_r) = b(r - \log(cr/a)).$$

It follows that $h(y_r) \geq 0$ is satisfied for r such that $r - \log r \geq \log(c/a)$. Applying the inequality $x - \log x \geq x/2$ for $x \geq 1$ yields that $r - \log r \geq \log(c/a)$ is satisfied when

$$r \geq 2 \log(c/a).$$

Therefore, $h(x)$ is positive and increasing for

$$x \geq 2 \frac{b}{a} \log \left(\frac{c}{a} \vee \sqrt{e} \right).$$
■

Proof [Proof of Lemma 16] Let $R_* = R^2/\sigma_\pi^2$, $\theta_0 \in E$ and let $W \sim \chi_d^2(\|\theta_0 - \theta_\pi\|^2/\sigma_\pi^2)$. We can lower bound $\pi(B_R(E))$ as follows:

$$\pi(B_R(E)) \geq \pi(\|X - \theta_0\| \leq R) = \Pr(W \leq R_*). \quad (\text{B.2})$$

Let $G(x, k, \lambda)$ be the cumulative distribution function for the non-central chi-squared measure with k degrees of freedom and non-centrality parameter λ . Recall that for all $x, k, \lambda > 0$, it holds that

$$G(x, k, \lambda) = e^{-\lambda^2} \sum_{j=0}^{\infty} \frac{(\lambda/2)^j}{j!} G(x, k+2j) \geq e^{-\lambda^2} G(x, k, 0).$$

Therefore,

$$\Pr(W \leq R_*) \geq e^{-\|\theta_0 - \theta_\pi\|^2/2\sigma_\pi^2} G(R_*, d, 0). \quad (\text{B.3})$$

We now lower bound $G(x, d, 0)$ for $x \leq 1$. Let γ be the lower incomplete gamma function. It follows from the properties of the gamma function that

$$-\log G(x, d, 0) = -\log \gamma(x/2, d/2) + \log \Gamma(d/2) \lesssim -\log \gamma(x/2, d/2) + d \log d.$$

Note that the assumption $d > 2$ implies that $\sup_y \gamma(y, d/2) \geq 1$. This implies that $y^{d/2-1}e^{-y/2}$ is increasing on $(0, x)$. Thus, for any $0 < r < 1$ it holds that

$$\begin{aligned} -\log G(x, d, 0) &\lesssim -\log \int_0^x (y/2)^{d/2-1} e^{-y/2} / 2 \, dy + d \log d \\ &\lesssim -\log(x-r)(r/2)^{d/2-1} / 2 + r/2 + d \log d. \end{aligned}$$

Setting $r = x/2$ and using the fact that $x \leq 1$ results in

$$-\log G(x, d, 0) \lesssim \frac{x}{4} + \frac{d}{2} \log \left(\frac{4}{x} \right) + d \log d \lesssim d \log \left(\frac{4}{x} \vee d \right), \quad (\text{B.4})$$

Note that $R_*^2 = R^2/\sigma_\pi^2 \leq 1$ by assumption. Thus, applying (B.4) results in

$$-\log G(R, d, 0) \lesssim d \log \left(\frac{\sigma_\pi}{R} \vee d \right). \quad (\text{B.5})$$

Combining (B.2), (B.3) and (B.5) implies that

$$-\log \pi(B_R(E)) \lesssim \frac{\|\theta_0 - \theta_\pi\|^2}{\sigma_\pi^2} + d \log \left(\frac{\sigma_\pi}{R} \vee d \right).$$

Now, note that this bound holds for all $\theta_0 \in E$, therefore, it holds that

$$-\log \pi(B_R(E)) \lesssim \frac{d(E, \theta_\pi)^2}{\sigma_\pi^2} + d \log \left(\frac{\sigma_\pi}{R} \vee d \right).$$

■

Proof [Proof of Lemma 17] First, note that for any ball $B_r(x)$ such that $B_r(x) \subset A_R(\theta_\pi)$, we have that

$$\pi(B_r(x)) = \frac{r^d \pi^{d/2}}{R^d \Gamma(d/2 + 1)} \geq \left(\frac{r}{R} \right)^d (d/2 + 1)^{d/2+1} \pi^{d/2} e^{-d/2} \geq \left(\frac{r}{R} \right)^d \left(\frac{\pi d}{2e} \right)^{d/2}.$$

Taking the negated logarithm of both sides implies that

$$-\log \pi(B_r(x)) \lesssim d \log \left(\frac{R}{r} \right) - d \log \pi d/2 + d \lesssim d \log \left(\frac{R}{r} \right). \quad (\text{B.6})$$

If there exists $\theta_0 \in E$ such that $B_r(\theta_0) \subset A_R(\theta_\pi)$ then (B.6) implies that

$$-\log \pi(B_r(\theta_0)) \lesssim d \log \left(\frac{R}{r} \right). \quad (\text{B.7})$$

If instead there exists $\theta_0 \in E$ such that $A_R(\theta_\pi) \subset B_r(\theta_0)$ then

$$-\log \pi(B_r(\theta_0)) = 0. \quad (\text{B.8})$$

Suppose finally that for all $\theta_0 \in E$ it holds that both $B_r(\theta_0) \not\subset A_R(\theta_\pi)$ and $A_R(\theta_\pi) \not\subset B_r(\theta_0)$. Let $\theta_{0,j}$ and $\theta_{p,j}$ be the j^{th} coordinate of θ_0 and θ_π respectively. where

$$r = \min_{1 \leq j \leq d} (|\theta_{0,j} - \theta_{p,j} - R/2| \wedge |\theta_{0,j} - \theta_{p,j} + R/2|) = d_{R,\theta_\pi}(\theta_0).$$

Now, by assumption, there exists $\theta_0 \in E$ such that $\theta_0 \in A_R$. For such $\theta_0 \in A_R$, the definition of d_{R,θ_π} immediately implies that $B_{d_{R,\theta_\pi}(\theta_0)}(\theta_0) \subset A_R(\theta_\pi)$. Therefore, (B.6) implies that

$$-\log \pi(B_r(E)) \leq -\log \pi(B_r(\theta_0)) \lesssim d \log \left(\frac{R}{d_{R,\theta_\pi}(\theta_0)} \right). \quad (\text{B.9})$$

It follows from (B.7), (B.8) and (B.9) that for all $d > 0$ and all r it holds that

$$-\log \pi(B_r(E)) \lesssim d \log \left(\frac{R}{d_{R,\theta_\pi}(\theta_0) \wedge r} \right). \quad (\text{B.10})$$

Note that (B.10) holds for all $\theta_0 \in E$, thus, for all $d > 0$ and all $r > 0$ it holds that

$$-\log \pi(B_r(E)) \lesssim d \log \left(\frac{R}{d_{R,\theta_\pi}(E) \wedge r} \right).$$

■

Lemma 35 *Suppose that Conditions 1 and 3 hold with $\phi = D$ for some $K > 0$. In addition, suppose that the map $x \mapsto D(x, \mu)$ is L -Lipschitz for some $L > 0$. There is a universal $c > 0$ such that for all $t \geq 0$, all $d \geq 1$ and all $0 < \gamma < 1$, we have that $\|\tilde{\theta}_n - E_{D,\mu}\| < t$ with probability at least $1 - \gamma$ provided that*

$$n \gtrsim \left[K^2 \frac{\log(1/\gamma) \vee (\text{VC}(\mathcal{F}) \log(\frac{K}{c\alpha(t)}) \vee 1)}{\alpha(t)^2} \right] \bigvee \left[K \frac{\log(1/\gamma) \vee -\log \pi(B_{\alpha(t)/4L}(E_{D,\mu}))}{\epsilon \alpha(t)} \right]. \quad (\text{B.11})$$

