

Average-Case Averages: Private Algorithms for Smooth Sensitivity and Mean Estimation

Supplementary material (full version)
for anonymous NeurIPS 2019 submission number 85.

Abstract

The simplest and most widely applied method for guaranteeing differential privacy is to add instance-independent noise to a statistic of interest that is scaled to its global sensitivity. However, global sensitivity is a worst-case notion that is often too conservative for realized dataset instances. We provide methods for scaling noise in an instance-dependent way and demonstrate that they provide greater accuracy under average-case distributional assumptions.

Specifically, we consider the basic problem of privately estimating the mean of a real distribution from i.i.d. samples. The standard empirical mean estimator can have arbitrarily-high global sensitivity. We propose the trimmed mean estimator, which interpolates between the mean and the median, as a way of attaining much lower sensitivity on average while losing very little in terms of statistical accuracy.

To privately estimate the trimmed mean, we revisit the smooth sensitivity framework of Nissim, Raskhodnikova, and Smith (STOC 2007), which provides a framework for using instance-dependent sensitivity. We propose three new additive noise distributions which provide concentrated differential privacy when scaled to smooth sensitivity. We provide theoretical and experimental evidence showing that our noise distributions compare favorably to others in the literature, in particular, when applied to the mean estimation problem.

1 Introduction

Consider a sensitive dataset $x \in \mathcal{X}^n$ consisting of the records of n individuals and a real-valued function $f : \mathcal{X}^n \rightarrow \mathbb{R}$. Our goal is to estimate $f(x)$ while protecting the privacy of the individuals whose data is being used. *Differential privacy* [DMNS06] gives a formal standard of individual privacy for this problem (and many others), requiring that for all pairs of datasets $x, y \in \mathcal{X}^n$ differing in one record (called *neighbouring* datasets and denoted $d(x, y) \leq 1$), the distribution of outputs should be similar for both inputs x and y .

The most basic technique in differential privacy is to release an answer $f(x) + \nu$, where ν is instance-independent additive noise (e.g., Laplace or Gaussian) with standard deviation proportional to the *global sensitivity* GS_f of the function f . Here, $\text{GS}_f = \max_{y, z \in \mathcal{X}^n : d(y, z) \leq 1} |f(y) - f(z)|$ measures the maximum amount that f can change across all pairs of datasets differing on one entry, including those which have nothing to do with x .

Calibrating noise to global sensitivity is optimal in the worst case, but may be overly pessimistic in the average case. This is the case in many statistical settings where x consists of i.i.d. samples from a reasonably structured distribution and the goal is to estimate a summary statistic of that distribution. The main example we consider in this work is that of estimating the mean of a distribution given i.i.d. samples from that distribution. The standard estimator for the distribution mean is the sample mean. However, for distributions with unbounded support, e.g., Gaussians, the global sensitivity of the sample mean is infinite. Thus we consider a different estimator and a different measure of its sensitivity.

A more fine-grained notion of sensitivity is the *local sensitivity* of f at the dataset x , which measures the variability of f in the neighbourhood of x . That is,

$$\text{LS}_f(x) = \max_{y \in \mathcal{X}^n : d(x, y) \leq 1} |f(y) - f(x)|.$$

However, naïvely calibrating noise to $\text{LS}_f(x)$ is not sufficient to guarantee differential privacy. The reason is that the local sensitivity may itself be highly variable between neighbouring datasets, and hence the magnitude of the noise observed in a statistical release may leak information about the underlying dataset x .

The work of Nissim, Raskhodnikova, and Smith [NRS07] addressed this issue by identifying *smooth sensitivity*, an intermediate notion between local and global sensitivity, with respect to which one can calibrate additive noise while guaranteeing differential privacy. Smooth sensitivity is a pointwise upper bound on local sensitivity which is itself “smooth” in that its multiplicative variation on neighboring datasets is small. More precisely, for a smoothing parameter $t > 0$, the *t -smoothed sensitivity* of a function f at a dataset x is defined as

$$S_f^t(x) = \max_{y \in \mathcal{X}^n} e^{-t \cdot d(x, y)} \cdot \text{LS}_f(y),$$

where $d(x, y)$ denotes the number of entries in which x and y disagree. Noise distributions which simultaneously do not change much under additive shifts and multiplicative dilations at scale t can be used with smooth sensitivity to give differential privacy.

In this work, we extend the smooth sensitivity framework by identifying three new distributions from which additive noise scaled to smooth sensitivity provides *concentrated differential privacy*. We apply these techniques to the problem of mean estimation, for which we propose the *trimmed mean* as an estimator that has both high accuracy and low smooth sensitivity.

1.1 Background

Before describing our results in more detail, we recall the definition of differential privacy.

Definition 1 (Differential Privacy (DP) [DMNS06; DKMMN06]). *A randomized algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is (ε, δ) -differentially private $((\varepsilon, \delta)$ -DP) if, for all neighboring datasets $x, y \in \mathcal{X}^n$ and all (measurable) sets $S \subset \mathcal{Y}$, $\mathbb{P}[M(x) \in S] \leq e^\varepsilon \mathbb{P}[M(y) \in S] + \delta$.*

We refer to $(\varepsilon, 0)$ -differential privacy as pure differential privacy (or pointwise differential privacy) and (ε, δ) -differential privacy with $\delta > 0$ as approximate differential privacy.

Given a statistic of interest $f : \mathcal{X}^n \rightarrow \mathbb{R}$ and a private dataset $x \in \mathcal{X}^n$, the randomized algorithm given by $M(x) = f(x) + \mathbf{GS}_f \cdot Z$ is $(\varepsilon, 0)$ -differentially private for Z sampled from a Laplace distribution scaled to have mean 0 and variance $2/\varepsilon^2$. We will use the smooth sensitivity in place of the global sensitivity. That is, we analyse algorithms of the form

$$M(x) = f(x) + \mathbf{S}_f^t(x) \cdot Z$$

for Z sampled from an admissible noise distribution.

The original work of Nissim, Raskhodnikova, and Smith [NRS07] proposed three admissible noise distributions. The first such distribution, the Cauchy distribution with density $\propto \frac{1}{1+z^2}$ (and its generalizations of the form $\frac{1}{1+|z|^\gamma}$ for a constant $\gamma > 1$) can be used with smooth sensitivity to guarantee pure $(\varepsilon, 0)$ -differential privacy. These distributions have polynomially decaying tails and finitely many moments, which means they may be appropriately concentrated around zero to guarantee accuracy for a single statistic, but can easily result in inaccurate answers when used to evaluate many statistics on the same dataset. Unfortunately, inverse polynomial decay is in fact essential to obtain pure differential privacy. The exponentially decaying Laplace and Gaussian distributions, which appear much more frequently in the differential privacy literature, were shown to yield approximate (ε, δ) -differential privacy with $\delta > 0$.

A recent line of work [DR16; BS16; Mir17; BDRS18] has developed variants of differential privacy which permit tighter analyses of privacy loss over multiple releases of statistics as compared to both pure and approximate differential privacy. In particular, the notion of concentrated differential privacy (CDP) [DR16; BS16] has a simple and tight composition theorem for analyzing how privacy degrades over many releases while accommodating most of the key algorithms in the differential privacy literature, including addition of Gaussian noise calibrated to global sensitivity.

Definition 2 (Concentrated Differential Privacy (CDP) [DR16; BS16]). *A randomized algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is $\frac{1}{2}\varepsilon^2$ -concentrated differentially private ($\frac{1}{2}\varepsilon^2$ -CDP) if, for all neighboring datasets $x, y \in \mathcal{X}^n$ and all $\alpha \in (1, \infty)$, $D_\alpha(M(x) \| M(y)) \leq \frac{1}{2}\varepsilon^2 \alpha$, where $D_\alpha(P \| Q) = \frac{1}{\alpha-1} \log \mathbb{E}_{X \leftarrow P} [(P(X)/Q(X))^{\alpha-1}]$ denotes the Rényi divergence of order α .*

It is natural to ask which distributions can guarantee concentrated differential privacy when calibrated to smooth sensitivity and what advantages they provide.

1.2 Our Contributions: Smooth Sensitivity and CDP

As CDP is a relaxation of pure differential privacy, Cauchy noise and its generalizations automatically guarantee CDP. But the inherency of polynomially-heavy tails does not carry over

to CDP and, in principle, admissible noise distributions for CDP could have quasi-polynomial tails (cf. Proposition 35). Nevertheless, it is not even clear *a priori* what distribution to conjecture would have these properties. In this work, we identify three such distributions with quasi-polynomial tails and show that they provide CDP when scaled to smooth sensitivity.

Laplace Log-Normal (§3.1): The first such distribution we identify, and term the “Laplace log-normal” $\text{LLN}(\sigma)$, is the distribution of the random variable $Z = X \cdot e^{\sigma Y}$ where X is a standard Laplace, Y is a standard Gaussian, and $\sigma > 0$ is a shape parameter. This distribution has mean zero, variance $2e^{2\sigma^2}$, and satisfies the quasi-polynomial tail bound $\mathbb{P}[|Z| > z] \leq e^{-\log^2(z)/3\sigma^2}$ for large z . The following result shows that scaling Laplace log-normal noise to smooth sensitivity gives concentrated differential privacy. (See also Theorem 18.)

Proposition 3. *Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ and let $Z \leftarrow \text{LLN}(\sigma)$ for some $\sigma > 0$. Then, for all $s, t > 0$, the randomized algorithm $M : \mathcal{X}^n \rightarrow \mathbb{R}$ given by $M(x) = f(x) + \frac{1}{s} \cdot \mathbf{S}_f^t(x) \cdot Z$ guarantees $\frac{1}{2}\varepsilon^2$ -CDP for $\varepsilon = t/\sigma + e^{3\sigma^2/2}s$.*

Uniform Log-Normal (§3.2): The second distribution we consider, the “uniform log-normal” distribution $\text{ULN}(\sigma)$, is the distribution of $Z = U \cdot e^{\sigma Y}$ where U is uniformly distributed over $[-1, 1]$, Y is a standard Gaussian, and $\sigma > 0$ is a shape parameter. It has mean zero and variance $\frac{1}{3}e^{2\sigma^2}$, and also has the tail bound $\mathbb{P}[|Z| > z] \leq \mathbb{P}[|Y| > \log(z)/\sigma] \leq e^{-\log^2(z)/2\sigma^2}$ for all $z \geq 1$. We show that it too gives CDP with smooth sensitivity. (See Theorem 22.)

Proposition 4. *Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ and let $Z \leftarrow \text{ULN}(\sigma)$ with $\sigma \geq \sqrt{2}$. Then, for all $s, t > 0$, the algorithm $M(x) = f(x) + \frac{1}{s} \cdot \mathbf{S}_f^t(x) \cdot Z$ guarantees $\frac{1}{2}\varepsilon^2$ -CDP for $\varepsilon = t/\sigma + e^{3\sigma^2/2} \cdot \sqrt{2/\pi\sigma^2} \cdot s$.*

Arsinh-Normal (§3.3): Our final new distribution is the “arsinh-normal” which is the distribution of $Z = \frac{1}{\sigma} \sinh(\sigma Y)$ where Y is a standard Gaussian and $\sinh(y) = (e^y - e^{-y})/2$ denotes the hyperbolic sine function. This distribution has mean zero and variance $\frac{e^{2\sigma^2} - 1}{2\sigma^2}$. Theorem 25 shows that it gives CDP, albeit with a worse dependence on the smoothing parameter (t) than the Laplace log-normal and uniform log-normal distributions.

Proposition 5. *Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ and let $Z = \sinh(Y)$ where Y is a standard Gaussian. Then, for all $s, t \in (0, 1)$, the algorithm $M(x) = f(x) + \frac{1}{s} \cdot \mathbf{S}_f^t(x) \cdot Z$ guarantees $\frac{1}{2}\varepsilon^2$ -CDP for $\varepsilon = 2(\sqrt{t} + s)$.*

1.3 Our Contributions: Private Mean Estimation

We study the following basic statistical estimation problem. Let \mathcal{D} be a distribution over \mathbb{R} with mean μ and let $X = (X_1, \dots, X_n)$ consist of i.i.d. samples from \mathcal{D} . Our goal is to design a differentially private algorithm M for estimating μ from the sample X .

The algorithmic framework we propose is as follows. We begin with a crude estimate on the range of the distribution mean, assuming $\mu \in [a, b]$.¹ Then we compute a trimmed and truncated mean of the sample X . That is,

$$f(x) = \left[\frac{x_{(m+1)} + x_{(m+2)} + \cdots + x_{(n-m)}}{n - 2m} \right]_{[a,b]},$$

where $x_{(1)} \leq x_{(2)} \leq \cdots \leq x_{(n)}$ denotes the sample in sorted order (a.k.a. the order statistics) and $[y]_{[a,b]} = \min\{\max\{y, a\}, b\}$ denotes truncation to the interval $[a, b]$. In other words, $f(x)$ first discards the largest m samples and the smallest m samples from x , then computes the mean of the remaining $n - 2m$ samples, and finally projects this to the interval $[a, b]$. Then we release $M(x) = f(x) + \mathbf{S}_f^t(x) \cdot Z$ for Z sampled from an admissible distribution for t -smoothed sensitivity. This framework requires picking a noise distribution Z , a trimming parameter $m \in \mathbb{Z}$ with $n > 2m \geq 0$, and a smoothing parameter $t > 0$.

Our mean estimation framework is versatile and may be applied to many distributions. We prove the following two illustrative results for this framework. The first gives a strong accuracy guarantee under a correspondingly strong distributional assumption, whereas the second gives a weaker accuracy guarantee under minimal distributional assumptions.

Theorem 6 (Mean Estimation for Symmetric, Subgaussian Distributions. Corollary 47). *Let $\varepsilon, \sigma > 0$, $a < b$, and $n \in \mathbb{Z}$ with $n \geq O(\log((b - a)/\sigma)/\varepsilon)$. There exists a $\frac{1}{2}\varepsilon^2$ -CDP algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ such that the following holds. Let \mathcal{D} be a σ -subgaussian distribution that is symmetric about its mean $\mu \in [a, b]$. Then*

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] \leq \frac{\sigma^2}{n} + \frac{\sigma^2}{n^2} \cdot O\left(\frac{\log((b - a)/\sigma)}{\varepsilon} + \frac{\log n}{\varepsilon^2}\right).$$

The first term $\frac{\sigma^2}{n}$ is exactly the non-private error required for this problem and the second term is the cost of privacy, which is a lower order term for large n .

Theorem 7 (Mean Estimation for General Distributions. Corollary 52). *Let $\varepsilon > 0$ and $\underline{\sigma} > 0$ and $a < b$. Let $n \geq O(\log(n(b - a)/\underline{\sigma})/\varepsilon)$. Then there exists a $\frac{1}{2}\varepsilon^2$ -CDP algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ such that the following holds. Let \mathcal{D} be a distribution with mean $\mu \in [a, b]$ and variance $\sigma^2 \geq \underline{\sigma}^2$. Then*

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] \leq \frac{\sigma^2}{n} \cdot O\left(\frac{\log(n(b - a)/\underline{\sigma})}{\varepsilon} + \frac{1}{\varepsilon^2}\right).$$

We stress that both Theorems 6 and 7 use the same algorithm; the only difference is in the setting of parameters and analysis. This illustrates the versatility of our algorithmic framework and that the distribution of the inputs directly translates to the accuracy of the estimate produced. We also emphasize that, while the accuracy guarantees above depend on

¹Assuming an a priori bound on the the mean is necessary to guarantee CDP. However, such a bound can, under reasonable assumptions, be discovered by an (ε, δ) -DP algorithm [KV18].

distributional assumptions on the dataset, the privacy guarantee requires no distributional assumptions, and holds without even assuming the data consists of i.i.d. draws.

In Section 7, we present an experimental evaluation of our approach when applied to Gaussian data.

1.4 Other Related Work

Prior work [BDRS18] showed that, when scaled to smooth sensitivity, Gaussian noise provides the relaxed notion of *truncated* CDP, but does *not* suffice to give CDP itself; see Section 3.5. Other than this and the original three distributions mentioned in Section 1.1, no other distributions have (to the best of our knowledge) been shown to provide differential privacy when scaled to smooth sensitivity.

We remark that smooth sensitivity is not the only way to exploit instance-dependent sensitivity. The most notable example is the “propose-test-release” framework of Dwork and Lei [DL09]. Roughly, their approach is as follows. First an upper bound on the local sensitivity is proposed. The validity of this bound is then tested in a differentially private manner.² If the test passes, then this bound can be used in place of the global sensitivity to release the desired quantity. If the test fails, the algorithm terminates without producing an estimate. This approach inherently requires relaxing to approximate differential privacy, to account for the small probability that the test passes erroneously.

Dwork and Lei [DL09] apply their method to the trimmed mean to obtain asymptotic results (rather than finite sample bounds like ours). They obtain a bound on the local sensitivity of the trimmed mean using the interquartile range of the data, which is itself estimated by a propose-test-release algorithm. Then they add noise proportional to this bound. This requires relaxing to approximate differential privacy and also requires assuming that the data distribution has sufficient density at the truncation points.

The mean and median (the extreme cases of the trimmed mean) have both been studied extensively in the differential privacy literature. We limit our discussion here to work in the central model of differential privacy, which is the model we study, though there has also been much work on mean estimation in the local model of privacy [JKMW18; GRS18; DR19].

Nissim, Raskhodnikova, and Smith [NRS07] analyze the smooth sensitivity of the median, but they do not apply it to mean estimation or give any average-case bounds for the smooth sensitivity of the median.

Smith [Smi08] gave a general method for private point estimation, with asymptotic efficiency guarantees for general “asymptotically normal” estimators. This method is ultimately based on global sensitivity and, in large part due to its generality, does not provide good finite sample complexity guarantees.

²Specifically, the differentially private test measures how far (in terms of the number of dataset elements that need to be changed) the input dataset is from any dataset with local sensitivity higher than the proposed bound. The test is only passed if the estimated distance is large enough to be confident that it is truly nonzero, despite the noise added for privacy. Thus, for the algorithm to succeed, we need the bound to hold not just for the realized dataset instance, but also all nearby datasets.

Karwa and Vadhan [KV18] consider confidence interval estimates for Gaussians. Although they work in a different setting, their guarantees are similar to Theorem 6 (although weaker in logarithmic factors). They propose a two-step algorithm: First a crude bound on the data is computed. Then the data is truncated using this bound and the mean is estimated via global sensitivity. Kamath, Li, Singhal, and Ullman [KLSU19] provide algorithms for learning multivariate Gaussians (extending the work of Karwa and Vadhan). We note that their algorithm does not readily extend to heavy-tailed distributions like ours does (cf. Theorem 7).

Feldman and Steinke [FS18] use a median-of-means algorithm to privately estimate the mean, yielding a guarantee similar to Theorem 7. Specifically, their algorithm partitions the dataset into evenly-sized subdatasets and computes the mean of each subdataset. Then a private approximate median is computed treating each subdataset mean as a single data point. This algorithm is simple and is applicable to any low-variance distribution, but is not as accurate as our algorithm when the input data distribution is well-behaved. Our results (Theorems 6 and 7) can be viewed as providing a unified approach that *simultaneously* matches the results of Karwa and Vadhan [KV18] for Gaussians and Feldman and Steinke [FS18] for general distributions (and also everything in between).

The smooth sensitivity framework has been applied to other problems. Some examples include learning decision forests [FI17], principal component analysis [GGB17], analysis of outliers [OFS15], and analysis of graphical data [KRSY11; KNRS13; WW13].

2 Preliminaries

2.1 Differential Privacy

We refer the reader to Section 1.1 for the definitions of differential privacy and concentrated differential privacy.

The canonical CDP algorithm is Gaussian noise addition. Adding $N(0, \sigma^2)$ to a sensitivity- Δ function attains $\frac{\Delta^2}{2\sigma^2}$ -CDP.

Lemma 8. *Let $\mu, \mu', \sigma \in \mathbb{R}$ with $\sigma > 0$. Then*

$$\forall \alpha \in (1, \infty) \quad D_\alpha(N(\mu, \sigma^2) \| N(\mu', \sigma^2)) = \alpha \frac{(\mu - \mu')^2}{2\sigma^2}.$$

The relationship between concentrated differential privacy and pure/approximate differential privacy is summarized by the following lemma.

Lemma 9 ([BS16]). *Let $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ be a randomized algorithm. If M satisfies $(\varepsilon, 0)$ -differential privacy, then it satisfies $\frac{1}{2}\varepsilon^2$ -CDP. If M satisfies $\frac{1}{2}\varepsilon^2$ -CDP, then it satisfies $(\frac{1}{2}\varepsilon^2 + \varepsilon \cdot \sqrt{2 \log(1/\delta)}, \delta)$ -differential privacy for all $\delta > 0$.*

We will make specific use of the following more precise statement of the conversion from pure to concentrated differential privacy.

Proposition 10 ([BS16, Prop. 3.3]). *Let P and Q be probability distributions satisfying $D_\infty(P\|Q) \leq \varepsilon$ and $D_\infty(Q\|P) \leq \varepsilon$. Then $D_\alpha(P\|Q) \leq \frac{1}{2}\varepsilon^2\alpha$ for all $\alpha \in (1, \infty)$.*

We also make use of the following group privacy property of concentrated differential privacy.

Lemma 11. *Let P, Q, R be probability distributions. Suppose $D_\alpha(P\|R) \leq a \cdot \alpha$ and $D_\alpha(R\|Q) \leq b \cdot \alpha$ for all $\alpha \in (1, \infty)$. Then, for all $\alpha \in (1, \infty)$,*

$$D_\alpha(P\|Q) \leq \alpha \cdot (\sqrt{a} + \sqrt{b})^2 \leq 2\alpha \cdot (a + b).$$

Proof. We use the following triangle-like inequality for Rényi divergence.

Lemma 12 ([BDRS18]). *Let P, Q, R be probability distributions and $\beta, \gamma \in (1, \infty)$. Set $\alpha = \beta\gamma/(\beta + \gamma - 1)$. Then*

$$D_\alpha(P\|Q) \leq \frac{\gamma}{\gamma - 1}D_\beta(P\|R) + D_\gamma(R\|Q).$$

We fix α and we must pick $\beta, \gamma \in (1, \infty)$ satisfying $\alpha = \beta\gamma/(\beta + \gamma - 1)$ to minimize

$$\begin{aligned} D_\alpha(P\|Q) &\leq \frac{\gamma}{\gamma - 1}D_\beta(P\|R) + D_\gamma(R\|Q) \\ &\leq \frac{\gamma}{\gamma - 1}a\beta + b\gamma \\ &= \frac{\alpha\gamma}{\gamma - \alpha}a + b\gamma \\ &= \alpha \cdot \left(\frac{u}{u - 1}a + bu \right), \end{aligned}$$

where the final equalities use the fact that $\beta = \alpha \cdot \frac{\gamma - 1}{\gamma - \alpha}$ and the substitution $\gamma = u\alpha$. We set $u = 1 + \sqrt{a/b}$ to minimize:

$$\begin{aligned} D_\alpha(P\|Q) &\leq \alpha \cdot \left(\frac{u}{u - 1}a + bu \right) \\ &= \alpha \cdot \left(\frac{1 + \sqrt{a/b}}{\sqrt{a/b}}a + b(1 + \sqrt{a/b}) \right) \\ &= \alpha \cdot (a + b + 2\sqrt{ab}) \\ &= \alpha \cdot (\sqrt{a} + \sqrt{b})^2 \\ &\leq \alpha \cdot (\sqrt{a} + \sqrt{b})^2 + \alpha \cdot (\sqrt{a} - \sqrt{b})^2 \\ &= \alpha \cdot (2a + 2b). \end{aligned}$$

□

2.2 Smooth Sensitivity

The smooth sensitivity framework was introduced by Nissim, Raskhodnikova, and Smith [NRS07]. We begin by recalling the definition of local (and global) sensitivity, as well as the local sensitivity at distance k .

Definition 13 (Global/Local Sensitivity). *Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ and $x \in \mathcal{X}^n$. The local sensitivity of f at x is defined as*

$$\text{LS}_f(x) = \sup_{x' \in \mathcal{X}^n: d(x, x') \leq 1} |f(x') - f(x)|,$$

where $d(x', x) = |\{i \in [n] : x'_i \neq x_i\}|$ is the number of entries on which x and x' differ. For $k \in \mathbb{N}$, the local sensitivity at distance k of f at x is defined as

$$\text{LS}_f^k(x) = \sup_{x' \in \mathcal{X}^n: d(x', x) \leq k} \text{LS}_f(x') = \sup_{x', x'' \in \mathcal{X}^n: d(x, x') \leq k, d(x', x'') \leq 1} |f(x'') - f(x')|.$$

The global sensitivity of f is defined as

$$\text{GS}_f = \sup_{x' \in \mathcal{X}^n} \text{LS}_f(x') = \sup_{x', x'' \in \mathcal{X}^n: d(x'', x') \leq 1} |f(x'') - f(x')|.$$

This allows us to define the smooth sensitivity.

Definition 14 (Smooth Sensitivity [NRS07]). *Let $f : \mathcal{X}^n \rightarrow \mathbb{R}$ and $x \in \mathcal{X}^n$. For $t > 0$, we define the t -smoothed sensitivity of f at x as*

$$\text{S}_f^x(t) = \max_{k \geq 0} e^{-tk} \text{LS}_f^k(x) = \sup_{x', x'' \in \mathcal{X}^n: d(x', x'') \leq 1} e^{-td(x, x')} |f(x'') - f(x')|,$$

where $d(x', x) = |\{i \in [n] : x'_i \neq x_i\}|$ is the number of entries on which x and x' differ.

The smooth sensitivity has two key properties: First, it is an upper bound on the local sensitivity. Second, it has multiplicative sensitivity bounded by t . Indeed, it can be shown that it is the smallest function with these two properties:

Lemma 15 ([NRS07]). *Let $f, g : \mathcal{X}^n \rightarrow \mathbb{R}$. Suppose that, for all $x, x' \in \mathcal{X}^n$ differing in a single entry,*

$$g(x) \geq \text{LS}_f(x) \quad \text{and} \quad g(x) \leq e^t g(x').$$

Then $g(x) \geq \text{S}_f^t(x)$.

For our applications, we could replace the smooth sensitivity with any function g satisfying the conditions of the above lemma. This is useful if, for example, computing the smooth sensitivity exactly is computationally expensive and we instead wish to use an approximation. Nissim, Raskhodnikova, and Smith [NRS07] show, for instance, that an upper bound on the smooth sensitivity of the median can be computed in sublinear time.

2.3 Subgaussian Distributions

We study the class of subgaussian distributions.

Definition 16. A distribution \mathcal{D} on \mathbb{R} is σ -subgaussian if

$$\exists \mu \in \mathbb{R} \forall t \in \mathbb{R} \quad \mathbb{E}_{X \leftarrow \mathcal{D}} [e^{t(X-\mu)}] \leq e^{t^2 \sigma^2 / 2}.$$

We say that a distribution is subgaussian if it is σ -subgaussian for some finite σ .

Note that $N(\mu, \sigma^2)$, a Gaussian with variance σ^2 , is σ -subgaussian. Furthermore, any distribution supported on the interval $[a, b]$ is $\frac{b-a}{2}$ -subgaussian.

If \mathcal{D} is σ -subgaussian and has mean zero, then $\mathbb{E}_{X \leftarrow \mathcal{D}} [X^2] \leq \sigma^2$ and, for all $\lambda > 0$,

$$\mathbb{P}_{X \leftarrow \mathcal{D}} [X \geq \lambda] = \mathbb{E}_{X \leftarrow \mathcal{D}} [\mathbb{I}[X - \lambda \geq 0]] \leq \mathbb{E}_{X \leftarrow \mathcal{D}} [e^{t(X-\lambda)}] = e^{t^2 \sigma^2 / 2 - t\lambda} = e^{-\lambda^2 / 2\sigma^2},$$

where the final equality follows from setting $t = \lambda/\sigma^2$.

3 Smooth Sensitivity Noise Distributions

3.1 Laplace Log-Normal

Definition 17. Let X and Y be independent random variables with X a standard Laplace (density $e^{-|x|}/2$) and Y a standard Gaussian (density $e^{-y^2/2}/\sqrt{2\pi}$). Let $\sigma > 0$ and $Z = X \cdot e^{\sigma Y}$. The distribution of Z is denoted $\text{LLN}(\sigma)$.

Note that $\text{LLN}(\sigma)$ is a symmetric distribution. It has mean 0 and variance $2e^{2\sigma^2}$. More generally, for all $p > 0$,

$$\mathbb{E}_{Z \leftarrow \text{LLN}(\sigma)} [|Z|^p] = \Gamma(p+1) \cdot e^{\sigma^2 p^2 / 2},$$

where Γ is the gamma function satisfying $\Gamma(p+1) = p!$ if p is an integer.

Theorem 18. Let $Z \leftarrow \text{LLN}(\sigma)$ and $s, t \in \mathbb{R}$. Then, for all $\alpha \in (1, \infty)$,

$$\left. \begin{aligned} D_\alpha(Z \| e^t Z + s) \\ D_\alpha(e^t Z + s \| Z) \end{aligned} \right\} \leq \frac{\alpha}{2} \cdot \left(\frac{|t|}{|\sigma|} + e^{\frac{3}{2}\sigma^2} \cdot |s| \right)^2 \leq \alpha \cdot \left(\frac{t^2}{\sigma^2} + e^{3\sigma^2} s^2 \right).$$

Lemma 19. Let $Z \leftarrow \text{LLN}(\sigma)$ for $\sigma > 0$. Let $t \in \mathbb{R}$ and $\alpha \in (1, \infty)$. Then

$$D_\alpha(Z \| e^t Z) \leq \frac{\alpha t^2}{2\sigma^2}.$$

Proof. Let X be a standard Laplace random variable and Y an independent standard Gaussian random variable. Let $Z = X e^{\sigma Y} \sim \text{LLN}(\sigma)$. By the quasi-convexity and postprocessing properties of Rényi divergence [BS16, Lem. 2.2], we have

$$D_\alpha(Z \| e^t Z) = D_\alpha(X e^{\sigma Y} \| X e^{\sigma Y + t}) \leq \sup_x D_\alpha(x e^{\sigma Y} \| x e^{\sigma Y + t}) \leq D_\alpha(\sigma Y \| \sigma Y + t).$$

Finally, we can calculate that $D_\alpha(\sigma Y \| \sigma Y + t) = \frac{\alpha t^2}{2\sigma^2}$ [BS16, Lem. 2.4]. \square

Lemma 20. Let $Z \leftarrow \text{LLN}(\sigma)$ for $\sigma > 0$. Let $s \in \mathbb{R}$ and $\alpha \in (1, \infty)$. Then

$$D_\alpha(Z \| Z + s) \leq \min \left\{ \frac{1}{2} e^{3\sigma^2} s^2 \alpha, e^{\frac{3}{2}\sigma^2} s \right\}.$$

Proof. Let $Z \leftarrow \text{LLN}(\sigma)$ for some $\sigma > 0$. First we compute the probability density function of Z : Let $z > 0$. (Note that Z is symmetric about the origin, so $f(-z) = f(z)$.)

$$\begin{aligned} f(z) &= \frac{d}{dz} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\mathbb{P}_{X \leftarrow \text{Lap}(1)} [X \cdot e^{\sigma Y} \leq z] \right] \\ &= \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\frac{d}{dz} \mathbb{P}_{X \leftarrow \text{Lap}(1)} [X \leq z \cdot e^{-\sigma Y}] \right] \\ &= \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\frac{1}{2} e^{-|z \cdot e^{-\sigma Y}|} \cdot e^{-\sigma Y} \right] \\ &= \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} e^{-y^2/2} \cdot \frac{1}{2} e^{-|z \cdot e^{-\sigma y}|} \cdot e^{-\sigma y} dy \\ &= \frac{1}{2\sqrt{2\pi}} \int_{\mathbb{R}} e^{-y^2/2 - \sigma y} e^{-z e^{-\sigma y}} dy \\ &= \frac{1}{2\sqrt{2\pi}} \int_{\mathbb{R}} e^{-(y-\sigma)^2/2 - \sigma(y-\sigma)} e^{-z e^{-\sigma(y-\sigma)}} dy \\ &= \frac{1}{2\sqrt{2\pi}} \int_{\mathbb{R}} e^{-y^2/2 + \sigma^2/2} e^{-z e^{\sigma^2} e^{-\sigma y}} dy \\ &= \frac{e^{\sigma^2/2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-z e^{\sigma^2} e^{\sigma Y} \right) \right]. \end{aligned}$$