Proof The proof is similar to that of Lemma 15. Note that all of the conditions of Theorem 12 hold, and so we can apply Theorem 12 with $\beta = n\epsilon/2K$. Proving the sample complexity bound then amounts to finding n which ensures that $I + II \leq \gamma$, where I, II are defined in the proof of Lemma 15. Now, following the same argument in the proof of Lemma 15, see the arguments between (4.3) and (4.5), we have that $I \leq \gamma/2$ whenever

$$n \gtrsim \frac{K^2 \log(2/\gamma)}{\alpha(t)^2} \bigvee K^2 \frac{\log(1/\gamma) \vee \text{VC}(\mathcal{F}) \log(\frac{K}{c\alpha(t)}) \vee 1}{\alpha(t)^2}. \quad (\text{B.12})$$

Furthermore, following (4.6), $II \leq \gamma/2$ when

$$n > 4K \frac{\log(2/\gamma) \vee \psi(n\epsilon/2K)}{\epsilon \alpha(t)}, \quad (\text{B.13})$$

from which it then remains to simplify the bound

$$\psi(n\epsilon/2K) \leq n\epsilon\alpha(t)/4K. \quad (\text{B.14})$$

Given that by assumption D is L -Lipschitz, we have that $\omega(r) \leq Lr$. Plugging this in to the definition of ψ from (3.1), we have that

$$\psi(n\epsilon/2K) \leq \min_{r>0} \left\{ \frac{n\epsilon}{2K} \cdot L \cdot r - \log \pi(B_r(E_{\phi,\mu})) \right\}.$$

By taking the specific choice $r = \alpha(t)/4L$, we see that a sufficient condition for (B.14) to hold is

$$\frac{n\epsilon\alpha(t)}{8K} - \log \pi(B_{\alpha(t)/4L}(E_{D,\mu})) \leq \frac{n\epsilon\alpha(t)}{4K}, \quad (\text{B.15})$$

which is equivalent to

$$n > \frac{-4K \log \pi(B_{\alpha(t)/4L}(E_{D,\mu}))}{\epsilon\alpha(t)}. \quad (\text{B.16})$$

Combining (B.12), (B.13) and (B.16) yields the desired result. \blacksquare

Appendix C. On the smoothing parameter in the smoothed integrated dual depth

We executed a small simulation study to assess the effect of the smoothing parameter s on the non-private estimation error of the smoothed integrated dual depth median, as well as the effect of s on the convergence of the gradient descent algorithm. The simulation set up was the same as described in Section 2.5. Figure 3 shows the empirical root-mean-square error for different values of s and the dimension d . We see that the choice of s does not particularly depend on the dimension of the data. Furthermore, we see that choosing $s \geq 10$ produces the best results in terms balancing convergence and estimation error. Given that larger s theoretically gives a closer approximation to the integrated dual depth, we recommend choosing $s = 100$ in practice.

Appendix D. Extra plots from simulations




Figure 3: ERMSE of the non-private smoothed integrated dual depth median for different values of s under Gaussian data (left) and contaminated Gaussian data (center) data. Convergence of one simulation run of the gradient descent algorithm in dimension 2 is also presented (right). It can be seen that choosing s large gives the lowest error and best convergence, though any value $s \geq 10$ produces similar results.

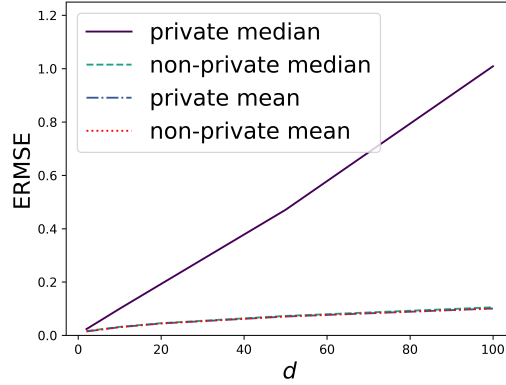


Figure 4: Empirical root mean squared error (ERMSE) of the location estimators under Gaussian data. In the uncontaminated setting, the ERMSE of the private median grows at a slightly faster rate than its non-private counterpart and the mean estimators. This is due to the cost of privacy.




Figure 5: Empirical root mean squared error (ERMSE) of the location estimators under (left) Gaussian data and (right) contaminated Gaussian data when $\epsilon = 5$.

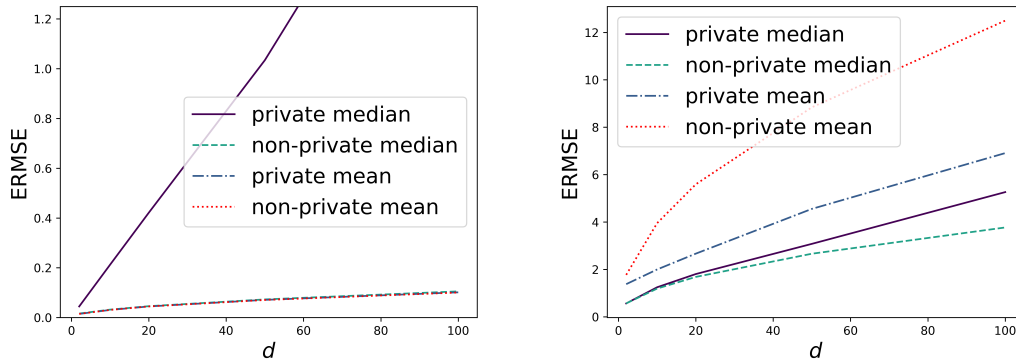


Figure 6: Empirical root mean squared error (ERMSE) of the location estimators under (left) Gaussian data and (right) contaminated Gaussian data when $\epsilon = 2$.

Appendix E. Auxiliary lemmas

Lemma 36 *Suppose U_n is an upper bound on $\text{GS}_n(\phi)$. Setting $\beta = \epsilon/2U_n$ implies $\tilde{\theta}_n$ is ϵ -differentially private.*

Proof From the definition of differential privacy, it is immediate that for any $0 < \delta_1 < \delta_2$, we have that δ_1 -differential privacy implies δ_2 -differential privacy. Now, it remains to show that $\tilde{\theta}_n$ is δ -differentially private for some $\delta < \epsilon$. To see this, note that $\epsilon/2U_n \leq \epsilon/2\text{GS}_n(\phi)$, which implies that there exists $\delta < \epsilon$ such that $\epsilon/2U_n = \delta/2\text{GS}_n(\phi)$. This implies that $\tilde{\theta}_n$ is δ -differentially private and the result follows. \blacksquare

Lemma 37 *If ϕ satisfies Condition 3 then $\text{GS}_n(\phi) \leq K/n$.*

Proof Definition 4 gives that

$$\begin{aligned} \sup_{\substack{\mathbb{X}_n \in \mathbf{D}_{n \times d}, \\ \tilde{\mu}_n \in \mathcal{M}(\mathbb{X}_n)}} \sup_x |\phi(x, \hat{\mu}_n) - \phi(x, \tilde{\mu}_n)| &\leq K \sup_{\substack{\mathbb{X}_n \in \mathbf{D}_{n \times d}, \\ \tilde{\mu}_n \in \mathcal{M}(\mathbb{X}_n)}} \sup_{g \in \mathcal{F}} \left| \int g d(\hat{\mu}_n - \tilde{\mu}_n) \right| \\ &\leq K \sup_{g \in \mathcal{F}} \|g\| / n = K/n. \end{aligned}$$

\blacksquare

Lemma 38 *Let D be the projection depth with location measure MED and scale measure MAD, i.e.,*

$$D(x, \mu) = \frac{1}{1 + O(x, \mu)}, \quad \text{with} \quad O(x, \mu) = \sup_{u \in \mathbb{S}^{d-1}} |x^\top u - \text{MED}(\mu_u)| / \text{MAD}(\mu_u).$$