Next we compute the derivative of the density:

$$\begin{aligned} f'(z) &= \frac{d}{dz} \frac{e^{\sigma^2/2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-z e^{\sigma^2} e^{\sigma Y} \right) \right] \\ &= \frac{e^{\sigma^2/2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\frac{d}{dz} \exp \left(-z e^{\sigma^2} e^{\sigma Y} \right) \right] \\ &= \frac{e^{\sigma^2/2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-z e^{\sigma^2} e^{\sigma Y} \right) \cdot (-1) e^{\sigma^2} e^{\sigma Y} \right] \\ &= -e^{\sigma^2} \cdot \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} \int_{\mathbb{R}} \exp \left(-\frac{y^2}{2} + \sigma y - z e^{\sigma^2} e^{\sigma y} \right) dy \\ &= -e^{\sigma^2} \cdot \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} \int_{\mathbb{R}} \exp \left(-\frac{(y+\sigma)^2}{2} + \sigma(y+\sigma) - z e^{\sigma^2} e^{\sigma(y+\sigma)} \right) dy \\ &= -e^{\sigma^2} \cdot \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} \int_{\mathbb{R}} \exp \left(-\frac{y^2}{2} + \frac{\sigma^2}{2} - z e^{2\sigma^2} e^{\sigma y} \right) dy \\ &= -\frac{e^{2\sigma^2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-z e^{2\sigma^2} e^{\sigma Y} \right) \right]. \end{aligned}$$

Now we can bound the derivative of the log density:

$$\begin{aligned}
\frac{d}{dz} \log f(z) &= \frac{f'(z)}{f(z)} \\
&= \frac{-\frac{e^{2\sigma^2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-ze^{2\sigma^2} e^{\sigma Y} \right) \right]}{\frac{e^{\sigma^2/2}}{2} \mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-ze^{\sigma^2} e^{\sigma Y} \right) \right]} \\
&= -e^{\frac{3}{2}\sigma^2} \frac{\mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-ze^{\sigma^2} e^{\sigma Y} \right) \cdot \exp \left(-z(e^{\sigma^2} - 1)e^{\sigma^2} e^{\sigma Y} \right) \right]}{\mathbb{E}_{Y \leftarrow \mathcal{N}(0,1)} \left[\exp \left(-ze^{\sigma^2} e^{\sigma Y} \right) \right]} \\
&\in \left[-e^{\frac{3}{2}\sigma^2}, 0 \right],
\end{aligned}$$

since $0 < \exp \left(-z(e^{\sigma^2} - 1)e^{\sigma^2} e^{\sigma Y} \right) \leq 1$. Thus $|\log(f(z)) - \log(f(z+s))| \leq s \cdot e^{\frac{3}{2}\sigma^2}$ for all $s, z \in \mathbb{R}$. Consequently, $D_\infty(Z \| Z+s) \leq e^{\frac{3}{2}\sigma^2} \cdot s$ for all $s \in \mathbb{R}$. Note that $D_\alpha(Z \| Z+s) \leq D_\infty(Z \| Z+s)$ [BS16].

By Proposition 10, $D_\alpha(X \| X+s) \leq \frac{1}{2}e^{3\sigma^2} \cdot s^2$ for all $s \in \mathbb{R}$ and all $\alpha \in (1, \infty)$. \square

Proof of Theorem 18. We use Lemma 11 to combine Lemmas 19 and 20. Specifically, to bound $D_\alpha(Z \| e^t Z + s) = D_\alpha(Z - s \| e^t Z)$, we use $D_\alpha(Z - s \| Z) = D_\alpha(Z \| Z + s)$ and $D_\alpha(Z \| e^t Z)$. To bound $D_\alpha(e^t Z + s \| Z) = D_\alpha(e^t Z \| Z - s)$, we use $D_\alpha(e^t Z \| Z) = D_\alpha(Z \| e^{-t} Z)$ and $D_\alpha(Z \| Z - s)$. \square

3.1.1 Optimizing Parameters

Let $f, g : \mathcal{X}^n \rightarrow \mathbb{R}$ and $s, t, \sigma > 0$. Suppose, for all neighbouring $x, x' \in \mathcal{X}^n$, we have

$$|f(x) - f(x')| \leq g(x) \quad \text{and} \quad e^{-t}g(x) \leq g(x') \leq e^t g(x).$$

Define a randomized algorithm $M : \mathcal{X}^n \rightarrow \mathbb{R}$ by

$$M(x) = f(x) + \frac{g(x)}{s} \cdot Z \quad \text{for} \quad Z \leftarrow \text{LLN}(\sigma).$$

Then, by Theorem 18, M is $\frac{1}{2}\varepsilon^2$ -CDP for

$$\varepsilon = \frac{t}{\sigma} + e^{\frac{3}{2}\sigma^2} \cdot s.$$

Namely, we have

$$D_\alpha(M(x) \| M(x')) = D_\alpha \left(Z \left\| \frac{f(x') - f(x)}{g(x)} \cdot s + \frac{g(x')}{g(x)} \cdot Z \right. \right).$$

We also have

$$\mathbb{E} \left[(M(x) - f(x))^2 \right] = \frac{g(x)^2}{s^2} 2e^{2\sigma^2}.$$

Our goal is to minimize this error for a fixed ε and t by setting s and σ . Clearly we should set $s = e^{-\frac{3}{2}\sigma^2}(\varepsilon - t/\sigma)$ to make the constraint tight. Thus our goal is to pick $\sigma > 0$ to minimize

$$\mathbb{E} [(M(x) - f(x))^2] = \frac{g(x)^2}{e^{-3\sigma^2}(\varepsilon - t/\sigma)^2} 2e^{2\sigma^2} = \frac{2g(x)^2}{e^{-5\sigma^2}(\varepsilon - t/\sigma)^2}.$$

Note that we need $\varepsilon\sigma > t$. Now we maximize the denominator:

$$\begin{aligned} h(\sigma) &= e^{-5\sigma^2}(\varepsilon - t/\sigma)^2 \\ h'(\sigma) &= e^{-5\sigma^2}(-10\sigma)(\varepsilon - t/\sigma)^2 + e^{-5\sigma^2}2(\varepsilon - t/\sigma)(+t/\sigma^2) \\ &= -e^{-5\sigma^2}\sigma^{-3}(10\sigma^2(\varepsilon\sigma - t) - 2t) \underbrace{(\varepsilon\sigma - t)}_{>0} \\ h'(\sigma) = 0 &\iff 10\sigma^2(\varepsilon\sigma - t) - 2t = 0 \\ &\iff \frac{5\varepsilon}{t} \cdot \sigma^3 - 5\sigma^2 - 1 = 0 \\ \frac{d}{d\sigma} \left(\frac{5\varepsilon}{t} \cdot \sigma^3 - 5\sigma^2 - 1 \right) &= \frac{15\varepsilon}{t} \cdot \sigma^2 - 10\sigma = 5\sigma \left(\frac{3\varepsilon}{t} \cdot \sigma - 2 \right) > 5\sigma > 0. \end{aligned}$$

The inequality in the underbrace follows from the fact that we restrict ourselves to $\varepsilon\sigma > t$ (as otherwise the problem is infeasible). So to minimize variance, we must pick σ such that $5\frac{\varepsilon}{t}\sigma^3 - 5\sigma^2 - 1 = 0$. Since the derivative of the right hand side of this with respect to σ is strictly positive, the cubic equation has exactly one real root.

Substituting $\sigma = t/\varepsilon$ into the cubic yields a strictly negative number. Thus the solution is strictly larger than this (as required). Substituting $\sigma = \max\{2t/\varepsilon, 1/2\}$ yields a strictly positive value for the cubic, yielding an upper bound on the solution. This means we can find the solution numerically using binary search.

3.2 Uniform Log-Normal

Definition 21. Let X and Y be independent random variables, where X is uniform on $[-1, 1]$ and Y is a standard Gaussian. Let $\sigma > 0$ and $Z = X \cdot e^{\sigma Y}$. The distribution of Z is denoted $\text{ULN}(\sigma)$.

Note that $\text{ULN}(\sigma)$ is a symmetric distribution. It has mean 0 and variance $\frac{1}{3}e^{2\sigma^2}$. More generally, for all $p > 0$,

$$\mathbb{E}_{Z \leftarrow \text{ULN}(\sigma)} [|Z|^p] = \frac{1}{p+1} e^{\sigma^2 p^2 / 2}.$$

Theorem 22. Let $Z \leftarrow \text{ULN}(\sigma)$ with $\sigma \geq \sqrt{2}$ and $s, t \in \mathbb{R}$. Then, for all $\alpha \in (1, \infty)$,

$$\left. \begin{aligned} D_\alpha(Z \| e^t Z + s) \\ D_\alpha(e^t Z + s \| Z) \end{aligned} \right\} \leq \frac{\alpha}{2} \cdot \left(\frac{|t|}{|\sigma|} + e^{\frac{3}{2}\sigma^2} \sqrt{\frac{2}{\pi\sigma^2}} \cdot |s| \right)^2 \leq \alpha \cdot \left(\frac{t^2}{\sigma^2} + e^{3\sigma^2} \frac{2}{\pi\sigma^2} s^2 \right).$$

The proof of Theorem 22 closely follows that of Theorem 18

Lemma 23. Let $Z \leftarrow \text{ULN}(\sigma)$ for $\sigma > 0$. Let $t \in \mathbb{R}$ and $\alpha \in (1, \infty)$. Then

$$D_\alpha(Z \| e^t Z) \leq \frac{\alpha t^2}{2\sigma^2}.$$

The proof is identical to that of Lemma 19.

Lemma 24. Let $Z \leftarrow \text{ULN}(\sigma)$ for $\sigma^2 \geq 2$. Let $s \in \mathbb{R}$ and $\alpha \in (1, \infty)$. Then

$$D_\alpha(Z \| Z + s) \leq \min \left\{ \frac{e^{3\sigma^2}}{\pi\sigma^2} s^2 \alpha, \frac{e^{\frac{3}{2}\sigma^2}}{\sigma\sqrt{\pi/2}} s \right\}.$$

Proof. We compute the probability density f_Z of $Z \leftarrow \text{ULN}(\sigma)$. Since the distribution is symmetric we can restrict our attention to $z > 0$.

$$\begin{aligned} f_Z(z) &= \frac{d}{dz} \mathbb{P}[Z \leq z] \\ &= \frac{d}{dz} \mathbb{E}_Y \left[\mathbb{P}_X[X \cdot e^{\sigma Y} \leq z] \right] \\ &= \mathbb{E}_Y \left[\frac{d}{dz} \mathbb{P}_X[X \leq z e^{-\sigma Y}] \right] \\ &= \mathbb{E}_Y \left[\frac{1}{2} \mathbb{I}[z e^{-\sigma Y} \leq 1] \cdot e^{-\sigma Y} \right] \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{1}{2} \mathbb{I}[z e^{-\sigma y} \leq 1] \cdot e^{-\sigma y - y^2/2} dy \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{1}{2} \mathbb{I}[z e^{-\sigma(y-\sigma)} \leq 1] \cdot e^{-\sigma(y-\sigma) - (y-\sigma)^2/2} dy \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \frac{1}{2} \mathbb{I} \left[y \geq \sigma + \frac{1}{\sigma} \log z \right] \cdot e^{\sigma^2/2 - y^2/2} dy \\ &= \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} \int_{\sigma + \frac{1}{\sigma} \log z}^{\infty} e^{-y^2/2} dy \\ &= \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} \int_{g(z)}^{\infty} e^{-y^2/2} dy \\ &= \frac{e^{\sigma^2/2}}{2} \mathbb{P}_{Y \leftarrow \mathcal{N}(0,1)} [Y \geq g(z)], \end{aligned}$$

where $g(z) := \sigma + \frac{1}{\sigma} \log z$. We can also calculate the derivative of the density from the

fundamental theorem of calculus:

$$\begin{aligned}
f'_Z(z) &= \frac{d}{dz} \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} \int_{g(z)}^{\infty} e^{-y^2/2} dy \\
&= -\frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} e^{-g(z)^2/2} \cdot g'(z) \\
&= -\frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} e^{-g(z)^2/2} \cdot \frac{1}{\sigma z}.
\end{aligned}$$

Clearly $f'_Z(z) < 0$ (for $z > 0$). This shows that the distribution is unimodal. We also have a second derivative:

$$\begin{aligned}
f''_Z(z) &= \frac{d}{dz} \left(-\frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} e^{-g(z)^2/2} \cdot \frac{1}{\sigma z} \right) \\
&= -\frac{e^{\sigma^2/2}}{2\sqrt{2\pi}} e^{-g(z)^2/2} \left(-g(z)g'(z) \cdot \frac{1}{\sigma z} - \frac{1}{\sigma z^2} \right) \\
&= \frac{e^{\sigma^2/2}}{2\sqrt{2\pi}\sigma} \frac{e^{-g(z)^2/2}}{z^2} \left(\frac{\log z}{\sigma^2} + 2 \right).
\end{aligned}$$

The second derivative is negative if $z < e^{-2\sigma^2}$ and positive if $z > e^{-2\sigma^2}$. Noting that the first derivative is always negative, this shows that the maximum magnitude of the first derivative is attained at $z = e^{-2\sigma^2}$. That is, for all $z > 0$, we have

$$-\frac{e^{2\sigma^2}}{2\sqrt{2\pi}\sigma} \leq f'_Z(z) \leq 0.$$

Our goal is to bound

$$\frac{d}{dz} \log f_Z(z) = \frac{f'_Z(z)}{f_Z(z)} = \frac{-1}{\sigma} \cdot \frac{\frac{1}{z} e^{-g(z)^2/2}}{\int_{g(z)}^{\infty} e^{-y^2/2} dy}.$$

We already have a uniform upper bound on the numerator: for all $z > 0$,

$$0 \leq \frac{1}{z} e^{-g(z)^2/2} \leq e^{\frac{3}{2}\sigma^2}.$$

Next we prove lower bounds on the denominator. Firstly, for any $a \leq 0$,

$$\int_a^{\infty} e^{-y^2/2} dy \geq \int_0^{\infty} e^{-y^2/2} dy = \sqrt{\frac{\pi}{2}}.$$

Thus, for $z \leq e^{-\sigma^2}$, we have $g(z) \leq 0$ and

$$\left| \frac{d}{dz} \log f_Z(z) \right| \leq \frac{e^{\frac{3}{2}\sigma^2}}{\sigma\sqrt{\pi/2}}.$$

For $a, b \geq 0$,

$$\begin{aligned}
\int_a^\infty e^{-y^2/2} dy &\geq \int_a^{a+b} e^{-y^2/2} dy \\
&= \int_0^b e^{-(y+a)^2/2} dy \\
&\geq \int_0^b \exp\left(\frac{-1}{2} \left(\frac{y}{b} \cdot (a+b)^2 + \left(1 - \frac{y}{b}\right)a^2\right)\right) dy \\
&= \int_0^b \exp\left(\frac{-1}{2} \left(a^2 + \frac{y}{b}(2ab + b^2)\right)\right) dy \\
&= e^{-a^2/2} \int_0^b e^{-(a+b/2)y} dy \\
&= e^{-a^2/2} \int_0^b \left(\frac{d}{dy} \frac{e^{-(a+b/2)y}}{-(a+b/2)}\right) dy \\
&= e^{-a^2/2} \cdot \frac{1 - e^{-(a+b/2)b}}{a + b/2},
\end{aligned}$$

where the second inequality uses the fact that $x \mapsto -x^2$ is concave. Hence, for any $z > e^{-\sigma^2}$ and $b > 0$, we have $g(z) > 0$ and

$$\begin{aligned}
\left| \frac{d}{dz} \log f_Z(z) \right| &= \frac{1}{\sigma} \cdot \frac{\frac{1}{z} e^{-g(z)^2/2}}{\int_{g(z)}^\infty e^{-y^2/2} dy} \\
&\leq \frac{1}{\sigma} \cdot \frac{\frac{1}{z} e^{-g(z)^2/2}}{e^{-g(z)^2/2} \cdot \frac{1 - e^{-(g(z)+b/2)b}}{g(z)+b/2}} \\
&= \frac{1}{\sigma} \cdot \frac{g(z) + b/2}{z} \frac{1}{1 - e^{-(g(z)+b/2)b}} \\
&= \left(\frac{1 + b/2\sigma}{z} + \frac{\log z}{\sigma^2 z} \right) \frac{1}{1 - e^{-(g(z)+b/2)b}} \\
&\leq \left(\frac{1 + b/2\sigma}{e^{-\sigma^2}} + \frac{1}{\sigma^2 e} \right) \frac{1}{1 - e^{-b^2/2}},
\end{aligned}$$

since $\log z \leq z/e$. It only remains to set b . Setting $b = 2$ gives, for $z > e^{-\sigma^2}$ and $\sigma \geq \sqrt{2}$,

$$\left| \frac{d}{dz} \log f_Z(z) \right| \leq \left(\frac{1 + 1/\sigma}{e^{-\sigma^2}} + \frac{1}{\sigma^2 e} \right) \frac{1}{1 - e^{-2}} \leq \frac{e^{\frac{3}{2}\sigma^2}}{\sigma \sqrt{\pi/2}}.$$