If $\epsilon = O(1)$ and π is absolutely continuous, and $\pi \not\propto \mathbb{1}\{x \in E_{D, \mu}\} dx$, then $\text{GS}_n(D) = 1$ and for a median drawn from the exponential mechanism based on D with $\beta = \epsilon/2\text{GS}_n(D)$, $\tilde{\theta}_n$, there exists $t > 0$ such that

$$\lim_{n \rightarrow \infty} \Pr \left(d(E_{\mu, \phi}, \tilde{\theta}_n) \geq t \right) \neq 0.$$

Proof The univariate sample median has infinite global sensitivity (Brunel and Avella-Medina, 2020), which implies that the projection depth with location measure MED and scale measure MAD has sensitivity 1. It follows that $\beta = \epsilon/2$. We first show that there exists $t > 0$ such that $Q_{\mu, \epsilon/2}(d(E_{D, \mu}, X) \geq t) > 0$. Since π is absolutely continuous and $\pi \not\propto \mathbb{1}\{x \in E_{D, \mu}\} dx$, there exists A and $t > 0$ such that $\inf_{x \in A} d(x, E_{D, \mu}) \geq t > 0$ and $\pi(A) > 0$. Since $D(x, \mu) \geq 0$, $Q_{\mu, \epsilon/2}(A) > 0$, so $Q_{\mu, \epsilon/2}(d(E_{\mu, \phi}, X) \geq t) > 0$ as well.

Next, we show that $\lim_{n \rightarrow \infty} \Pr \left(d(E_{D, \mu}, \tilde{\theta}_n) \geq t \right) = Q_{\mu, \epsilon/2}(d(E_{D, \mu}, X) \geq t)$. To this end, we have that Zuo (2003) gives that $D(x, \hat{\mu}_n) \rightarrow D(x, \mu)$ μ -a.s. as $n \rightarrow \infty$. Thus,

$dQ_{\hat{\mu}_n, \epsilon/2}/d\pi \rightarrow dQ_{\mu, \epsilon/2}/d\pi$ μ -a.s. as $n \rightarrow \infty$. This implies that, μ -a.s., $Q_{\hat{\mu}_n, \epsilon/2} \rightarrow Q_{\mu, \epsilon/2}$ in total variation as $n \rightarrow \infty$. We have that this implies

$$\begin{aligned} \lim_{n \rightarrow \infty} \Pr \left(d(E_{\mu, \phi}, \tilde{\theta}_n) \geq t \right) &= \lim_{n \rightarrow \infty} \mathbb{E}(Q_{\hat{\mu}_n, \epsilon/2}(d(E_{\mu, \phi}, X) \geq t)) \\ &= \mathbb{E}(\lim_{n \rightarrow \infty} Q_{\hat{\mu}_n, \epsilon/2}(d(E_{\mu, \phi}, X) \geq t)) \\ &= Q_{\mu, \epsilon/2}(d(E_{\mu, \phi}, X) \geq t) \neq 0, \end{aligned}$$

which completes the proof. ■

Lemma 39 *Suppose that D is non-negative and satisfies properties (2) and (4) in Definition 2. Then there exists a universal constant $C > 0$ such that $\text{GS}_n(\hat{\mathbf{D}}(\hat{\mu}_n)) \geq C$. Furthermore, if D satisfies Condition 3, then $\text{GS}_n(\hat{\mathbf{D}}(\hat{\mu}_n)) = O(1)$.*

Proof Consider $\tilde{\mu}_n$ to be centrally symmetric about X_1 . Property (2) and non-negativity imply that $D(X_1, \tilde{\mu}_n) = C > 0$. Let $\tilde{\mu}_n(Y_1)$ be the empirical measure corresponding to Y_1, \dots, X_n with $Y_1 \neq X_1$. Given that $\tilde{\mu}_n(Y_1) \in \widetilde{\mathcal{M}}(\mathbb{X}_n)$, property (4) implies that

$$\text{GS}_n(\hat{\mathbf{D}}(\hat{\mu}_n)) \geq \sup_{Y_1 \in \mathbb{R}^d} |D(Y_1, \tilde{\mu}_n(Y_1)) - D(X_1, \hat{\mu}_n)| \geq D(X_1, \hat{\mu}_n) - \inf_{Y_1 \in \mathbb{R}^d} D(Y_1, \tilde{\mu}_n(Y_1)) \geq C.$$

Now, if D satisfies Condition 3, then Lemma 37 implies that

$$\text{GS}_n(\hat{\mathbf{D}}(\hat{\mu}_n))^2 \leq \frac{K^2}{n} + \sup_{\hat{\mu}_n, y} |D(y, \hat{\mu}_n) - D(X_1, \hat{\mu}_n)|^2 \leq \frac{K^2}{n} + \|D\| = O(1). \quad \blacksquare$$

Lemma 40 *If $\mu = \mathcal{N}(\theta_0, \Sigma)$ then for all $t > 0$ and positive definite Σ , $\sup_{\|x\|=t} \text{HD}(x, \mu) = \Phi(-t/\sqrt{\lambda_1})$.*

Proof From the fact that HD is affine invariant, without loss of generality we can set $\theta_0 = 0$. From the fact that Φ is increasing, it holds that

$$\begin{aligned} \sup_{\|x\|=t} \text{HD}(x, \mu) &= \sup_{\|x\|=t} \inf_{\|u\|=1} \Phi\left(\frac{x^\top u}{\sqrt{u^\top \Sigma u}}\right) \\ &= 1 - \inf_{\|v\|=1} \sup_{\|u\|=1} \Phi\left(\frac{tv^\top u}{\sqrt{u^\top \Sigma u}}\right) = 1 - \Phi\left(t \inf_{\|v\|=1} \sup_{\|u\|=1} \frac{v^\top u}{\sqrt{u^\top \Sigma u}}\right). \end{aligned}$$

It suffices to compute

$$\inf_{\|v\|=1} \sup_{\|u\|=1} \frac{(v^\top u)^2}{u^\top \Sigma u}.$$

First, note that

$$\inf_{\|v\|=1} \sup_{\|u\|=1} \frac{(v^\top u)^2}{u^\top \Sigma u} \geq \inf_{\|v\|=1} \frac{(v^\top v)^2}{v^\top \Sigma v} = 1/\lambda_1.$$

Showing the reverse inequality, setting $v = e_1$ and setting $z = \Sigma^{1/2}u$ gives that

$$\inf_{\|v\|=1} \sup_{\|u\|=1} \frac{(v^\top u)^2}{u^\top \Sigma u} \leq \sup_{\|u\|=1} \frac{(e_1^\top u)^2}{u^\top \Sigma u} = \sup_{\|\Sigma^{-1/2}z\|=1} \frac{(e_1^\top \Sigma^{-1/2}z)^2}{z^\top z} = \sup_{\|\Sigma^{-1/2}z\|=1} (e_1^\top z / \|z\|)^2 / \lambda_1 \leq \frac{1}{\lambda_1}.$$

Thus,

$$\sup_{\|x\|=t} \text{HD}(x, \mu) = 1 - \Phi(t/\sqrt{\lambda_1}) \geq \Phi(-t/\sqrt{\lambda_1}).$$

■

Appendix F. Proofs of the results from Section 6

We now prove the various results presented in Section 6.

Proof [Proof of Proposition 27] With the exception of the modified spatial depth, these results were shown by (Zuo and Serfling, 2000; Liu, 1988; Serfling, 2002; Cuevas and Fraiman, 2009; Ramsay et al., 2019), respectively. For the modified spatial depth, properties (1), (2) and (4) follow from the fact that properties (1), (2) and (4) for spatial depth. ■

Proof [Proof of Proposition 32] We prove properties (1)-(4) in order. Let $A \in \mathbb{R}^{d \times d}$ be any orthogonal matrix and $b \in \mathbb{R}^d$. For any integrable function $h: \mathbb{S}^{d-1} \rightarrow \mathbb{R}^+$, it holds that

$$\int_{\mathbb{S}^{d-1}} h(u) d\nu = \int_{\mathbb{S}^{d-1}} h(Au) d\nu. \quad (\text{F.1})$$

It also holds that

$$(Ax + b - AX - b)^\top u = A^\top (x - X)^\top u = A^\top u^\top (x - X).$$

Let $Y \sim A\mu + b$. Applying the above identity implies that

$$\mathbb{E}\sigma\left(s(Ax + b - Y)^\top u\right) = \mathbb{E}\sigma\left(sA^\top u(x - X)\right).$$

The result follows from the definition of the smoothed integrated dual depth and (F.1).