To justify the final inequality above, note that it is equivalent to

$$\frac{\sqrt{\pi/2}}{1 + e^{-2}} \left(\frac{1 + \sigma}{e^{\sigma^2/2}} + \frac{1}{\sigma e^{1+\sigma^2/2}} \right) \leq 1,$$

which can be verified numerically to hold for $\sigma = \sqrt{2}$. It is also easy to show that the left hand side is a decreasing function of σ for $\sigma \geq 1$, which establishes it for all $\sigma \geq \sqrt{2}$. This bound on $\left| \frac{d}{dz} \log f_Z(z) \right|$ entails a pure differential privacy guarantee for additive distortion, which by Proposition 10 gives the concentrated differential privacy guarantee. \square

3.3 Arsinh-Normal

Theorem 25. *Let Y be a standard Gaussian and $\sigma > 0$. Let $X = \frac{1}{\sigma} \sinh(\sigma Y)$. Let $s, t \in \mathbb{R}$. Then, for all $\alpha \in (1, \infty)$,*

$$\left. \begin{array}{l} D_\alpha(X \| e^t X + s) \\ D_\alpha(e^t X + s \| X) \end{array} \right\} \leq \frac{\alpha}{2} \cdot \left(\sqrt{|t| \cdot \left(\frac{|t|}{\sigma^2} + \frac{1}{\sigma} + 2 \right)} + |s| \cdot \left(\frac{2}{3\sigma} + \frac{\sigma}{2} \right) \right)^2.$$

We can calculate that $\mathbb{E}_{Y \leftarrow N(0,1)} \left[\left(\frac{1}{\sigma^2} \sinh(\sigma Y) \right)^2 \right] = \frac{e^{2\sigma^2} - 1}{2\sigma^2}$. In particular, setting $\sigma = 2/\sqrt{3}$ and requiring $|t| \leq 1/2$, yields

$$\left. \begin{array}{l} D_\alpha(X \| e^t X + s) \\ D_\alpha(e^t X + s \| X) \end{array} \right\} \leq \frac{\alpha}{2} \cdot \left(1.81\sqrt{|t|} + 1.16|s| \right)^2$$

and $\mathbb{E}[X^2] \leq 5.03$.

Lemma 26. *Let Y be a standard Gaussian and $\sigma > 0$. Let $X = \frac{1}{\sigma} \sinh(\sigma Y)$. Then the density of X is given by*

$$f_X(x) = \frac{1}{\sqrt{2\pi}} e^{-(\operatorname{arsinh}(\sigma x))^2 / 2\sigma^2} \cdot \frac{1}{\sqrt{1 + (\sigma x)^2}}.$$

This follows from the change-of-variables lemma

Lemma 27. *Let Y be a distribution on \mathbb{R} with density f_Y . Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be an increasing and differentiable function. Let $X = g^{-1}(Y)$. Then the density of X is given by*

$$f_X(x) = f_Y(g(x)) \cdot g'(x).$$

Lemma 28. *Let Y be a standard Gaussian and $\sigma > 0$. Let $X = \frac{1}{\sigma} \sinh(\sigma Y)$. Then*

$$D_\infty(X \| X + s) \leq |s| \cdot \left(\frac{2}{3\sigma} + \frac{\sigma}{2} \right)$$

for all $s \in \mathbb{R}$.

Note that setting $\sigma = 2/\sqrt{3} \approx 1.15$ yields $D_\infty(X \| X + s) \leq |s| \cdot \frac{2}{\sqrt{3}}$ for all $s \in \mathbb{R}$.

Proof. To prove this we simply need to show that

$$\left| \frac{d}{dx} \log(f_X(x)) \right| = \left| \frac{f'_X(x)}{f_X(x)} \right| \leq \frac{2}{3\sigma} + \frac{\sigma}{2}$$

for all $x \in \mathbb{R}$, where f_X is the density of X given by

$$f_X(x) = \frac{1}{\sqrt{2\pi}} e^{-(\operatorname{arsinh}(\sigma x))^2/2\sigma^2} \cdot \frac{1}{\sqrt{1+(\sigma x)^2}}.$$

We have

$$\begin{aligned} f'_X(x) &= \frac{1}{\sqrt{2\pi}} e^{-(\operatorname{arsinh}(\sigma x))^2/2\sigma^2} \cdot \frac{-\operatorname{arsinh}(\sigma x)}{\sigma^2} \cdot \frac{\sigma}{\sqrt{1+(\sigma x)^2}} \cdot \frac{1}{\sqrt{1+(\sigma x)^2}} \\ &\quad + \frac{1}{\sqrt{2\pi}} e^{-(\operatorname{arsinh}(\sigma x))^2/2\sigma^2} \cdot \frac{-1}{2(1+(\sigma x)^2)^{3/2}} \cdot 2\sigma^2 x \\ &= -f_X(x) \cdot \left(\frac{\operatorname{arsinh}(\sigma x)}{\sigma \sqrt{1+(\sigma x)^2}} + \frac{\sigma^2 x}{1+(\sigma x)^2} \right) \\ &= -f_X(x) \cdot \left(\frac{1}{\sigma} \cdot \frac{\operatorname{arsinh}(u)}{\sqrt{1+u^2}} + \sigma \cdot \frac{u}{1+u^2} \right), \end{aligned}$$

where $u = \sigma x$. We have $\left| \frac{u}{1+u^2} \right| \leq \frac{1}{2}$ and $\left| \frac{\operatorname{arsinh}(u)}{\sqrt{1+u^2}} \right| \leq \frac{2}{3}$ for all $u \in \mathbb{R}$. This yields the result. \square

Lemma 29. *Let Y be a standard Gaussian and $\sigma > 0$. Let $X = \frac{1}{\sigma} \sinh(\sigma Y)$. Then*

$$D_\alpha(X \| e^t X) \leq \alpha \cdot \frac{t^2}{2\sigma^2} + \frac{|t|}{2\sigma} + \max\{0, t\} \leq \alpha |t| \left(\frac{|t|}{2\sigma^2} + \frac{1}{2\sigma} + 1 \right)$$

for all $t \in \mathbb{R}$ and all $\alpha \in (1, \infty)$.

Proof. Fix $t \in \mathbb{R}$ and all $\alpha \in (1, \infty)$. We have

$$D_\alpha(X \| e^t X) = D_\alpha \left(\frac{1}{\sigma} \sinh(\sigma Y) \left\| \frac{e^t}{\sigma} \sinh(\sigma Y) \right. \right) = D_\alpha \left(Y \left\| \frac{1}{\sigma} \operatorname{arsinh}(e^t \sinh(\sigma Y)) \right. \right) = D_\alpha(Y \| g^{-1}(Y)),$$

where $g(x) = \frac{1}{\sigma} \operatorname{arsinh}(e^{-t} \sinh(\sigma x))$. By Lemma 27, the density of $g^{-1}(Y)$ is given by

$f_{g^{-1}(Y)}(y) = f_Y(g(y)) \cdot g'(y)$. Hence

$$\begin{aligned}
e^{(\alpha-1)\text{D}_\alpha(X||e^t X)} &= e^{(\alpha-1)\text{D}_\alpha(Y||g^{-1}(Y))} \\
&= \mathbb{E}_{Y \leftarrow N(0,1)} \left[\left(\frac{f_Y(Y)}{f_{g^{-1}(Y)}(Y)} \right)^{\alpha-1} \right] \\
&= \mathbb{E}_{Y \leftarrow N(0,1)} \left[\left(\frac{f_Y(Y)}{f_Y(g(Y)) \cdot g'(Y)} \right)^{\alpha-1} \right] \\
&= \mathbb{E}_{Y \leftarrow N(0,1)} \left[e^{\frac{\alpha-1}{2}(g(Y)^2 - Y^2)} \cdot \frac{1}{(g'(Y))^{\alpha-1}} \right] \\
&\leq \mathbb{E}_{Y \leftarrow N(0,1)} \left[e^{\frac{\alpha-1}{2}(g(Y)^2 - Y^2)} \right] \cdot \left(\sup_{y \in \mathbb{R}} \frac{1}{g'(y)} \right)^{\alpha-1}.
\end{aligned}$$

We now bound these terms one by one.

We have, for all $y \in \mathbb{R}$,

$$g'(y) = \frac{e^{-t} \cosh(\sigma y)}{\sqrt{1 + e^{-2t} \sinh(\sigma y)^2}} > 0.$$

Thus, for all $y \in \mathbb{R}$,

$$\begin{aligned}
\left(\frac{1}{g'(y)} \right)^2 &= \frac{1 + e^{-2t} \sinh(\sigma y)^2}{e^{-2t} \cosh(\sigma y)^2} \\
&= \frac{1 + e^{-2t} (\cosh(\sigma y)^2 - 1)}{e^{-2t} \cosh(\sigma y)^2} \\
&= 1 + \frac{1 - e^{-2t}}{e^{-2t} \cosh(\sigma y)^2} \\
&= 1 + \frac{e^{2t} - 1}{\cosh(\sigma y)^2} \\
&\leq \max\{1, e^{2t}\}.
\end{aligned}$$

This implies that $\sup_{y \in \mathbb{R}} \frac{1}{g'(y)} \leq \max\{1, e^t\} = e^{\max\{0, t\}}$.

Now we consider $y \in \mathbb{R}$ fixed and let $t \in \mathbb{R}$ vary: Define

$$h(t) = g(y) = \frac{1}{\sigma} \text{arsinh}(e^{-t} \sinh(\sigma y)).$$

Then $h(0) = y$ and

$$\begin{aligned}
h'(t) &= \frac{1}{\sigma} \frac{1}{\sqrt{1 + (e^{-t} \sinh(\sigma y))^2}} \cdot e^{-t} \sinh(\sigma y) (-1) \\
&= \frac{-1}{\sigma} \frac{1}{\sqrt{1 + (e^{-t} \sinh(\sigma y))^{-2}}} \\
&\in \left[\frac{-1}{\sigma}, 0 \right].
\end{aligned}$$

This implies that, for all $y \in \mathbb{R}$,

$$\frac{\min\{0, -t\}}{\sigma} \leq g(y) - y \leq \frac{\max\{0, -t\}}{\sigma}$$

and

$$g(y)^2 - y^2 = (g(y) - y)^2 + (g(y) - y) \cdot 2y \leq \frac{t^2}{\sigma^2} + \frac{2}{\sigma} \max\{0, -ty\}.$$

Now we calculate:

$$\begin{aligned} \mathbb{E}_{Y \leftarrow N(0,1)} \left[e^{\frac{\alpha-1}{2}(g(Y)^2 - Y^2)} \right] &\leq \mathbb{E}_{Y \leftarrow N(0,1)} \left[e^{\frac{\alpha-1}{2}(t^2/\sigma^2 + 2 \max\{0, -tY\}/\sigma)} \right] \\ &= e^{(\alpha-1)t^2/2\sigma^2} \cdot \mathbb{E}_{Y \leftarrow N(0,1)} \left[e^{(\alpha-1) \max\{0, |t|Y\}/\sigma} \right] \\ (v := (\alpha-1)|t|/\sigma \geq 0) &= e^{(\alpha-1)t^2/2\sigma^2} \cdot \left(\frac{1}{2} + \int_0^\infty e^{vy} \cdot \frac{e^{-y^2/2}}{\sqrt{2\pi}} dy \right) \\ (y \mapsto \hat{y} + v) &= e^{(\alpha-1)t^2/2\sigma^2} \cdot \left(\frac{1}{2} + \int_{-v}^\infty e^{v(\hat{y}+v)} \cdot \frac{e^{-(\hat{y}+v)^2/2}}{\sqrt{2\pi}} d\hat{y} \right) \\ &= e^{(\alpha-1)t^2/2\sigma^2} \cdot \left(\frac{1}{2} + \int_{-v}^\infty e^{v^2/2} \cdot \frac{e^{-\hat{y}^2/2}}{\sqrt{2\pi}} d\hat{y} \right) \\ &= e^{(\alpha-1)t^2/2\sigma^2} \cdot \left(\frac{1}{2} + e^{v^2/2} \cdot \mathbb{P}_{Y \leftarrow N(0,1)} [Y \geq -v] \right) \\ &\leq e^{(\alpha-1)t^2/2\sigma^2} \cdot \left(\frac{1}{2} + e^{v^2/2} \cdot \left(\frac{1}{2} + \frac{v}{\sqrt{2\pi}} \right) \right) \\ &\leq e^{(\alpha-1)t^2/2\sigma^2} \cdot e^{v^2/2+v/2} \\ &= \exp \left((\alpha-1) \left(\alpha \cdot \frac{t^2}{2\sigma^2} + \frac{|t|}{2\sigma} \right) \right). \end{aligned}$$

Now we complete the calculation:

$$\begin{aligned} e^{(\alpha-1)D_\alpha(X \| e^t X)} &\leq \mathbb{E}_{Y \leftarrow N(0,1)} \left[e^{\frac{\alpha-1}{2}(g(Y)^2 - Y^2)} \right] \cdot \left(\sup_{y \in \mathbb{R}} \frac{1}{g'(y)} \right)^{\alpha-1} \\ &\leq e^{(\alpha-1)(\alpha t^2/2\sigma^2 + |t|/2\sigma)} \cdot e^{(\alpha-1) \max\{0, t\}}. \end{aligned}$$

This implies the result:

$$D_\alpha(X \| e^t X) \leq \alpha \cdot \frac{t^2}{2\sigma^2} + \frac{|t|}{2\sigma} + \max\{0, t\}.$$

□

3.4 Student's T

Nissim, Raskhodnikova, and Smith [NRS07] showed that distributions with density $\propto \frac{1}{1+|z|^\gamma}$ for any constant $\gamma > 0$ can be scaled to smooth sensitivity to guarantee pure differential privacy. Here we show that the same is true for the family of Student's T distributions, which are defined similarly and have a number of uses in statistics.

Definition 30. *The student's T distribution with $d > 0$ degrees of freedom is denoted $\mathsf{T}(d)$ and has a probability density of*

$$f_{\mathsf{T}(d)}(x) = \frac{\Gamma\left(\frac{d+1}{2}\right)}{\sqrt{\pi d} \Gamma\left(\frac{d}{2}\right)} \left(\frac{1}{1 + \frac{x^2}{d}}\right)^{\frac{d+1}{2}}$$

The Cauchy distribution corresponds to $\mathsf{T}(1)$. The T distribution is centered at zero and has variance $\frac{d}{d-2}$ for $d > 2$. For integral $d \geq 1$ we can sample $Z \leftarrow \mathsf{T}(d)$ by letting $Z = \frac{X_0}{\sqrt{\frac{1}{d} \sum_{i=1}^d X_i^2}}$, where X_0, X_1, \dots, X_d are independent standard Gaussians.