For property (2), by definition $\sup_x \text{IDD}(x, \mu, s) \leq 1/4$. It remains to show $\text{IDD}(\theta_0, \mu, s) = 1/4$. We will use the fact that

$$(1 - \sigma(x - X)) = \sigma(X - x). \quad (\text{F.2})$$

Now, consider the integrand in the definition of smoothed integrated dual depth at any fixed u . Central symmetry of μ and (F.2) yield that

$$\mathbb{E}\sigma\left(s(\theta_0 - X)^\top u\right) \left(1 - \mathbb{E}\sigma\left(s(\theta_0 - X)^\top u\right)\right) = \left(\mathbb{E}\sigma\left(s(X - \theta_0)^\top u\right)\right)^2.$$

Central symmetry of μ and (F.2) also yield that

$$\mathbb{E}\sigma\left(s(\theta_0 - X)^\top u\right) \left(1 - \mathbb{E}\sigma\left(s(\theta_0 - X)^\top u\right)\right) = \left(1 - \mathbb{E}\sigma\left(s(X - \theta_0)^\top u\right)\right)^2.$$

These two equalities yield that

$$\mathbb{E}\sigma\left(s(X - \theta_0)^\top u\right) = 1/2.$$

Given that this holds for any $u \in \mathbb{S}^{d-1}$, it follows that $\text{IDD}(\theta_0; \mu, s) = 1/4 = \sup_x \text{IDD}(x; \mu, s)$.

For property (3), at any fixed u , central symmetry of μ and the fact that σ is increasing yields that either

$$\mathbb{E}\sigma\left(s(x - X)^\top u\right) \leq \mathbb{E}\sigma\left(s(p\theta_0 + (1-p)x - X)^\top u\right) \leq 1/2,$$

holds or

$$1/2 \leq \mathbb{E}\sigma\left(s(p\theta_0 + (1-p)x - X)^\top u\right) \leq \mathbb{E}\sigma\left(s(x - X)^\top u\right),$$

holds. The result follows immediately from the fact that $f(x) = x(1-x)$ is monotone on both $[0, 1/2]$ and $[1/2, 1]$.

Lastly, we show property (4). For any fixed u , the fact that $\lim_{x \rightarrow \infty} \sigma(x)\sigma(-x) = 0$ and the bounded convergence theorem imply that

$$\begin{aligned} \lim_{c \rightarrow \infty} \text{IDD}(cu; \mu, s) &= \lim_{c \rightarrow \infty} \int_{\mathbb{S}^{d-1}} \left(\mathbb{E}\sigma\left(s(cu - X)^\top u\right) \right) \left(1 - \mathbb{E}\sigma\left(s(cu - X)^\top u\right) \right) d\nu(u) \\ &= \int_{\mathbb{S}^{d-1}} \lim_{c \rightarrow \infty} \left(\mathbb{E}\sigma\left(s(cu - X)^\top u\right) \right) \left(1 - \mathbb{E}\sigma\left(s(cu - X)^\top u\right) \right) d\nu(u) \\ &= 0. \end{aligned}$$

■

Proof [Proof of Proposition 33] Let

$$Z = \sup_x \left| \widehat{\text{IDD}}(x, \hat{\mu}_n, s) - \text{IDD}(x, \hat{\mu}_n, s) \right|.$$

We first prove a concentration result for Z . To this end, let

$$\begin{aligned} G_n(x, u, s) &= \frac{1}{n} \sum_{i=1}^n \sigma\left(s(x - X_i)^\top u\right) \\ W_1 &= \sup_x \left| \frac{1}{M} \sum_{m=1}^M G_n(x, u_m, s) - \int_{\mathbb{S}^{d-1}} G_n(x, u, s) d\nu(u) \right| \\ W_2 &= \sup_x \left| \frac{1}{M} \sum_{m=1}^M G_n(x, u_m, s)^2 - \int_{\mathbb{S}^{d-1}} G_n(x, u, s)^2 d\nu(u) \right|. \end{aligned}$$

From the fact that $Z \leq W_1 + W_2$, we need only prove concentration results for W_1 and W_2 . Let

$$\mathcal{G} = \left\{ \sigma(y^\top u + b)/n : y \in \mathbb{R}^d, b \in \mathbb{R} \right\}.$$

Note that $\text{VC}(\mathcal{G}) \leq d+2$. Therefore, it follows from Theorem 7.12 in (Sen, 2022) that there exists a universal constant $K > 0$ such that for all $\epsilon \in (0, 1)$, it holds that

$$\sup_Q N(\epsilon/n, \mathcal{G}, L_2(Q)) \leq \left(\frac{K}{\epsilon}\right)^{2d+4}.$$

Now, define

$$\sum_{i=1}^n \mathcal{G}_i = \left\{ \frac{1}{n} \sum_{i=1}^n \sigma(y_i^\top u + b_i) : y_1, \dots, y_n \in \mathbb{R}^d, b_1, \dots, b_n \in \mathbb{R} \right\}.$$

It follows from Lemma 7.21 in (Sen, 2022) that for all $\epsilon \in (0, 1)$, it holds that

$$\sup_Q N\left(2\epsilon, \sum_{i=1}^n \mathcal{G}_i, L_2(Q)\right) \leq \left(\frac{K}{\epsilon}\right)^{2dn+4n}.$$

It follows from (Talagrand, 1994; Sen, 2022), that there exists a universal constant $K' > 0$ such that for all $t > 0$, it holds that

$$\Pr\left(\sqrt{M}W_1 > t\right) \leq \left(\frac{K't}{\sqrt{2dn+4n}}\right)^{2dn+4n} e^{-2t^2}.$$

However, note that $|W_1| \leq 1/4$, therefore, for $t > \sqrt{M}/4$, it holds that $\Pr\left(\sqrt{M}W_1 > t\right) = 0$. Therefore, it holds that

$$\left(\frac{K't}{\sqrt{2dn+4n}}\right)^{2dn+4n} \leq \left(\frac{K'\sqrt{M}}{4\sqrt{2dn+4n}}\right)^{2dn+4n},$$

with probability 1, and so,

$$\Pr(W_1 > t) \leq \left(\frac{K'M}{4dn}\right)^{dn} e^{-2Mt^2}.$$

The same logic applied to W_2 gives that there exists a universal constant $K'' > 0$ such that for all $t > 0$, it holds that

$$\Pr(W_2 > t) \leq \left(\frac{K''M}{dn}\right)^{dn} e^{-2Mt^2}.$$

Therefore, there exists universal constants $K, K' > 0$ such that for all $t > 0$, it holds that

$$\Pr(Z > t) \leq \left(\frac{KM}{dn}\right)^{dn} e^{-K'Mt^2}.$$

Now, to prove the sample complexity bound, we want to find M such that

$$\left(\frac{KM}{dn}\right)^{dn} e^{-K'Mt^2} \leq \gamma,$$

which is equivalent to the inequality

$$-dn \log(M/nd) + K' M t^2 \geq \log(1/\gamma) + dn \log(K). \quad (\text{F.3})$$

First, we show that $dn \log(M/nd) \leq K' M t^2/2$. Note that the function $h(M) = K' M t^2/2 - dn \log(M/nd)$ satisfies the conditions of Lemma 14. Applying Lemma 14 yields that $h(M) \geq 0$ for

$$M \geq 2 \frac{2dn}{K' t^2} \log\left(\frac{2}{K' t^2} \vee e\right) \gtrsim \frac{dn}{t^2} \log\left(\frac{1}{t} \vee e\right). \quad (\text{F.4})$$

When (F.4) is satisfied, we have that (F.3) is satisfied for

$$M \geq 2 \frac{\log(1/\gamma) + dn \log K}{K' t^2}. \quad (\text{F.5})$$

Thus, for all $t > 0$ and all $0 < \gamma < 1$, there exists a universal constant $c_1 > 0$ such that

$$\sup_x \left| \widehat{\text{IDD}}(x, \hat{\mu}_n, s) - \text{IDD}(x, \hat{\mu}_n, s) \right| \leq t,$$

with probability $1 - \gamma$, provided that

$$M \geq c_1 \frac{\log(1/\gamma) \vee dn \log\left(\frac{1}{t} \vee e\right)}{t^2}.$$

■

We prove Theorem 30 with a series of Lemmas.

Lemma 41 *Condition 1 holds for each of the depth functions defined in Section 6.*

Proof Recall that all of the depth functions satisfy property (4) given in Definition 2. Property (4) implies that there exists a compact set E such that $\sup_{x \in \mathbb{R}^d} D(x, \mu) = \sup_{x \in E} D(x, \mu)$. It follows from boundedness of $D(x, \mu)$ and compactness of E that there exists a point in $y \in E$ such that $D(y, \mu) = \sup_{x \in E} D(x, \mu) = \sup_{x \in \mathbb{R}^d} D(x, \mu)$. ■

Lemma 42 *Suppose that π is absolutely continuous and that for all $u \in \mathbb{S}^{d-1}$, the measure μ_u is absolutely continuous with density f_u and $\sup_{\|u\|=1} \|f_u\|_\infty = L < \infty$. Then there exists a universal constant $C > 0$ such that for all $s \in (0, \infty]$, halfspace, simplicial, integrated rank-weighted and smoothed integrated dual depth satisfy Condition 4 with $\omega(r) = C \cdot L \cdot r$.*

Proof We prove the result for each depth in turn.

Halfspace depth: Consider two points $x, y \in \mathbb{R}^d$. By assumption, $F(\cdot, u, \mu)$ is L -Lipschitz function a.e. . It follows that

$$\begin{aligned} |\text{HD}(x, \mu) - \text{HD}(y, \mu)| &= \left| \inf_u F(x, u, \mu) - \inf_u F(y, u, \mu) \right| \\ &\leq \sup_u |F(x, u, \mu) - F(y, u, \mu)| \\ &\leq L \cdot \|x - y\| \quad \text{a.e.} \end{aligned}$$

Integrated rank-weighted depth: For integrated rank-weighted depth, note that absolute continuity of μ_u implies that $F(x, u, \mu) = F(x-, u, \mu)$, where one recalls that $F(x-, u, \mu) = \mu(\{X^\top u < x^\top u\})$. With this in mind, for any two points $x, y \in \mathbb{R}^d$ the reverse triangle inequality yields that

$$|\text{IRW}(x, \mu) - \text{IRW}(y, \mu)| \leq \int_{\mathbb{S}^{d-1}} 2|F(x, u, \mu) - F(y, u, \mu)| d\nu(u) \leq 2L \cdot \|x - y\| \text{ a.e. } .$$

Integrated dual depth: For any $h : S \rightarrow [0, 1]$, $h' : S \rightarrow [0, 1]$, for any $x \in S$ it holds that

$$|h(x)(1 - h(x)) - h'(x)(1 - h'(x))| \leq 3|h(x) - h'(x)|. \quad (\text{F.6})$$

Using this inequality,

$$|\text{IDD}(x, \mu, s) - \text{IDD}(y, \mu, s)| \leq 3 \sup_u \left| \mathbb{E}_\mu \sigma(s(x - X)^\top u) - \mathbb{E}_\mu \sigma(s(y - X)^\top u) \right|. \quad (\text{F.7})$$

It remains to bound $|\mathbb{E}_\mu \sigma(s(x - X)^\top u) - \mathbb{E}_\mu \sigma(s(y - X)^\top u)|$. To this end, first assume that $s \leq \infty$ and let $w, z \in \mathbb{R}$. Using the fact that σ is a 1/4-Lipschitz function, it holds that

$$\begin{aligned} |\mathbb{E}_{\mu_u} \sigma(s(w - X)) - \mathbb{E}_{\mu_u} \sigma(s(z - X))| &= \left| \int \sigma(s(w - t)) - \sigma(s(z - t)) f_u(t) dt \right| \\ &= \left| \frac{1}{s} \int \sigma(v) - \sigma(sz - sw + v) f_u(-v/s - w) dv \right| \\ &\leq |z - w|/4 \int |f_u(-v/s - w)| dv \\ &\leq |z - w| \|f_u\| / 4, \end{aligned}$$

where in the second line we use the substitution $v = s(w - t)$. The preceding inequality and (F.7) imply that

$$\begin{aligned} |\text{IDD}(x, \mu, s) - \text{IDD}(y, \mu, s)| &\leq 3 \sup_u \left| \mathbb{E}_\mu \sigma(s(x - X)^\top u) - \mathbb{E}_\mu \sigma(s(y - X)^\top u) \right| \\ &\leq \frac{3}{4} L \|x - y\|. \end{aligned}$$

Now assume that $s = \infty$. It follows from absolute continuity of μ_u and (F.7) that

$$|\text{IDD}(x, \mu) - \text{IDD}(y, \mu)| \leq 3 \sup_u |F(x, u, \mu) - F(y, u, \mu)| \leq 3L \|x - y\| \text{ a.e.}$$

Simplicial depth: For simplicial depth, we must show that $\Pr(x \in \Delta(X_1, \dots, X_{d+1}))$ is Lipschitz continuous. It is easy to begin with two dimensions. Consider $\Pr(x \in \Delta(X_1, X_2, X_3)) - \Pr(y \in \Delta(X_1, X_2, X_3))$, as per (Liu, 1990), we need to show that $\Pr(\overline{X_1 X_2} \text{ intersects } \overline{xy}) \leq L \cdot \|x - y\|$. In order for this event to occur, we must have that X_1 is above \overline{xy} and X_2 is below \overline{xy} , but both are projected onto the line segment \overline{xy} when projected onto the line running through \overline{xy} . Affine invariance of simplicial depth implies we can assume, without

loss of generality, that x and y lie on the axis of the first coordinate. Let x_1 and y_1 be the first coordinates of x and y . Suppose that X_{11} is the first coordinate of X_1 . It then follows from a.e. Lipschitz continuity of $F(x, u, \mu)$ that

$$\Pr(\overline{X_1 X_2} \text{ intersects } \overline{xy}) \leq \Pr(x_1 < X_{11} < y_1) \leq L \cdot |x_1 - y_1| \leq L \|x - y\| \text{ a.e. .}$$

In dimensions greater than two, a similar line of reasoning can be used. We can again assume, without loss of generality, that x and y lie on the axis of the first coordinate. It holds that

$$\Pr(x \in \Delta(X_1, X_2, X_3)) - \Pr(y \in \Delta(X_1, X_2, X_3)) \leq \binom{d+1}{d} \Pr(A_d),$$

where A_d is the event that the $d - 1$ -dimensional face of the random simplex, formed by d points randomly drawn from F , intersects the line segment \overline{xy} . It is easy to see that

$$\Pr(A_d) \leq \Pr(x_1 < X_{11} < y_1) \leq \|f_u\|_\infty |x_1 - y_1| \leq L \cdot \|x - y\| \text{ a.e. .}$$

This completes the proof. ■

Lemma 43 *Suppose that μ has a bounded density and that $\sup_y \mathbb{E} \|y - X\|^{-1} = L' \leq \infty$. Then the modified spatial depth satisfies Condition 4 with $\omega(r) = 2\sqrt{d}L' \cdot r$ and spatial depth satisfies Condition 4 with $\omega(r) = 2L' \cdot r$.*

Proof Taking the derivative of MSD and using the assumption of a bounded density results in

$$\frac{d \text{MSD}(y, \mu)}{dy} = -2\mathbb{E}_\mu(\text{SR}(y - X)) \mathbb{E}_\mu \left[\frac{1_d}{\|y - X\|} - \frac{(y - X) * (y - X)^\top 1_d}{\|y - X\|^3} \right].$$

Now,

$$\left\| \frac{d \text{MSD}(y, \mu)}{dy} \right\| \leq 2\sqrt{d} \sup_y \mathbb{E} \|y - X\|^{-1} = 2\sqrt{d}L'.$$

Thus, the map $y \mapsto \text{MSD}(y, \mu)$ is $2\sqrt{d}L'$ -Lipschitz.