Theorem 31. *Let $Z \leftarrow \mathsf{T}(d)$ with $d > 0$ and $s, t \in \mathbb{R}$. Then, for all $\alpha \in (1, \infty)$,*

$$\left. \begin{array}{l} D_\alpha(Z \| e^t Z + s) \\ D_\alpha(e^t Z + s \| Z) \end{array} \right\} \leq \min \left\{ |t| \cdot (d+1) + |s| \cdot \frac{d+1}{2\sqrt{d}}, \frac{\alpha}{2} \left(|t| \cdot (d+1) + |s| \cdot \frac{d+1}{2\sqrt{d}} \right)^2 \right\}.$$

Proof. Let $Z \leftarrow \mathsf{T}(d)$. We handle the multiplicative distortion first and then the additive distortion. Define

$$g(t) = \log \left(\frac{f_{\mathsf{T}}(e^t x)}{f_{\mathsf{T}}(x)} \right) = \frac{d+1}{2} (\log(d+x^2) - \log(d+e^{2t}x^2)).$$

Then

$$g'(t) = -\frac{d+1}{2} \frac{2e^{2t}x^2}{d+e^{2t}x^2}$$

and

$$|g'(t)| \leq d+1.$$

Thus $D_\infty(Z \| e^t Z) \leq |t| \cdot (d+1)$. Next we handle the additive distortion. Define

$$h(s) = \log \left(\frac{f_{\mathsf{T}}(x+s)}{f_{\mathsf{T}}(x)} \right) = \frac{d+1}{2} (\log(d+x^2) - \log(d+(x+s)^2)).$$

Then

$$h'(s) = -\frac{d+1}{2} \frac{2(x+s)}{d+(x+s)^2} = -\frac{d+1}{d/y+y}$$

for $y = x+s$. The magnitude is maximized for $y = \sqrt{d}$, yielding the bound

$$|h'(s)| \leq \frac{d+1}{2\sqrt{d}}.$$

Thus $D_\infty(Z \| Z+s) \leq |s| \cdot \frac{d+1}{2\sqrt{d}}$. Combining the multiplicative and additive bounds via a triangle-like inequality for max divergence and Proposition 10 yields the result. \square

3.5 Gaussian & tCDP

For comparison, we consider the Normal distribution under tCDP [BDRS18] – a relaxation of CDP. First we state the definition of tCDP.

Definition 32 (Truncated CDP [BDRS18]). *A randomized algorithm $M : \mathcal{X}^n \rightarrow \mathcal{Y}$ is $(\frac{1}{2}\varepsilon^2, \omega)$ -truncated CDP ($(\frac{1}{2}\varepsilon^2, \omega)$ -tCDP) if, for all $x, x' \in \mathcal{X}^n$ differing in a single entry,*

$$\sup_{\alpha \in (1, \omega)} \frac{1}{\alpha} D_\alpha(M(x) \| M(x')) \leq \frac{1}{2} \varepsilon^2,$$

where $D_\alpha(\cdot \| \cdot)$ denotes the Rényi divergence of order α .

Gaussian noise with variance σ^2 scaled to t -smooth sensitivity provides $(\frac{1}{2\sigma^2(1-\omega(1-e^{-t}))} + \frac{t^2}{4(1-\omega(1-e^{-t}))^2}, \omega)$ -tCDP for all $\omega < \frac{1}{1-e^{-t}}$.

Lemma 33 ([BDRS18]). *Let $\mu_0, \mu_1, \sigma, t, \gamma, \alpha \in \mathbb{R}$. If $\alpha > 1$ and $\alpha(e^t - 1) + 1 \geq \gamma > 0$, then*

$$D_\alpha(\mathcal{N}(\mu_0, \sigma^2) \| \mathcal{N}(\mu_1, e^t \sigma^2)) \leq \alpha \cdot \left(\frac{(\mu_0 - \mu_1)^2}{2\gamma \cdot \sigma^2} + \frac{t^2}{4\gamma^2} \right).$$

3.6 Laplace & Approximate DP

As a final comparison, we consider the approximate DP guarantees of Laplace noise. This is a sharpening of the analysis of Nissim, Raskhodnikova, and Smith [NRS07].

Theorem 34. *Let X be a standard Laplace random variable and $s, t \in \mathbb{R}$. Then, for all measurable $E \subset \mathbb{R}$,*

$$\mathbb{P}[e^t X + s \in E] \leq e^\varepsilon \mathbb{P}[X \in E] + \delta$$

for any $\delta \in (0, e^{-2})$ and

$$\varepsilon \geq |s| + (e^{|t|} - 1) \log(1/\delta) - |t|.$$

Proof. Fix $s, t \in \mathbb{R}$, $c > 0$, and a measurable set E . Since $-|x - s|e^{-t} = -|x - s| + (1 - e^{-t})|x - s| \leq -|x| + |s| + (1 - e^{-t})|x - s|$, we have

$$\begin{aligned} \mathbb{P}[e^t X + s \in E] &= \int_E \frac{1}{2e^t} e^{-|x-s|e^{-t}} dx \\ &= \int_{E \cap [s-c, s+c]} \frac{1}{2e^t} e^{-|x-s|e^{-t}} dx + \int_{E \setminus [s-c, s+c]} \frac{1}{2e^t} e^{-|x-s|e^{-t}} dx \\ &\leq \int_{E \cap [s-c, s+c]} \frac{1}{2e^t} e^{-|x|+|s|+\max\{0, c(1-e^{-t})\}} dx + \int_{\mathbb{R} \setminus [s-c, s+c]} \frac{1}{2e^t} e^{-|x-s|e^{-t}} dx \\ &= e^{|s|-t+\max\{0, c(1-e^{-t})\}} \mathbb{P}[X \in E \cap [s-c, s+c]] + \frac{1}{e^t} \int_c^\infty e^{-ye^{-t}} dy \\ &\leq e^{|s|-t+\max\{0, c(1-e^{-t})\}} \mathbb{P}[X \in E] + e^{-ce^{-t}}. \end{aligned}$$

Setting $c = e^t \log(1/\delta)$ yields

$$\mathbb{P}[e^t X + s \in E] \leq e^{|s|-t+\max\{0, (e^t-1)\log(1/\delta)\}} \mathbb{P}[X \in E] + \delta.$$

□

4 Lower Bounds

We provide a lower bound on the tail of any distribution providing concentrated differential privacy for smooth sensitivity. We note that this roughly matches the tail bounds attained by our three distributions.

Proposition 35. *Let $s, t, \varepsilon > 0$. Let Z be a real random variable satisfying, for all $\alpha \in (1, \infty)$,*

$$D_\alpha(e^t Z + s \| Z) \leq \frac{1}{2} \varepsilon^2 \alpha \quad \text{and} \quad D_\alpha(Z \| e^t Z - s) \leq \frac{1}{2} \varepsilon^2 \alpha.$$

Then, for all $x > 0$,

$$\mathbb{P}[|Z| > x] \geq \frac{1}{4} e^{-\varepsilon^2 \lceil \frac{1}{t} \log(1 + \frac{x}{s}(e^t - 1)) \rceil^2} = e^{-\Theta(\frac{\varepsilon}{t} \log(\frac{xt}{s}))^2}.$$

Proof. We may assume $\mathbb{P}[Z \geq 0] \geq \frac{1}{2}$. If not, we replace Z with $-Z$ in the argument below.

Fix $x > 0$ and Let $p = \mathbb{P}[Z \geq x]$. We may assume $p < \frac{1}{2}$, as otherwise the result is trivial.

By postprocessing and group privacy (Lemma 11), for all integers $k \geq 0$ and all $\alpha \in (1, \infty)$,

$$D_\alpha\left(\mathbb{I}\left[e^t Z + s \frac{e^{kt} - 1}{e^t - 1} \geq x\right] \middle\| \mathbb{I}[Z \geq x]\right) \leq D_\alpha\left(e^t Z + s \frac{e^{kt} - 1}{e^t - 1} \middle\| Z\right) \leq \frac{1}{2} \varepsilon^2 k^2 \alpha.$$

Note that the indicators above are simply Bernoulli random variables. Let $q_k = \mathbb{P}\left[e^t Z + s \frac{e^{kt} - 1}{e^t - 1} \geq x\right]$.

Then we have $D_1(q_k \| p) := D_1(\text{Bern}(q_k) \| \text{Bern}(p)) \leq \frac{1}{2} \varepsilon^2 k^2$.

Suppose $k \geq \frac{1}{t} \log\left(1 + \frac{x}{s}(e^t - 1)\right)$. Then $s \frac{e^{kt} - 1}{e^t - 1} \geq x$, whence

$$q_k = \mathbb{P}\left[e^t Z + s \frac{e^{kt} - 1}{e^t - 1} \geq x\right] \geq \mathbb{P}[Z \geq 0] \geq \frac{1}{2}.$$

Thus

$$\frac{1}{2} \varepsilon^2 k^2 \geq D_1(q_k \| p) \geq D_1\left(\frac{1}{2} \middle\| p\right) = \frac{1}{2} \log\left(\frac{1}{4p(1-p)}\right).$$

Setting $k = \lceil \frac{1}{t} \log\left(1 + \frac{x}{s}(e^t - 1)\right) \rceil = \Theta(\log(xt/s)/t)$ and rearranging yields the result:

$$p \geq p(1-p) \geq \frac{1}{4} e^{-\varepsilon^2 k^2} = e^{-\Theta(\frac{\varepsilon}{t} \log(\frac{xt}{s}))^2}.$$

□

We also provide a lower bound on the variance of the distributions providing concentrated differential privacy with smooth sensitivity.

Proposition 36. *Let $s, t, \varepsilon > 0$. Let Z be a real random variable satisfying $\mathbb{E}[Z] = 0$ and, for all $\alpha \in (1, \infty)$,*

$$D_\alpha(e^t Z + s \| Z) \leq \frac{1}{2} \varepsilon^2 \alpha \quad \text{and} \quad D_\alpha(Z \| e^{-t} Z - s) \leq \frac{1}{2} \varepsilon^2 \alpha.$$

Then, for all $k, \ell \in \mathbb{Z}$ with $k, \ell \geq 0$ and $k + \ell > 0$,

$$\mathbb{E}[Z^2] \geq \frac{s^2}{(1 - e^{-t})^2} \left(\frac{(e^{(k+\ell-1)t} - e^{(\ell-1)t} + e^{\ell t} - 1)^2}{e^{\varepsilon^2(k+\ell)^2} - 1} - (e^{\ell t} - 1)^2 \right).$$

Setting $\ell = 0$ yields

$$\mathbb{E}[Z^2] \geq \frac{s^2}{(e^t - 1)^2} \cdot \frac{(e^{kt} - 1)^2}{e^{\varepsilon^2 k^2} - 1}.$$

Proof. By group privacy (Lemma 11), for all $k, \ell \in \mathbb{Z}$ with $k, \ell \geq 0$ and all $\alpha \in (1, \infty)$,

$$D_\alpha \left(e^{kt} Z + s \frac{e^{kt} - 1}{e^t - 1} \left\| e^{-\ell t} Z - s \frac{1 - e^{-\ell t}}{1 - e^{-t}} \right. \right) \leq \frac{1}{2} \varepsilon^2 (k + \ell)^2 \alpha.$$

We next use a Rényi version of Pinsker's inequality [BS16, Lem. C.2]:

$$\forall X \forall Y \quad |\mathbb{E}[X] - \mathbb{E}[Y]| \leq \sqrt{\mathbb{E}[Y^2] (e^{D_2(X \| Y)} - 1)}.$$

Setting $X = e^{kt} Z + s \frac{e^{kt} - 1}{e^t - 1}$ and $Y = e^{-\ell t} Z - s \frac{1 - e^{-\ell t}}{1 - e^{-t}}$ in the above, we have

$$s \left(\frac{e^{kt} - 1}{e^t - 1} + \frac{1 - e^{-\ell t}}{1 - e^{-t}} \right) \leq \sqrt{\mathbb{E} \left[\left(e^{-\ell t} Z - s \frac{1 - e^{-\ell t}}{1 - e^{-t}} \right)^2 \right] (e^{\varepsilon^2(k+\ell)^2} - 1)}.$$

This rearranges to

$$e^{-2\ell t} \mathbb{E}[Z^2] + s^2 \left(\frac{1 - e^{-\ell t}}{1 - e^{-t}} \right)^2 \geq s^2 \left(\frac{e^{kt} - 1}{e^t - 1} + \frac{1 - e^{-\ell t}}{1 - e^{-t}} \right)^2 \frac{1}{e^{\varepsilon^2(k+\ell)^2} - 1}$$

and

$$\begin{aligned} \mathbb{E}[Z^2] &\geq s^2 e^{2\ell t} \left(\left(\frac{e^{kt} - 1}{e^t - 1} + \frac{1 - e^{-\ell t}}{1 - e^{-t}} \right)^2 \frac{1}{e^{\varepsilon^2(k+\ell)^2} - 1} - \left(\frac{1 - e^{-\ell t}}{1 - e^{-t}} \right)^2 \right) \\ &= \frac{s^2}{(1 - e^{-t})^2} \left(\frac{(e^{(k+\ell-1)t} - e^{(\ell-1)t} + e^{\ell t} - 1)^2}{e^{\varepsilon^2(k+\ell)^2} - 1} - (e^{\ell t} - 1)^2 \right). \end{aligned}$$

□

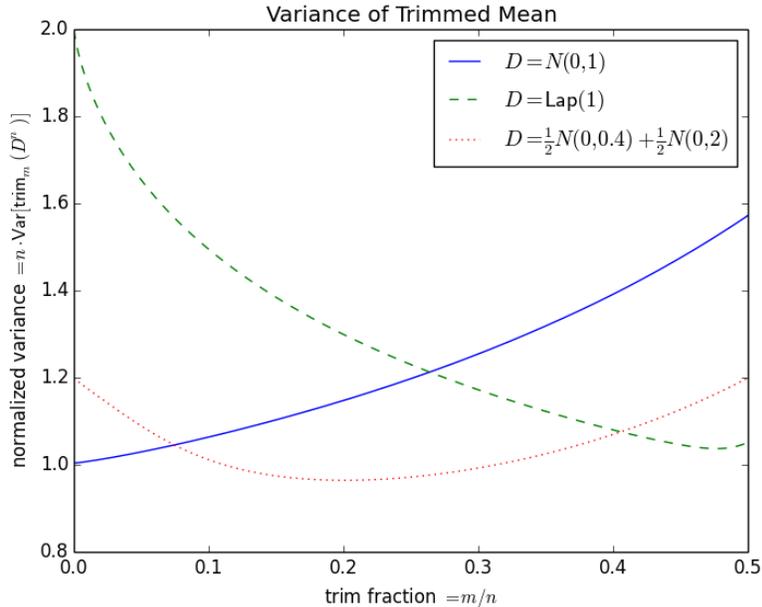


Figure 1: Variance of the Trimmed Mean for Various distributions as the trimming fraction is varied. The plot depicts experimental values for $n = 1001$ averaged over 10^6 repetitions.

5 Trimmed Mean

For the problem of mean estimation, we use the trimmed mean as our estimator.

Definition 37 (Trimmed Mean). For $n, m \in \mathbb{Z}$ with $n > 2m \geq 0$, define $\text{trim}_m : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$\text{trim}_m(x) = \frac{x_{(m+1)} + x_{(m+2)} + \cdots + x_{(n-m)}}{n - 2m},$$

where $x_{(1)} \leq x_{(2)} \leq \cdots \leq x_{(n)}$ denote the order statistics of x .

Intuitively, the trimmed mean interpolates between the mean ($m = 0$) and the median ($m = \frac{n-1}{2}$).

5.1 Error of the Trimmed Mean

Before we consider privatising the trimmed mean, we look at the error introduced by the trimming itself. We focus on mean squared error relative to the mean. That is,

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(\text{trim}_m(X) - \mu)^2],$$

where $\mu = \mathbb{E}_{X \leftarrow \mathcal{D}} [X]$ is the mean of the distribution \mathcal{D} .

We remark that mean squared error may not be the most relevant error metric for many applications. For example, the length of a confidence interval may be more relevant [KV18].

Similarly, the mean may not be the most relevant parameter to estimate. We pick this error metric as it is simple, widely-applicable, and does not require picking additional parameters (such as the confidence level).

The error of the trimmed mean depends on both the trimming fraction and also the data distribution. Figure 1 illustrates this. For Gaussian data, the optimal estimate is the empirical mean, corresponding to trimming $m = 0$ elements. This has mean squared error $\frac{1}{n}$ for n samples. As the trimming fraction is increased, the error does too. At the extreme, the median of Gaussian data has asymptotic variance $\frac{\pi}{2n} \approx \frac{1.57}{n}$. However, if the data has slightly heavier tails than Gaussian data, such as Laplacian data, then trimming actually reduces variance. The Laplacian Mean has variance $\frac{2}{n}$, while the median has asymptotic variance $\frac{1}{n}$. In between these two cases is a mixture of two Gaussians with the same mean and differing variances. Here a small amount of trimming reduces the error, but a large amount of trimming increases it again, and there is an optimal trimming fraction in between.

For our main theorems we use the following two analytic bounds. The first is a strong bound for symmetric distributions, while the second is a weaker bound for asymmetric distributions.

Proposition 38. *Let \mathcal{D} be a symmetric distribution on \mathbb{R} . Let $n, m \in \mathbb{Z}$ satisfy $n > 2m \geq 0$. Then $\text{trim}_m(X)$ is also symmetric for $X \leftarrow \mathcal{D}^n$. Moreover,*

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(\text{trim}_m(X))^2] \leq \frac{n}{(n - 2m)^2} \cdot \mathbb{E}_{X \leftarrow \mathcal{D}} [X^2].$$

This lemma follows from the symmetry of both the trimming and the distribution.

Proposition 39. *Let $n, m \in \mathbb{Z}$ satisfy $n > 2m \geq 0$. Let X_1, \dots, X_n be i.i.d. samples from a distribution with mean μ and variance σ^2 . Then*

$$\mathbb{E} [(\text{trim}_m(X) - \mu)^2] \leq \frac{n(1 + \sqrt{8m})}{(n - 2m)^2} \sigma^2 = O\left(\frac{m}{n} \sigma^2\right).$$

We first remark about the tightness of this bound: Consider the Bernoulli distribution with mean $\mu = m/2n$. With high probability ($1 - 2^{-\Omega(m)}$), the trimming removes all the 1s. Thus $\mathbb{E} [(\text{trim}_m(X) - \mu)^2] \geq 0.8\mu^2 = \frac{m^2}{5n^2}$. And we have $\frac{m}{n} \sigma^2 = \frac{m^2}{2n^2} (1 - \mu) \leq \frac{m^2}{2n^2}$. Thus the bound is tight up to constant factors (in the regime where $n - 2m = \Omega(n)$).

Proof. By affine scaling, we may assume that $\mu = 0$ and $\sigma^2 = 1$. Let X_1, \dots, X_n be i.i.d. samples from some distribution with mean zero and variance one. Let σ be a permutation on $[n]$ such that $X_{\sigma(1)} \leq X_{\sigma(2)} \leq \dots \leq X_{\sigma(n)}$ and σ is uniform (that is, ties are broken randomly).

We make σ explicit in this proof so that we can reason about its uniform distribution and we can talk about $\sigma^{-1}(i)$, which is the index of X_i in the sorted order.

For $s \subset [n]$, let σ_s denote the restriction of σ to s . That is, $s = \{\sigma_s(1), \sigma_s(2), \dots, \sigma_s(|s|)\}$ and $\sigma_s^{-1}(i) \leq \sigma_s^{-1}(j) \iff \sigma^{-1}(i) \leq \sigma^{-1}(j)$ for all $i, j \in s$.

Let

$$Y = (n - 2m)\text{trim}_m(X) = \sum_{i=m+1}^{n-m} X_{\sigma(i)} = \sum_{i=1}^n (1 - A_i) X_i,$$

where $A_i \in \{0, 1\}$ is defined by

$$1 - A_i = 1 \iff m + 1 \leq \sigma^{-1}(i) \leq n - m.$$

Note that $\mathbb{E}[A_i] = \frac{2m}{n}$ for all $i \in [n]$.

Our goal is to bound

$$\begin{aligned} \mathbb{E}[Y^2] &= \sum_{i,j \in [n]} \mathbb{E}[(1 - A_i)X_i(1 - A_j)X_j] \\ &= n - \sum_{i,j \in [n]} \mathbb{E}[X_i X_j (A_i + A_j - A_i A_j)] \\ &= n - \sum_{i \in [n]} \mathbb{E}[X_i^2 A_i] \\ &\quad - \sum_{i,j \in [n]: i \neq j} \mathbb{E}[X_i X_j \tilde{A}_i^j] + \mathbb{E}[X_i X_j \tilde{A}_j^i] + \mathbb{E}[X_i X_j (A_i - \tilde{A}_i^j + A_j - \tilde{A}_j^i - A_i A_j)], \end{aligned}$$

where \tilde{A}_i^j will be defined next. For $i, j \in [n]$ with $i \neq j$,

$$1 - \tilde{A}_i^j = 1 \iff m + 1 \leq \sigma_{[n] \setminus \{j\}}^{-1}(i) \leq (n - 1) - m.$$

Since $\sigma^{-1}(i) - 1 \leq \sigma_{[n] \setminus \{j\}}^{-1}(i) \leq \sigma^{-1}(i)$, we have $1 - \tilde{A}_i^j \leq 1 - A_i$ for all $i, j \in [n]$ with $i \neq j$.

We also have $\sigma_{[n] \setminus \{j\}}^{-1}(i) = \sigma^{-1}(i) \iff \sigma^{-1}(j) > \sigma^{-1}(i)$ for all $i, j \in [n]$ with $i \neq j$. Now, for $i, j \in [n]$ with $i \neq j$, we have

$$\begin{aligned} \tilde{A}_i^j \neq A_i &\iff \sigma_{[n] \setminus \{j\}}^{-1}(i) = m = \sigma^{-1}(i) - 1 \text{ OR } \sigma_{[n] \setminus \{j\}}^{-1}(i) = n - m = \sigma^{-1}(i) \\ &\iff (\sigma^{-1}(i) = m + 1 \text{ AND } \sigma^{-1}(j) < \sigma^{-1}(i)) \text{ OR } (\sigma^{-1}(i) = n - m \text{ AND } \sigma^{-1}(j) > \sigma^{-1}(i)) \\ &\iff (\sigma^{-1}(i) = m + 1 \text{ AND } \sigma^{-1}(j) < m + 1) \text{ OR } (\sigma^{-1}(i) = n - m \text{ AND } \sigma^{-1}(j) > n - m). \end{aligned}$$

Since σ is a uniformly random permutation, we conclude that, for $i, j \in [n]$ with $i \neq j$,

$$\mathbb{P}[\tilde{A}_i^j \neq A_i] \leq \mathbb{P}[\sigma^{-1}(j) < \sigma^{-1}(i) = m + 1] + \mathbb{P}[\sigma^{-1}(j) > \sigma^{-1}(i) = n - m] \leq \frac{2m}{n(n-1)}.$$

Now we observe that (by construction) the pair (X_i, \tilde{A}_i^j) is independent from X_j for all $i, j \in [n]$ with $i \neq j$.

We return to our calculation:

$$\begin{aligned}
\mathbb{E} [Y^2] &\leq n - \sum_{i \in [n]} \mathbb{E} [X_i^2 0] \\
&\quad - \sum_{i, j \in [n]: i \neq j} \mathbb{E} [X_i \tilde{A}_i^j] \mathbb{E} [X_j] + \mathbb{E} [X_i] \mathbb{E} [X_j \tilde{A}_j^i] + \mathbb{E} [X_i X_j (A_i - \tilde{A}_i^j + A_j - \tilde{A}_j^i - A_i A_j)] \\
&= n - \sum_{i, j \in [n]: i \neq j} \mathbb{E} [X_i X_j (A_i - \tilde{A}_i^j + A_j - \tilde{A}_j^i - A_i A_j)] \\
&\leq n + \sum_{i, j \in [n]: i \neq j} \sqrt{\mathbb{E} [(X_i X_j)^2] \mathbb{E} [(A_i - \tilde{A}_i^j + A_j - \tilde{A}_j^i - A_i A_j)^2]} \\
&\leq n + \sum_{i, j \in [n]: i \neq j} \sqrt{2 \mathbb{E} [\tilde{A}_i^j - A_i + \tilde{A}_j^i - A_j + A_i A_j]} \\
&\leq n + n(n-1) \sqrt{2 \left(\frac{2m}{n(n-1)} + \frac{2m}{n(n-1)} + \frac{2m}{n} \cdot \frac{2m-1}{n-1} \right)} \\
&\leq n + \sqrt{8} nm.
\end{aligned}$$

Finally,

$$\mathbb{E} [(\text{trim}_m(X))^2] = \frac{1}{(n-2m)^2} \mathbb{E} [Y^2] \leq \frac{n(1 + \sqrt{8}m)}{(n-2m)^2}$$

□

5.2 Sensitivity of Trimmed Mean

The other key property we need is that the trimmed mean has low local and smooth sensitivity.

Proposition 40. *Let $a, b, t \in \mathbb{R}$ with $a < b$ and $t \geq 0$ and $n, m, k \in \mathbb{Z}$ with $n > 2m \geq 0$ and $k \geq 0$ and $x \in \mathbb{R}^n$. Denote x_1, x_2, \dots, x_n in sorted order as $x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(n)}$. The local sensitivity of the trimmed mean at x is*

$$\text{LS}_{\text{trim}_m}(x) = \frac{\max\{x_{(n-m+1)} - x_{(m+1)}, x_{(n-m)} - x_{(m)}\}}{n-2m}$$

and at distance k it is

$$\text{LS}_{\text{trim}_m}^x(k) = \begin{cases} \frac{\max_{\ell=0}^{k+1} x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)}}{n-2m} & \text{if } k < m \\ \infty & \text{if } k \geq m \end{cases}.$$

The t -smooth sensitivity of the trimmed mean restricted to inputs in $[a, b]$ – that is, $\text{trim}_m : [a, b]^n \rightarrow [a, b]$ – is

$$\text{S}_{\text{trim}_m}^t(x) = \frac{1}{n-2m} \max_{k=0}^n e^{-kt} \max_{\ell=0}^{k+1} x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)},$$

where we define $x_{(i)} = a$ for $i \leq 0$ and $x_{(i)} = b$ for $i > n$.

The proof of this is a direct extension of the analysis of the smooth sensitivity of the median by Nissim, Raskhodnikova, and Smith [NRS07]. There is a $O(n \log n)$ -time algorithm for computing the smooth sensitivity.

6 Average-Case Mean Estimation via Smooth Sensitivity of Trimmed Mean

Having compiled the relevant tools in the previous sections, we turn to applying them to the problem of mean estimation. We consider an average-case distributional setting. We have an unknown distribution \mathcal{D} on \mathbb{R} and our goal is to estimate the mean $\mu = \mathbb{E}_{X \leftarrow \mathcal{D}} [X]$, given n independent samples X_1, \dots, X_n from \mathcal{D} .

Our non-private comparison point is the (un-trimmed) empirical mean $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$. This is unbiased – that is, $\mathbb{E}_{X \leftarrow \mathcal{D}^n} [\bar{X}] = \mu$ and has variance $\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(\bar{X} - \mu)^2] = \frac{\sigma^2}{n}$, where $\sigma^2 = \mathbb{E}_{X \leftarrow \mathcal{D}} [(X - \mu)^2]$.

We make the assumption that some loose bound $\mu \in [a, b]$ is known. Our results will only pay logarithmically in $b - a$, so this bound need not be tight. In general, some dependence on this range is required.

In our situation the inputs may be unbounded. This means the trimmed mean has infinite global sensitivity and thus infinite smooth sensitivity. Thus we must do something to control the sensitivity. We apply truncation:

Definition 41 (Truncation). *For $a, b, x \in \mathbb{R}$ with $a < b$, define*

$$[x]_{[a,b]} = \begin{cases} b & \text{if } x > b \\ x & \text{if } a \leq x \leq b \\ a & \text{if } x < a \end{cases} .$$

For $x \in \mathbb{R}^n$ and $a < b$, define $[x]_{[a,b]} = ([x_1]_{[a,b]}, [x_2]_{[a,b]}, \dots, [x_n]_{[a,b]})$.

6.1 Truncation of Inputs

By truncating inputs before applying the trimmed mean, we obtain the following error bound. This holds for symmetric and subgaussian distributions.

The key is that, if we know that the distribution \mathcal{D} is σ -subgaussian and its mean lies within $[a, b]$, then we can truncating the inputs to the range $[a - O(\sigma \log n), b + O(\sigma \log n)]$ without significantly affecting them.

Proposition 42. *Let \mathcal{D} be a symmetric $\bar{\sigma}$ -subgaussian distribution on \mathbb{R} . Let $\mu = \mathbb{E}_{X \leftarrow \mathcal{D}} [X]$*

and $\sigma^2 = \mathbb{E}_{X \leftarrow \mathcal{D}} [(X - \mu)^2]$. Let $a < \mu < b$. Let $n, m \in \mathbb{Z}$ satisfy $n > 2m \geq 0$. Then

$$\begin{aligned} \mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(\text{trim}_m([X]_{[a,b]}) - \mu \right)^2 \right] &\leq \left(\sqrt{\frac{n\sigma^2}{(n-2m)^2}} + \sqrt{\frac{2n\bar{\sigma}^2}{n-2m} (e^{-(b-\mu)^2/2\bar{\sigma}^2} + e^{-(\mu-a)^2/2\bar{\sigma}^2})} \right)^2 \\ &= \frac{\sigma^2}{n} \left(1 + \frac{2m}{n-2m} + \sqrt{\frac{\bar{\sigma}^2}{\sigma^2} \frac{2n^2}{n-2m} (e^{-(b-\mu)^2/2\bar{\sigma}^2} + e^{-(\mu-a)^2/2\bar{\sigma}^2})} \right)^2 \\ &= \frac{\sigma^2}{n} \left(1 + O\left(\frac{m}{n}\right) \right), \end{aligned}$$

where the final asymptotic statement assumes $n - 2m = \Omega(n)$ and $\min\{b - \mu, \mu - a\} \geq O(\bar{\sigma} \log n)$ and $\bar{\sigma}^2 = O(\sigma^2)$.

We remark that if \mathcal{D} is not subgaussian, but rather subexponential then a similar bound can be proved. This result is simply meant to be indicative of what is possible.

Proof. We may assume $\mu = 0$ without loss of generality. The result follows by combining Proposition 38 and Lemma 43 with the following reformulation of the Cauchy-Schwartz inequality.

$$\forall X, Y \quad \mathbb{E} [(X + Y)^2] \leq \left(\sqrt{\mathbb{E} [X^2]} + \sqrt{\mathbb{E} [Y^2]} \right)^2.$$

□

Lemma 43. Let \mathcal{D} be a centered σ -subgaussian distribution on \mathbb{R} . Let $n, m \in \mathbb{Z}$ satisfy $n > 2m \geq 0$. Let $a < 0 < b$. Then

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(\text{trim}_m([X]_{[a,b]}) - \text{trim}_m(X) \right)^2 \right] \leq \frac{2n}{n-2m} \cdot \sigma^2 \cdot \left(e^{-b^2/2\sigma^2} + e^{-a^2/2\sigma^2} \right)$$

Proof. We begin with the standard tail bound of subgaussians: For $t, x \geq 0$,

$$\mathbb{P}_{X \leftarrow \mathcal{D}} [X \geq x] = \mathbb{E}_{X \leftarrow \mathcal{D}} [\mathbb{I}[X \geq x]] \leq \mathbb{E}_{X \leftarrow \mathcal{D}} [e^{t(X-x)}] \leq e^{\sigma^2 t^2/2 - tx} = e^{-x^2/2\sigma^2},$$

where the final inequality follows by setting $t = x/\sigma^2$ to minimize the expression. Similarly

$\mathbb{P}_{X \leftarrow \mathcal{D}} [X \leq x] \leq e^{-x^2/2\sigma^2}$ for $x \leq 0$. Next we apply this to the quantity at hand:

$$\begin{aligned}
\mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(\text{trim}_m([X]_{[a,b]}) - \text{trim}_m(X) \right)^2 \right] &\leq \frac{1}{n-2m} \sum_{\ell=m+1}^{n-m} \mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left([X_{(\ell)}]_{[a,b]} - X_{(\ell)} \right)^2 \right] \\
&\leq \frac{1}{n-2m} \sum_{\ell=1}^n \mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left([X_{\ell}]_{[a,b]} - X_{\ell} \right)^2 \right] \\
\text{(} f_{\mathcal{D}} \text{ is the density of } \mathcal{D} \text{)} &= \frac{n}{n-2m} \left(\int_b^{\infty} (x-b)^2 f_{\mathcal{D}}(x) dx + \int_{-\infty}^a (a-x)^2 f_{\mathcal{D}}(x) dx \right) \\
\text{(integration by parts)} &= \frac{n}{n-2m} \left(\int_b^{\infty} 2(x-b) \mathbb{P}_{X \leftarrow \mathcal{D}} [X \geq x] dx \right. \\
&\quad \left. + \int_{-\infty}^a 2(a-x) \mathbb{P}_{X \leftarrow \mathcal{D}} [X \leq x] dx \right) \\
&\leq \frac{2n}{n-2m} \left(\int_b^{\infty} (x-b) e^{-x^2/2\sigma^2} dx \right. \\
&\quad \left. + \int_{-\infty}^a (a-x) e^{-x^2/2\sigma^2} dx \right) \\
&= \frac{2n}{n-2m} \int_0^{\infty} x \left(e^{-(x+b)^2/2\sigma^2} + e^{-(x-a)^2/2\sigma^2} \right) dx \\
&\leq \frac{2n}{n-2m} \int_0^{\infty} x \cdot e^{-x^2/2\sigma^2} \cdot \left(e^{-b^2/2\sigma^2} + e^{-a^2/2\sigma^2} \right) dx \\
&= \frac{2n}{n-2m} \cdot \sigma^2 \cdot \left(e^{-b^2/2\sigma^2} + e^{-a^2/2\sigma^2} \right).
\end{aligned}$$

□

Next we turn to analyzing the smooth sensitivity of the trimmed mean with truncated inputs.