Define the spatial rank function $h(y, \mu) = \mathbb{E}_\mu(\text{SR}(y - X))$. Proposition 4.1 of (Koltchinskii, 1997) says that if μ has a bounded density, then the map $y \mapsto h(y, \mu)$ is continuously differentiable in \mathbb{R}^d with derivative

$$\frac{dh}{dy} = \int_{\{x \neq y\}} \frac{1}{\|y - x\|} \left[I_d - \frac{(y - x)(y - x)^\top}{\|y - x\|^2} \right] d\mu(x).$$

Now,

$$\left\| \frac{dh}{dy} \right\| \leq 2 \sup_y \mathbb{E} \|y - X\|^{-1} = 2L',$$

where $\|\cdot\|$ denotes the operator norm. Thus, the map $y \mapsto h(y, \mu)$ is $2L'$ -Lipschitz, which implies that the map $y \mapsto \text{SD}(y, \mu)$ is also $2L'$ -Lipschitz. ■

Lemma 44 *Condition 3 holds for each of the depth functions defined in Section 6.*

Proof We prove the result for each depth in turn.

Halfspace depth: Let

$$\mathcal{F} = \left\{ \mathbb{1} \left\{ X^\top u \leq y \right\} : u \in \mathbb{S}^{d-1}, y \in \mathbb{R} \right\}$$

and recall that the set of subgraphs of \mathcal{F} is the set of closed halfspaces in \mathbb{R}^d . Furthermore, it holds that

$$\sup_x |\text{HD}(x, \mu_1) - \text{HD}(x, \mu_2)| \leq \sup_{g \in \mathcal{F}} |\mathbb{E}_{\mu_1} g(X) - \mathbb{E}_{\mu_2} g(X)|.$$

Thus, it follows HD is $(1, \mathcal{F})$ -regular with $\text{VC}(\mathcal{F}) = d + 2$.

Integrated rank-weighted-depth: First observe that

$$\begin{aligned} \sup_x |\text{IRW}(x, \mu_1) - \text{IRW}(x, \mu_2)| &\leq 4 \sup_{u, x} [|F(x, u, \mu_1) - F(x, u, \mu_2)| \\ &\quad \vee |F(x-, u, \mu_1) - F(x-, u, \mu_2)|]. \end{aligned}$$

Define

$$\mathcal{G} = \left\{ g : g(X) = \mathbb{1} \left\{ X^\top u < y \right\}, u \in \mathbb{S}^{d-1}, y \in \mathbb{R} \right\}.$$

We see that the integrated rank-weighted depth function is $(4, \mathcal{G} \vee \mathcal{F})$ -regular. Furthermore, since the subgraphs of \mathcal{G} are contained in those of \mathcal{F} , we have that $\text{VC}(\mathcal{G} \vee \mathcal{F}) = d + 2$.

Integrated dual depth: If $s = \infty$ then it follows immediately from (F.6) and the analysis of halfspace depth that the integrated dual depth is $(3, \mathcal{F})$ -regular with $\text{VC}(\mathcal{F}) = d + 2$. Suppose that $s < \infty$. It follows from (F.6) that

$$\sup_x |\text{IDD}(x, \mu_2) - \text{IDD}(x, \mu_1)| \leq 3 \sup_{y \in \mathbb{R}, u \in \mathbb{S}^{d-1}} |\mathbb{E}_{\mu_1} \sigma(s(x - X)^\top u) - \mathbb{E}_{\mu_2} \sigma(s(x - X)^\top u)|.$$

We have that

$$\sup_{y \in \mathbb{R}, u \in \mathbb{S}^{d-1}} |\mathbb{E}_{\mu_1} \sigma(s(x - X)^\top u) - \mathbb{E}_{\mu_2} \sigma(s(x - X)^\top u)| \leq \sup_{g \in \mathcal{F}'} |\mathbb{E}_{\mu_1} g(X) - \mathbb{E}_{\mu_2} g(X)|,$$

where

$$\mathcal{F}' = \left\{ \sigma(AX_i + b) : A \in \mathbb{R}^d, b \in \mathbb{R} \right\}.$$

The class of functions \mathcal{F}' is constructed from a monotone function applied to a finite-dimensional vector space of measurable functions. It follows from Lemma 7.15 and Lemma 7.19 of Sen (2022) that $\text{VC}(\mathcal{F}') = d + 2$. Therefore, the smoothed integrated dual depth is $(3, \mathcal{F}')$ -regular with $\text{VC}(\mathcal{F}') = d + 2$.

Simplicial Depth: Dümbgen (1992) gives that

$$\sup_x |\text{SMD}(x, \mu_1) - \text{SMD}(x, \mu_2)| \leq (d+1) \sup_{A \in \mathcal{A}} |\mu_1(A) - \mu_2(A)|,$$

where \mathcal{A} is the set of intersections of d open half-spaces in \mathbb{R}^d . Define the set of Boolean functions $\mathcal{A}' = \{g(x) = \mathbb{1}\{x \in A\} : A \in \mathcal{A}\}$. Thus, simplicial depth is $(d+1, \mathcal{A}')$ -regular with $\text{VC}(\mathcal{A}') = O(d^2 \log d)$.

Spatial Depth: Applying the reverse triangle inequality yields

$$\begin{aligned} \sup_x \left| \left\| \mathbb{E}_{\mu_1} \frac{x - X}{\|x - X\|} \right\| - \left\| \mathbb{E}_{\mu_2} \frac{x - X}{\|x - X\|} \right\| \right| &\leq \sup_x \left\| \mathbb{E}_{\mu_1} \frac{x - X}{\|x - X\|} - \mathbb{E}_{\mu_2} \frac{x - X}{\|x - X\|} \right\| \\ &\leq \sup_x \sum_{j=1}^d \left| \mathbb{E}_{\mu_1} \frac{x_j - X_j}{\|x - X\|} - \mathbb{E}_{\mu_2} \frac{x_j - X_j}{\|x - X\|} \right| \\ &\leq d \sup_{x, 1 \leq j \leq d} \left| \mathbb{E}_{\mu_1} \frac{x_j - X_j}{\|x - X\|} - \mathbb{E}_{\mu_2} \frac{x_j - X_j}{\|x - X\|} \right|. \end{aligned}$$

It remains to compute the VC-dimension of

$$\mathcal{G} = \left\{ g(X) = \frac{x_j - X_j}{\|x - X\|} : x \in \mathbb{R}^d, j \in [d] \right\}.$$

Now, let E be the set of standard basis vectors for \mathbb{R}^d . We can then write

$$\mathcal{G} = \left\{ g(X) = \frac{e^\top (x - X)}{\|x - X\|} : x \in \mathbb{R}^d, e \in E \right\}.$$

It is now easy to see that

$$\mathcal{G} \subset \left\{ g(X) = y^\top (x - X) : x \in \mathbb{R}^d, y \in \mathbb{R}^d, \|y\| = 1 \right\}.$$

It follows that $\text{VC}(\mathcal{G}) \leq O(d)$ and that spatial depth is (d, \mathcal{G}) regular.