Lemma 44. *Let \mathcal{D} be a σ -subgaussian distribution on \mathbb{R} . Let $a < 0 < b$. Then*

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(S_{\text{trim}_m([\cdot]_{[a,b]})}^t(X) \right)^2 \right] \leq \frac{8\sigma^2 \log n + e^{-2mt}(b-a)^2}{(n-2m)^2}.$$

Proof. By Proposition 40,

$$S_{\text{trim}_m([\cdot]_{[a,b]})}^t(x) = \frac{1}{n-2m} \max_{k=0}^n e^{-kt} \max_{\ell=0}^{k+1} x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)} \leq \frac{\max\{x_{(n)} - x_{(1)}, e^{-mt} \cdot (b-a)\}}{n-2m},$$

where the inequality follows from the fact that $x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)} \leq x_{(n)} - x_{(1)}$ when $k < m$ and $x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)} \leq b - a$ when $k \geq m$. Thus

$$\begin{aligned}
\mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(S_{\text{trim}_m([\cdot]_{[a,b]})}^t(X) \right)^2 \right] &\leq \frac{1}{(n-2m)^2} \mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[(X_{(n)} - X_{(1)})^2 + e^{-2mt}(b-a)^2 \right] \\
&\leq \frac{8\sigma^2 \log(2n) + e^{-2mt}(b-a)^2}{(n-2m)^2},
\end{aligned}$$

where the final inequality follows from Lemma 45 and the fact that $(x-y)^2 \leq 4 \max\{x^2, y^2\}$ for all $x, y \in \mathbb{R}$. □

Lemma 45 ([FS18, Lem. 4.5]). *Let \mathcal{D} be a σ -subgaussian distribution. Then $\mathbb{E}_{X \leftarrow \mathcal{D}^n} [\max_{i=1}^n X_i^2] \leq 2\sigma^2 \log(2n)$.*

Combining Proposition 42 and Lemma 44 yields the following result.

Theorem 46. *Let $s, t, \varepsilon > 0$. Let Z be a centered distribution with the property that*

$$\left. \begin{array}{l} D_\alpha(Z \| e^{\pm t} Z \pm s) \\ D_\alpha(e^{\pm t} Z \pm s \| Z) \end{array} \right\} \leq \frac{1}{2} \alpha \varepsilon^2$$

for all $\alpha \in (1, \infty)$.

Let $n, m \in \mathbb{Z}$ with $n > 2m \geq 0$. Let $a < b$ and $c, \sigma, \bar{\sigma} > 0$.

Define a randomized algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$M(x) = \text{trim}_m([x]_{[a,b]}) + \frac{1}{s} S_{\text{trim}_m([\cdot]_{[a,b]})}^t(X) \cdot Z.$$

Then M is $\frac{1}{2}\varepsilon^2$ -CDP and has the following property. Let \mathcal{D} be a distribution that is symmetric about its mean $\mu \in [a + \bar{\sigma}c, b - \bar{\sigma}c]$ and has variance σ^2 and is $\bar{\sigma}$ -subgaussian. Then

$$\begin{aligned} \mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] &\leq \frac{\sigma^2}{n} \left(1 + \frac{2m}{n - 2m} + \sqrt{\frac{\bar{\sigma}^2}{\sigma^2} \frac{2n^2}{n - 2m} (e^{-(b-\mu)^2/2\bar{\sigma}^2} + e^{-(\mu-a)^2/2\bar{\sigma}^2})} \right)^2 \\ &\quad + \frac{8\bar{\sigma}^2 \log(2n) + e^{-2mt}(b-a)^2}{(n-2m)^2 s^2} \cdot \text{Var}[Z]. \end{aligned}$$

Combining Theorem 46 with the distributions from Section 3 yields the following results. The first is a simpler case when the variance is known and the second is for when it is unknown.

Corollary 47. *Let $\varepsilon, \sigma > 0$, $a < b$, and $n \in \mathbb{Z}$ with $n \geq O(\log((b-a)/\sigma)/\varepsilon)$. There exists a $\frac{1}{2}\varepsilon^2$ -CDP algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ such that the following holds. Let \mathcal{D} be a σ -subgaussian distribution that is symmetric about its mean $\mu \in [a, b]$. Then*

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] \leq \frac{\sigma^2}{n} + \frac{\sigma^2}{n^2} \cdot O\left(\frac{\log((b-a)/\sigma)}{\varepsilon} + \frac{\log n}{\varepsilon^2}\right).$$

Corollary 48. *Let $\varepsilon \in (0, 1)$, $0 < \underline{\sigma} \leq \bar{\sigma}$, $a < b$, and $n \in \mathbb{Z}$ with $n \geq O(\log((b-a)\bar{\sigma}/\underline{\sigma}^2)/\varepsilon)$. There exists a $\frac{1}{2}\varepsilon^2$ -CDP algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ such that the following holds. Let \mathcal{D} be a distribution that is symmetric about its mean $\mu \in [a, b]$ and is σ -subgaussian with $\underline{\sigma} \leq \sigma \leq \bar{\sigma}$. Then*

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] \leq \frac{\sigma^2}{n} + \frac{\sigma^2}{n^2} \cdot O\left(\frac{\log((b-a)/\underline{\sigma}) + \log(\bar{\sigma}/\underline{\sigma})}{\varepsilon} + \frac{\log n}{\varepsilon^2}\right).$$

6.2 Truncation of Outputs

Rather than truncating the inputs to the trimmed mean, we can truncate the output. This is useful for heavier-tailed distributions and is also simpler to analyze. Indeed, the truncation only reduces error:

Lemma 49. *Let $a \leq \mu \leq b$ and let X be a random variable. Then*

$$\mathbb{E} \left[[X - \mu]_{[a,b]}^2 \right] \leq \mathbb{E} \left[(X - \mu)^2 \right].$$

Truncation of the output also controls the smooth sensitivity:

Lemma 50. *Let $n, m \in \mathbb{Z}$ with $n > 2m \geq 0$. Let $t > 0$ and $a < b$. For $x \in \mathbb{R}^n$, we have*

$$\mathbf{S}_{[\text{trim}_m(\cdot)]_{[a,b]}}^t(x) \leq \max \left\{ \frac{x_{(n)} - x_{(1)}}{n - 2m}, e^{-mt}(b - a) \right\}.$$

Let \mathcal{D} be a distribution with variance σ^2 . Then

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(\mathbf{S}_{[\text{trim}_m(\cdot)]_{[a,b]}}^t(X) \right)^2 \right] \leq \sigma^2 \frac{2n}{(n - 2m)^2} + e^{-2mt}(b - a)^2.$$

This should be contrasted with Lemma 44. The proof is also nearly identical.

Proof. Proposition 40 and the fact that $\mathbf{LS}_{[\text{trim}_m(\cdot)]_{[a,b]}}^k(x) \leq \min\{\mathbf{LS}_{\text{trim}_m}^k(x), b - a\}$ give

$$\mathbf{LS}_{[\text{trim}_m(\cdot)]_{[a,b]}}^x(k) \leq \begin{cases} \frac{\max_{\ell=0}^{k+1} x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)}}{n-2m} & \text{if } k < m \\ b - a & \text{if } k \geq m \end{cases}.$$

We thus obtain the first part of the result from Definition 14 and the fact that $x_{(n-m+1+k-\ell)} - x_{(m+1-\ell)} \leq x_{(n)} - x_{(1)}$ for $0 \leq \ell \leq k + 1 \leq m$, namely

$$\mathbf{S}_f^x(t) = \max_{k \geq 0} e^{-tk} \mathbf{LS}_f^k(x) \leq \max \left\{ \frac{x_{(n)} - x_{(1)}}{n - 2m}, e^{-mt}(b - a) \right\}.$$

Next we use the inequality $(x_{(n)} - x_{(1)})^2 \leq 2 \sum_{i=1}^n x_i^2$ to obtain

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} \left[\left(\mathbf{S}_{[\text{trim}_m(\cdot)]_{[a,b]}}^t(X) \right)^2 \right] \leq \frac{\mathbb{E} \left[(X_{(n)} - X_{(1)})^2 \right]}{(n - 2m)^2} + e^{-2mt}(b - a)^2 \leq \frac{2n\sigma^2}{(n - 2m)^2} + e^{-2mt}(b - a)^2.$$

□

Combining Proposition 39, Lemma 49, and Lemma 50 yields the following result.

Theorem 51. Let $s, t, \varepsilon > 0$. Let Z be a centered distribution with the property that

$$\left. \begin{array}{l} D_\alpha(Z \| e^{\pm t} Z \pm s) \\ D_\alpha(e^{\pm t} Z \pm s \| Z) \end{array} \right\} \leq \frac{1}{2} \alpha \varepsilon^2$$

for all $\alpha \in (1, \infty)$.

Let $n, m \in \mathbb{Z}$ with $n > 2m \geq 0$. Let $a < b$ and $\sigma > 0$. Define a randomized algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ by

$$M(x) = [\text{trim}_m(x)]_{[a,b]} + \frac{1}{s} \mathbf{S}_{[\text{trim}_m(\cdot)]_{[a,b]}}^t(x) \cdot Z.$$

Then M is $\frac{1}{2}\varepsilon^2$ -CDP and has the following property. Let \mathcal{D} be a distribution with mean $\mu \in [a, b]$ and variance σ^2 . Then

$$\begin{aligned} \mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] &\leq \frac{n(1 + \sqrt{8m})}{(n - 2m)^2} \sigma^2 + \left(\sigma^2 \frac{2n}{(n - 2m)^2} + e^{-2mt}(b - a)^2 \right) \cdot \frac{1}{s^2} \cdot \text{Var}[Z] \\ &\leq \frac{\sigma^2}{n} \cdot O\left(m + \frac{\text{Var}[Z]}{s^2}\right) + e^{-2mt}(b - a)^2. \end{aligned}$$

Combining Theorem 51 with the distributions from Section 3 yields the following result.

Corollary 52. Let $\varepsilon > 0$ and $0 < \underline{\sigma}$ and $a < b$. Let $n \geq O(\log(n(b - a)/\underline{\sigma})/\varepsilon)$. Then there exists a $\frac{1}{2}\varepsilon^2$ -CDP algorithm $M : \mathbb{R}^n \rightarrow \mathbb{R}$ such that the following holds. Let \mathcal{D} be a distribution with mean $\mu \in [a, b]$ and variance $\sigma^2 \geq \underline{\sigma}^2$. Then

$$\mathbb{E}_{X \leftarrow \mathcal{D}^n} [(M(X) - \mu)^2] \leq \frac{\sigma^2}{n} \cdot O\left(\frac{\log(n(b - a)/\underline{\sigma})}{\varepsilon} + \frac{1}{\varepsilon^2}\right).$$

7 Experimental Results

We perform an experimental evaluation of our methods, specifically the combination of the trimmed mean with various smooth sensitivity distributions applied to Gaussian data. The results are shown in Figures 2, 3, 4, 5, & 6.

7.1 Experimental Setup

We explain the experimental setup and parameter choices below.

- **Data:** Our data is sampled from a univariate Gaussian distribution. Gaussian is a natural data distribution. However, as shown in Figure 1, the trimmed mean performs better on heavier tailed distributions. That is to say, we would expect our results to be better for non-Gaussian data.

The variance of our data is set to $\sigma^2 = 1$ and we set the mean to $\mu = 0$. The truncation interval is set conservatively to $[a, b] = [-50, 1050]$. The data is truncated before applying the trimmed mean.

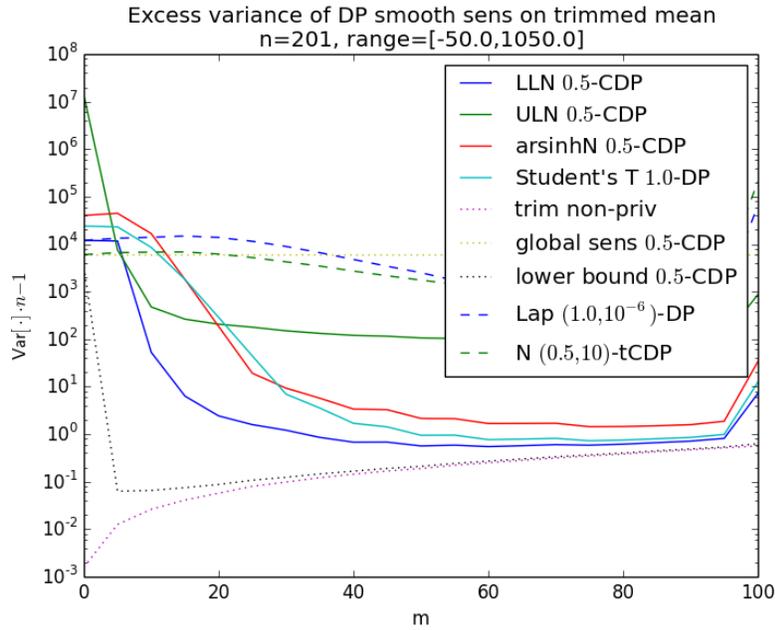


Figure 2: Excess variance of the private trimmed mean with smooth sensitivity for $N(0, 1)$ data. Here $n = 201$ and results are averaged over 10^6 repetitions.

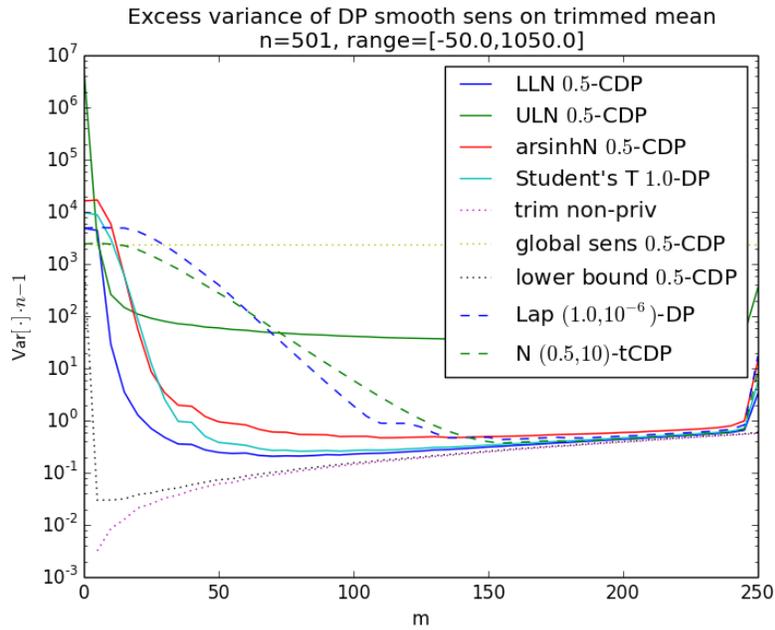


Figure 3: Excess variance of the private trimmed mean with smooth sensitivity for $N(0, 1)$ data. Here $n = 501$ and results are averaged over 10^6 repetitions.

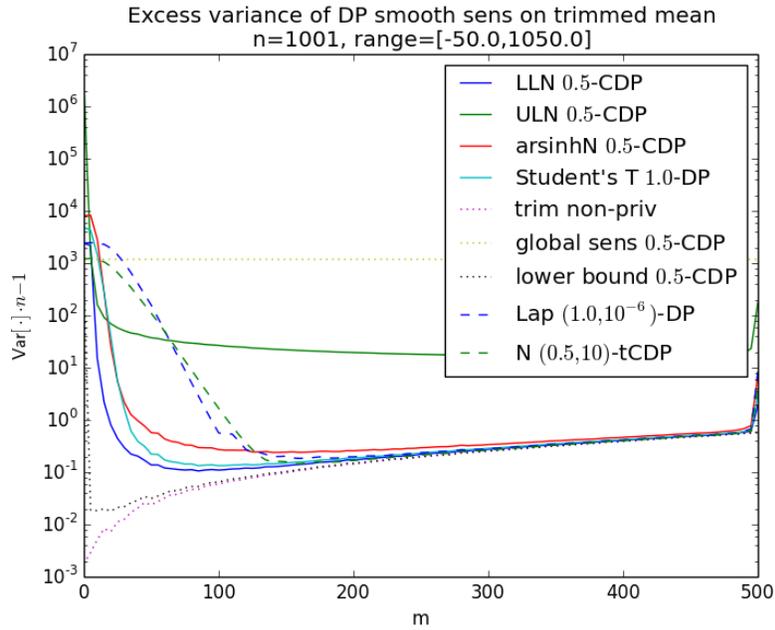


Figure 4: Excess variance of the private trimmed mean with smooth sensitivity for $N(0, 1)$ data. Here $n = 1001$ and results are averaged over 10^6 repetitions.

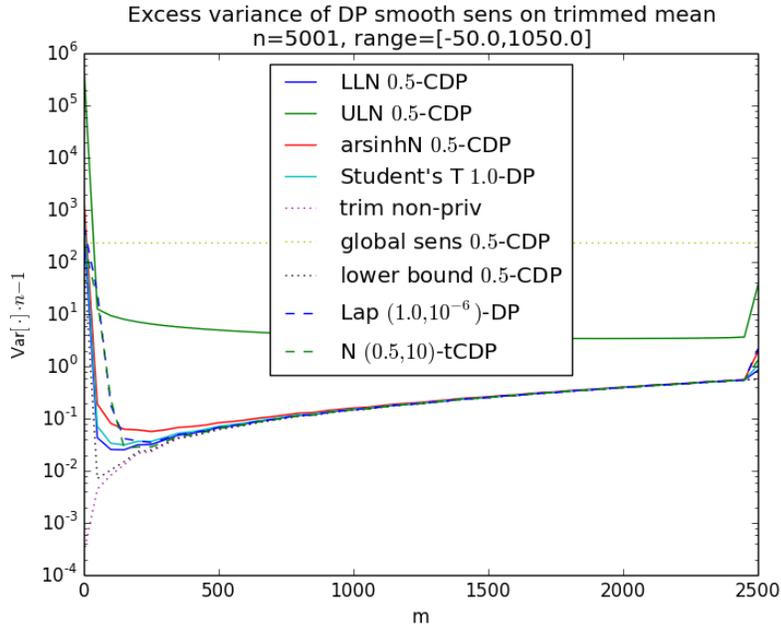


Figure 5: Excess variance of the private trimmed mean with smooth sensitivity for $N(0, 1)$ data. Here $n = 5001$ and results are averaged over 10^6 repetitions.

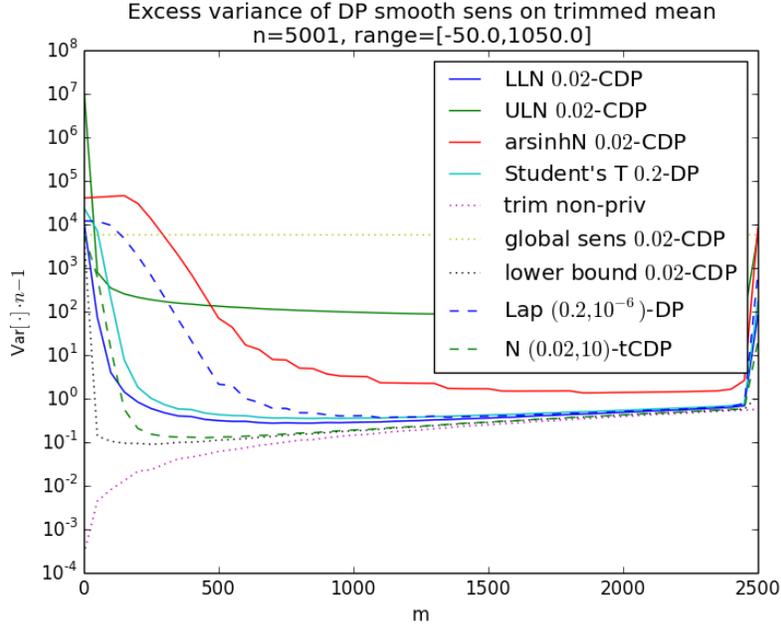


Figure 6: Excess variance of the private trimmed mean with smooth sensitivity for $N(0, 1)$ data. Here $n = 5001$ and results are averaged over 10^6 repetitions. This plot has smaller privacy parameters than the others.

- **Error:** We measure the variance or mean squared error of the various algorithms. That is,

$$\sigma^2 = \mathbb{E}_{X \leftarrow N(\mu, 1)^n} \left[\left(\text{trim}_m([X]_{[a,b]}) + \mathbf{S}_{\text{trim}_m([\cdot]_{[a,b]})}^t(X) \cdot Z - \mu \right)^2 \right],$$

where Z is an appropriately-scaled distribution suited for providing differential privacy when scaled to t -smooth sensitivity. For scaling, we multiply by n and subtract 1. Subtracting 1 corresponds to the variance of the sample mean, which is the optimal non-private error. Multiplying by n allows for a comparison of different values of n , as it normalizes by the correct convergence rate. So we plot the normalized excess variance $\sigma^2 \cdot n - 1$ (on a logarithmic scale).

- **Algorithms:** We compare our three noise distributions against three other algorithms. Three further comparison points are provided: global sensitivity with truncation, our lower bound, and the non-private error of the trimmed mean. We explain each of the lines below.
 - LLN: We evaluate the Laplace Log-Normal distribution from Section 3.1. The plot uses the privacy analysis given in Theorem 18.
 - ULN: We evaluate the Uniform Log-Normal distribution from Section 3.2. The plot uses the privacy analysis given in Theorem 22.

- **arshinhN**: We evaluate the Arsinh-Normal distribution from Section 3.3. The plot uses the privacy analysis given in Theorem 25.
- **Student’s T**: We evaluate the Student’s T distribution from Section 3.4. The plot uses the privacy analysis given in Theorem 31. We set the degrees of freedom parameter to $d = 3$.
- **trim non-priv**: We plot the line where zero noise is added for privacy. The only source of error is the trimmed mean itself. This comparison point is useful as it illustrates the fact that, in many cases, most of the error is not coming from the privacy-preserving noise.
- **global sens**: We plot the error that would be attained by truncating the data and then using this to bound global sensitivity and add Gaussian noise. This is the baseline algorithm which we compare to. Note that the comparison here depends significantly on the truncation interval $[a, b]$.
- **lower bound**: We plot the lower bound on variance given by Proposition 36. No smooth sensitivity-based algorithm can beat this (although a completely different approach might).
- **Lap**: We compare to Laplace noise. This was suggested in the original work of Nissim, Raskhodnikova, and Smith [NRS07]. The plot uses the privacy analysis given in Theorem 34.
- **N**: We compare to Gaussian noise. This was analyzed in prior work [BDRS18]; see Lemma 33.

We note that the Cauchy distribution is not included in this comparison because it has infinite variance.

- **Privacy**: The algorithms we compare satisfy different variants of differential privacy. As such, it is not possible to give a completely fair comparison. Our new distributions satisfy concentrated differential privacy, whereas the Student’s T distribution satisfies pure differential privacy. Laplace and Gaussian noise satisfy approximate differential privacy or truncated concentrated differential privacy.

To provide the fairest possible comparison, we pick a ϵ value and then compare $(\epsilon, 0)$ -differential privacy with relaxations thereof. Namely, we compare $(\epsilon, 0)$ -differential privacy with $\frac{1}{2}\epsilon^2$ -CDP, $(\frac{1}{2}\epsilon^2, 10)$ -tCDP, and $(\epsilon, 10^{-6})$ -differential privacy. Each of these is implied by $(\epsilon, 0)$ -differential privacy and the implication is fairly tight in the sense that these definitions intuitively seem to provide a roughly similar level of privacy.

Our plots use the values $\epsilon = 1$ or $\epsilon = 0.2$.

- **Other parameters**: Aside from the privacy parameters (ϵ etc.) and the dataset size (n), we must choose the trimming level (m) and the smoothing parameter (t). (Note that, given the privacy parameters and smoothing parameter, the scale parameter (s) is chosen to be as large as possible in order to minimize the noise magnitude.)

Our plots show a range of trimming levels on the horizontal axis. We numerically optimized the smoothing parameter. Specifically, the smooth sensitivity is evaluated for 150 values of the smoothing parameter (ranging from $t = 9$ to $t = 10^{-9}$, roughly evenly spaced on a logarithmic scale) and whichever attains the lowest variance for the given algorithm and other parameter values is used.

Finally, several of our distributions have a shape parameter, which we set as follows. For Laplace Log-Normal, we numerically optimize σ ; see Section 3.1.1. For Uniform Log-Normal, we set $\sigma = \sqrt{2}$, which is the smallest value permitted by our analysis (Theorem 22). For Arsinh-Normal, we set $\sigma = 2/\sqrt{3}$, which minimizes one of the terms in the analytical bound (Theorem 25). For Student’s T, we set the degrees of freedom to 3 (the smallest integer with finite variance).

7.2 Experimental Discussion & Comparison

Overall Performance: The experimental results demonstrate that for relatively moderate parameter settings ($n = 201$ and $\varepsilon = 1$ depicted in Figure 2) it is possible to privately estimate the mean with variance that is only a factor of two higher than non-privately. For $n = 1001$, it is possible to drive this excess variance down to 10%. Indeed, in these settings, the additional error introduced by trimming is more significant than that introduced by the privacy-preserving noise.

We remark that the data for these experiments is perfectly Gaussian. If the data deviates from this ideal, the robustness of the trimmed mean may actually be beneficial for accuracy (and not just privacy). Figure 1 shows that for some natural distributions the trimming does reduce variance.

Comparison of Algorithms: The results show that different algorithms perform better in different parameter regimes. However, generally, the Laplace Log-Normal distribution has the lowest variance, closely followed by the Student’s T distribution. The Arsinh-Normal distribution performs adequately, but the Uniform Log-Normal distribution performs poorly. The Laplace and Gaussian distributions from prior work often perform substantially worse than our distributions, but are better or similar in many parameter settings.

Note that the different algorithms satisfy slightly different privacy guarantees and also have very different tail behaviours. Since the variance of many of the algorithms is broadly similar, the choice of which algorithm is truly best will depend on these factors.

If the stronger pure differential privacy guarantee is preferable, the Student’s T distribution is likely best. However, this has no third moment and consequently heavy tails. This makes it bad if, for example, the goal is a confidence interval, rather than a point estimate of the mean. The lightest tails are provided by the Gaussian, but this only satisfies the weaker truncated CDP or approximate differential privacy definitions. Laplace Log-Normal is in between – it satisfies the strong concentrated differential privacy definition and has quasipolynomial tails and all its moments are finite.

References

- [BDRS18] M. Bun, C. Dwork, G. N. Rothblum, and T. Steinke. “Composable and Versatile Privacy via Truncated CDP”. In: *STOC*. 2018. URL: <http://doi.acm.org/10.1145/3188745.3188946>.
- [BS16] M. Bun and T. Steinke. “Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds”. In: *TCC*. <https://arxiv.org/abs/1605.02065>. 2016.
- [DKMMN06] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. “Our Data, Ourselves: Privacy Via Distributed Noise Generation”. In: *EUROCRYPT*. 2006.
- [DL09] C. Dwork and J. Lei. “Differential privacy and robust statistics”. In: *STOC*. 2009.
- [DMNS06] C. Dwork, F. McSherry, K. Nissim, and A. Smith. “Calibrating Noise to Sensitivity in Private Data Analysis”. In: *TCC*. <http://repository.cmu.edu/jpc/vol17/iss3/2>. 2006.
- [DR16] C. Dwork and G. Rothblum. “Concentrated Differential Privacy”. In: *CoRR* abs/1603.01887 (2016). <https://arxiv.org/abs/1603.01887>. URL: <http://arxiv.org/abs/1603.01887>.
- [DR19] J. Duchi and R. Rogers. “Lower Bounds for Locally Private Estimation via Communication Complexity”. In: *arXiv preprint arXiv:1902.00582* (2019).
- [FI17] S. Fletcher and M. Z. Islam. “Differentially private random decision forests using smooth sensitivity”. In: *Expert Systems with Applications* 78 (2017), pp. 16–31.
- [FS18] V. Feldman and T. Steinke. “Calibrating Noise to Variance in Adaptive Data Analysis”. In: *COLT*. 2018. URL: <http://proceedings.mlr.press/v75/feldman18a.html>.
- [GGB17] A. Gonen and R. Gilad-Bachrach. “Smooth Sensitivity Based Approach for Differentially Private Principal Component Analysis”. In: *arXiv preprint arXiv:1710.10556* (2017).
- [GRS18] M. Gaboardi, R. Rogers, and O. Sheffet. “Locally private mean estimation: Z-test and tight confidence intervals”. In: *arXiv preprint arXiv:1810.08054* (2018).
- [JKMW18] M. Joseph, J. Kulkarni, J. Mao, and Z. S. Wu. “Locally private gaussian estimation”. In: *arXiv preprint arXiv:1811.08382* (2018).
- [KLSU19] G. Kamath, J. Li, V. Singhal, and J. Ullman. “Privately Learning High-Dimensional Distributions”. In: *COLT*. 2019. URL: <https://arxiv.org/abs/1805.00216>.

- [KNRS13] S. P. Kasiviswanathan, K. Nissim, S. Raskhodnikova, and A. Smith. “Analyzing graphs with node differential privacy”. In: *Theory of Cryptography Conference*. Springer. 2013, pp. 457–476.
- [KRSY11] V. Karwa, S. Raskhodnikova, A. Smith, and G. Yaroslavtsev. “Private analysis of graph structure”. In: *Proceedings of the VLDB Endowment* 4.11 (2011), pp. 1146–1157.
- [KV18] V. Karwa and S. Vadhan. “Finite Sample Differentially Private Confidence Intervals”. In: *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2018.
- [Mir17] I. Mironov. “Rényi Differential Privacy”. In: *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*. 2017, pp. 263–275.
- [NRS07] K. Nissim, S. Raskhodnikova, and A. Smith. “Smooth sensitivity and sampling in private data analysis”. In: *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM. 2007, pp. 75–84.
- [OFS15] R. Okada, K. Fukuchi, and J. Sakuma. “Differentially private analysis of outliers”. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer. 2015, pp. 458–473.
- [Smi08] A. Smith. “Efficient, differentially private point estimators”. In: *arXiv preprint arXiv:0809.4794* (2008).
- [WW13] Y. Wang and X. Wu. “Preserving differential privacy in degree-correlation based graph generation”. In: *Transactions on data privacy* 6.2 (2013), p. 127.