Modified Spatial Depth: We may write

$$\begin{aligned} \sup_x |\text{SD}(X, \mu_1) - \text{SD}(X, \mu_2)| &= \sup_x \left| \left\| \mathbb{E}_{\mu_1} \left(\frac{x - X}{\|x - X\|} \right) \right\|^2 - \left\| \mathbb{E}_{\mu_2} \left(\frac{x - X}{\|x - X\|} \right) \right\|^2 \right| \\ &= \sup_x \left| \sup_u \left[\mathbb{E}_{\mu_1} \left(\frac{x - X}{\|x - X\|} \right)^\top u \right] - \sup_u \left[\mathbb{E}_{\mu_2} \left(\frac{x - X}{\|x - X\|} \right)^\top u \right] \right| \\ &\leq \sup_u \sup_x \left| \mathbb{E}_{\mu_1} \left(\frac{x - X}{\|x - X\|} \right)^\top u - \mathbb{E}_{\mu_2} \left(\frac{x - X}{\|x - X\|} \right)^\top u \right| \\ &\leq \sup_u \sup_x \left| \mathbb{E}_{\mu_1} \left(\frac{x - X}{\|x - X\|} \right)^\top u - \mathbb{E}_{\mu_2} \left(\frac{x - X}{\|x - X\|} \right)^\top u \right|. \end{aligned}$$

It remains to compute the VC-dimension of

$$\mathcal{G} = \left\{ g(X) = \frac{u^\top (x - X)}{\|x - X\|} : x \in \mathbb{R}^d, u \in \mathbb{S}^{d-1} \right\}.$$

It is now easy to see that

$$\mathcal{G} \subset \left\{ g(X) = y^\top X + b : b \in \mathbb{R}, y \in \mathbb{R}^d, j \in [d] \right\}.$$

It follows that $\text{VC}(\mathcal{G}) \leq d + 2$ and that the modified spatial depth is $(1, \mathcal{G})$ regular. \blacksquare

Proof [Proof of Theorem 30] The result follows from Lemmas 41-44. \blacksquare

Proof [Proof of Corollary 34] Together, Lemma 37 and the definition of the Laplace mechanism imply that $\tilde{D}(x, \hat{\mu}_n)$ is ϵ -differentially private. Furthermore, the properties of the Laplace measure give that

$$\begin{aligned} \Pr \left(\left| \tilde{D}(x, \hat{\mu}_n) - D(x, \mu) \right| > t \right) &\leq \Pr (|D(x, \hat{\mu}_n) - D(x, \mu)| > t/2) + \Pr \left(\left| W_1 \frac{K}{n\epsilon} \right| > t/2 \right) \\ &\leq \Pr (|D(x, \hat{\mu}_n) - D(x, \mu)| > t/2) + e^{-n\epsilon t/2K}. \end{aligned}$$

Applying (3.4) yields

$$\Pr \left(\left| \tilde{D}(x, \hat{\mu}_n) - D(x, \mu) \right| > t \right) \leq e^{\text{VC}(\mathcal{F}) \log \left(c_1 \frac{n}{\text{VC}(\mathcal{F})} \right) - c_2 n t^2} + e^{-n\epsilon t/2K}.$$

\blacksquare

References

- John Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajjhala, Brett Moran, William Sexton, Matthew Spence, and Pavel Zhuravlev. The 2020 census disclosure avoidance system TopDown algorithm. *Harvard Data Science Review*, (Special Issue 2), 6 2022. <https://hdrs.mitpress.mit.edu/pub/7evz361i>.
- Peyman Afshani, Donald R. Sheehy, and Yannik Stein. Approximating the simplicial depth in high dimensions. In *European Workshop on Computational Geometry*, 2016.
- Hilal Asi, Jonathan Ullman, and Lydia Zakyntinou. From robustness to privacy and back. In *International Conference on Machine Learning*, 2023.
- Marco Avella-Medina. The role of robust statistics in private data analysis. *CHANCE*, 33 (4):37–42, 2020. doi: 10.1080/09332480.2020.1847958.
- Marco Avella-Medina and Victor-Emmanuel Brunel. Differentially private sub-Gaussian location estimators. *arXiv e-prints*, 2019. arXiv:1906.11923.

- Jordan Awan and Aleksandra Slavković. Structure and sensitivity in differential privacy: Comparing k -norm mechanisms. *Journal of the American Statistical Association*, 116(534):935–954, 2021. doi: 10.1080/01621459.2020.1773831.
- Amos Beimel, , Shay Moran, Kobbi Nissim, and Uri Stemmer. Private center points and learning of halfspaces. In *Conference on Learning Theory*, volume 99, pages 269–282. PMLR, 6 2019.
- Omri Ben-Eliezer, Dan Mikulincer, and Ilias Zadik. Archimedes meets privacy: on privately estimating quantiles in high dimensions under minimal assumptions. In *Advances in Neural Information Processing Systems*, NIPS ’22, Red Hook, NY, USA, 2022. Curran Associates Inc. ISBN 9781713871088.
- Sourav Biswas, Yihe Dong, Gautam Kamath, and Jonathan Ullman. CoinPress: Practical private mean and covariance estimation. *Advances in Neural Information Processing Systems*, 33:14475–14485, 2020.
- Gavin Brown, Marco Gaboardi, Adam Smith, Jonathan Ullman, and Lydia Zakynthinou. Covariance-aware private mean estimation without private covariance estimation. *Advances in Neural Information Processing Systems*, 34:7950–7964, 2021.
- Victor-Emmanuel Brunel and Marco Avella-Medina. Propose, test, release: Differentially private estimation with high probability. *arXiv e-prints*, 2020. arXiv:2002.08774.
- Mark Bun and Thomas Steinke. Average-case averages: Private algorithms for smooth sensitivity and mean estimation. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mark Bun, Gautam Kamath, Thomas Steinke, and Steven Z Wu. Private hypothesis selection. *Advances in Neural Information Processing Systems*, 32, 2019.
- Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5): 2825 – 2850, 2021. doi: 10.1214/21-AOS2058.
- Kamalika Chaudhuri and Daniel Hsu. Convergence rates for differentially private statistical estimation. In *International Conference on Machine Learning.*, volume 2012, page 1327, 2012.
- Probal Chaudhuri. On a geometric notion of quantiles for multivariate data. *Journal of the American Statistical Association*, 91:862–872, 1996.
- Mengjie Chen, Chao Gao, and Zhao Ren. Robust covariance and scatter matrix estimation under huber’s contamination model. *The Annals of Statistics*, 46(5):1932–1960, 2018.
- Zhiqiang Chen and David E. Tyler. The influence function and maximum bias of Tukey’s median. *Annals of Statistics*, 30(6):1737–1759, 2002. doi: 10.1214/aos/1043351255.
- Shojaeddin Chenouri, Christopher G. Small, and Thomas J. Farrar. Data depth-based nonparametric scale tests. *Canadian Journal of Statistics*, 39(2):356–369, 2011. doi: 10.1002/cjs.10099.

- Antonio Cuevas and Ricardo Fraiman. On depth measures and dual statistics. a methodology for dealing with general data. *Journal of Multivariate Analysis*, 100(4):753–766, 2009.
- Differential Privacy Team. Learning with privacy at scale. *Apple*, 2017.
- David L Donoho. Breakdown properties of multivariate location estimators. Technical report, Harvard University, 1982.
- Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. *ACM symposium on theory of computing - STOC '09*, page 371, 2009.
- Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014. ISSN 1551-305X. doi: 10.1561/04000000042.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284, 2006.
- 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, 4 2017. doi: 10.1146/annurev-statistics-060116-054123.
- Rainer Dyckerhoff and Pavlo Mozharovskiy. Exact computation of the halfspace depth. *Computational Statistics & Data Analysis*, 98:19–30, 2016.
- Lutz Dümbgen. Limit theorems for the simplicial depth. *Statistics & Probability Letters*, 14(2):119 – 128, 1992. doi: [https://doi.org/10.1016/0167-7152\(92\)90075-G](https://doi.org/10.1016/0167-7152(92)90075-G).
- Rina Foygel Barber and John C. Duchi. Privacy and statistical risk: Formalisms and minimax bounds. *arXiv e-prints*, 2014. arXiv:1412.4451.
- Ricardo Fraiman and Graciela Muniz. Trimmed means for functional data. *Test*, 10(2): 419–440, 12 2001. ISSN 1133-0686. doi: 10.1007/BF02595706.
- Yue Gao and Or Sheffet. Private approximations of a convex hull in low dimensions. *arXiv e-prints*, 2020. arXiv:2007.08110.
- Kristian Georgiev and Samuel Hopkins. Privacy induces robustness: Information-computation gaps and sparse mean estimation. *Advances in neural information processing systems*, 35:6829–6842, 2022.
- Miguel Guevara. How we’re helping developers with differential privacy. Technical report, Google, 2021.
- Samuel B Hopkins, Gautam Kamath, and Mahbod Majid. Efficient mean estimation with pure differential privacy via a sum-of-squares exponential mechanism. In *ACM SIGACT Symposium on Theory of Computing*, pages 1406–1417, 2022.

- Samuel B Hopkins, Gautam Kamath, Mahbod Majid, and Shyam Narayanan. Robustness implies privacy in statistical estimation. In *ACM Symposium on Theory of Computing*, pages 497–506, 2023.
- Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan Ullman. Privately learning high-dimensional distributions. In *Conference on Learning Theory*, pages 1853–1902. PMLR, 2019.
- Gautam Kamath, Vikrant Singhal, and Jonathan Ullman. Private mean estimation of heavy-tailed distributions. In *Conference on Learning Theory*, pages 2204–2235. PMLR, 2020.
- Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. *arXiv e-prints*, 2017. arXiv:1711.03908.
- Vladimir Koltchinskii. M-estimation, convexity and quantiles. *The Annals of Statistics*, 25: 435–477, 1997.
- Michael R Kosorok. *Introduction to empirical processes and semiparametric inference*. Springer, 2008.
- Clément Sébastien Lalanne, Clément Gastaud, Nicolas Grislain, Aurélien Garivier, and Rémi Gribonval. Private quantiles estimation in the presence of atoms, 2023.
- Jun Li and Regina Y. Liu. New nonparametric tests of multivariate locations and scales using data depth. *Statistical Science*, 19(4):686–696, 11 2004. doi: 10.1214/088342304000000594.
- Regina Y. Liu. On a notion of simplicial depth. *National Academy of Sciences*, 85(6): 1732–1734, 1988. doi: 10.1073/pnas.85.6.1732.
- Regina Y. Liu. On a notion of data depth based on random simplices. *Annals of Statistics*., 18(1):405–414, 3 1990. doi: 10.1214/aos/1176347507.
- Regina Y Liu, Robert Joseph Serfling, and Diane L Souvaine. *Data depth: Robust multivariate analysis, computational geometry, and applications*, volume 72. American Mathematical Soc., 2006.
- Xiyang Liu, Weihao Kong, Sham Kakade, and Sewoong Oh. Robust and differentially private mean estimation. *Advances in neural information processing systems*, 34:3887–3901, 2021.
- Xiyang Liu, Weihao Kong, and Sewoong Oh. Differential privacy and robust statistics in high dimensions. In *Conference on Learning Theory*, pages 1167–1246. PMLR, 2022.
- Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *IEEE Symposium on Foundations of Computer Science*, pages 94–103. IEEE, 10 2007. doi: 10.1109/FOCS.2007.4389483.
- Karl Mosler. *Data depth*, pages 105–131. Springer New York, New York, NY, 2002. ISBN 978-1-4613-0045-8. doi: 10.1007/978-1-4613-0045-8_4.

- Stanislav Nagy. Halfspace depth does not characterize probability distributions. *Statistical Papers*, 62(3):1135–1139, 2021.
- Kelly Ramsay and Dylan Spicker. Differentially private projection-depth-based medians. *Electron. J. Statist.*, 19(2):5457–5506, 2025. ISSN 1935-7524. doi: 10.1214/25-EJS2464.
- Kelly Ramsay, Stéphane Durocher, and Alexandre Leblanc. Integrated rank-weighted depth. *Journal of Multivariate Analysis*, 173:51 – 69, 2019. doi: <https://doi.org/10.1016/j.jmva.2019.02.001>.
- Kelly Ramsay, Aukosh Jagannath, and Shojaeddin Chenouri. An elementary concentration bound for gibbs measures arising in statistical learning theory. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL <https://openreview.net/forum?id=ZInwrlkQ3f>.
- Jeremy Seeman, Matthew Reimherr, and Aleksandra Slavković. Exact privacy guarantees for markov chain implementations of the exponential mechanism with artificial atoms. *Advances in Neural Information Processing Systems*, 34:13125–13136, 2021.
- Bodhisattva Sen. A gentle introduction to empirical process theory and applications. *Lecture Notes, Columbia University*, 2022.
- Robert Serfling. A depth function and a scale curve based on spatial quantiles. In *Statistical Data Analysis Based on the L_1 -Norm and Related Methods*, pages 25–38. Birkhäuser Basel, 2002.
- Robert Serfling. Depth functions on general data spaces, II. Formulation and maximality, with consideration of the Tukey, projection, spatial, and “contour” depths. Technical report, Serfling & Thompson Statistical Consulting, 2019.
- Robert J Serfling. Depth functions in nonparametric multivariate inference. *Data Depth: Robust Multivariate Analysis, Computational Geometry, and Applications*, pages 1–16, 2006.
- Christopher G. Small. A survey of multidimensional medians. *International Statistical Review / Revue Internationale de Statistique*, 58(3):263, 12 1990. doi: 10.2307/1403809.
- Werner Stahel. Breakdown of covariance estimators. Technical report, Eidgenössische technische Hochschule, 1981.
- Anja Struyf and Peter J Rousseeuw. Halfspace depth and regression depth characterize the empirical distribution. *Journal of Multivariate Analysis*, 69:1355153, 1999.
- Michel Talagrand. Sharper bounds for gaussian and empirical processes. *The Annals of Probability*, 22(1):28–76, 1994. ISSN 00911798.
- John W Tukey. Mathematics and the picturing of data. In *International Congress of Mathematicians*, 1974.

- Christos Tzamos, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Ilias Zadik. Optimal private median estimation under minimal distributional assumptions. *Advances in Neural Information Processing Systems*, 33:3301–3311, 2020.
- Yehuda Vardi and Cun-Hui Zhang. The multivariate L_1 -median and associated data depth. *National Academy of Sciences of the United States of America*, 97:1423–6, 03 2000. doi: 10.1073/pnas.97.4.1423.
- Banghua Zhu, Jiantao Jiao, and Jacob Steinhardt. When does the Tukey median work? *2020 IEEE International Symposium on Information Theory*, pages 1201–1206, 2020. doi: 10.1109/ISIT44484.2020.9173995.
- Yijun Zuo. Projection-based depth functions and associated medians. *The Annals of Statistics*, 31(5):1460–1490, 10 2003. doi: 10.1214/aos/1065705115.
- Yijun Zuo. Influence function and maximum bias of projection depth based estimators. *Annals of Statistics*, 32(1):189–218, 2004.
- Yijun Zuo and Robert Serfling. General notions of statistical depth function. *Annals of Statistics*, 28(2):461–482, 2000.