Optimizing DP Gaussian Mean Estimation in $\mathbb{R}$

Jingyi Gao [*1] **Zhifan Jing** [*1]

Abstract

This report surveys Karwa and Vadhan's finite-sample framework for differentially private confidence intervals for the mean of a one-dimensional Gaussian distribution. We explain how their algorithms overcome the unbounded sensitivity of the empirical mean by first privately localizing the data, then clipping to a data-dependent but private range, and finally adding calibrated noise to form a conservative confidence interval. We summarize both the known-variance and unknown-variance settings, including the use of differentially private histogram learners for range and variance estimation. We then compare their upper bounds with the corresponding statistical and privacy lower bounds. The main message is that their confidence intervals achieve the classical non-private statistical rate together with an additional privacy cost that is optimal up to polylogarithmic factors.

1. Introduction

Univariate Gaussian mean estimation has a long history and has been extensively studied in statistics. Formally, given independent samples $X_1, \ldots, X_n \sim N(\mu, \sigma^2)$, where the variance σ^2 can be either known or unknown, the goal is to estimate the unknown mean μ. In (Karwa & Vadhan, 2018), this task is further studied and the novelty comes from a combination of three aspects:

1. Their algorithm is differentially private.

2. Their estimator is conservative: the algorithm outputs a confidence interval whose coverage probability is at least $1 - \alpha$, even in finite samples.

3. They studied finite-sample accuracy instead of only asymptotics, and showed that their algorithm is optimal up to some polylog terms.

[1]Department of Computer Sciences, University of Wisconsin. Correspondence to: Zhifan Jing <zjing24@wisc.edu>.

Preprint. May 1, 2026.

2. Related Work

First, without the concern of privacy, it is very a classical result (see, for example, (Lehmann & Romano, 2005)) that under the assumption of known variance, the confidence interval for a normal mean is

$$I(X_1, .., X_n) = \overline{X} \pm \frac{\sigma}{\sqrt{n}} \cdot z_{1-\alpha/2}$$

where $\overline{X}$ is the sample mean and z_a represents the a^{th} quantile of a standard normal distribution. This interval has the smallest expected size among all $1-\alpha$ level confidence sets for a normal mean. And this interval has a deterministic length of

$$|I(X_1, .., X_n)| = (2\sigma z_{1-\alpha/2})/\sqrt{n} = \Theta(\sigma\sqrt{\log(1/\alpha)/n})$$

When the variance σ^2 is unknown, the interval is (see (Lehmann & Romano, 2005))

$$I(X_1, .., X_n) = \overline{X} \pm \frac{s}{\sqrt{n}} \cdot t_{n-1,1-\alpha/2}$$

where s^2 is the sample variance ($s^2 := \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \overline{X})^2$) and $t_{n-1,a}$ is the a^{th} quantile of a t-distribution with $n-1$ degrees of freedom. Now the length is a random variable, and its expected length is

$$\mathbb{E}\left[|I(X_1, \ldots, X_n)|\right] = \frac{2\sigma}{\sqrt{n}} k_n t_{n-1,1-\alpha/2} = \Theta\left(\sigma\sqrt{\frac{\log(1/\alpha)}{n}}\right),$$

for an appropriate constant $k_n = 1 - O(1/n)$. for some appropriate constant $k_n = 1 - O(1/n)$.

Conservative hypothesis testing with differential privacy was advocated by (Gaboardi et al., 2016). There are numerous results but many of them were evaluated empirically or the conservativeness only holds in a particular asymptotic regime. In particular, both (Sheffet, 2015) and (Cai et al., 2017) did give a finite-sample analyses; but they don't have a matching lower bound and the algorithm (Sheffet, 2015) even requires the data to be bounded.

Remark that the paper ((Karwa & Vadhan, 2018)) we study does provide a nearly matching lower bound, and have no assumptions about boundedness of the data.

3. Results

In this section, we summarize the main results in (Karwa & Vadhan, 2018).

The overall structure of their approach is algorithmic rather than purely statistical. The main difficulty is that the Gaussian distribution is supported on the whole real line, so the usual empirical mean has unbounded global sensitivity. Therefore, one cannot directly add Laplace or Gaussian noise to the empirical mean without first controlling the possible effect of a single data point.

Karwa and Vadhan solve this by using a multi-stage procedure. First, they privately locate a short interval that contains the relevant part of the data. Second, they clip or truncate the data to this interval, so that the empirical mean has bounded sensitivity. Finally, they add calibrated noise and widen the resulting interval enough to guarantee coverage of the true mean.

A useful way to view the paper is that the authors separate the problem into several smaller private subroutines. The first subroutine is a differentially private histogram learner, which is used to locate a region of high probability mass. The second subroutine privately estimates a range containing most Gaussian samples when the variance is known. The third subroutine privately estimates the scale when the variance is unknown. These pieces are then combined to build differentially private confidence intervals for the Gaussian mean.

3.1. DP histogram learner

The first technical tool is a differentially private histogram learner. Suppose we divide the sample space into bins $B_1, \dots, B_K$. For a distribution $\mathcal{D}$, its histogram with respect to these bins is the vector

$$(p_1, \dots, p_K), \qquad p_i = \mathbb{P}_{X\sim\mathcal{D}}(X \in B_i).$$

If the data were not private, one could simply count how many samples fall into each bin and choose the bin with the largest empirical count. However, these counts are data-dependent, so releasing them directly may violate privacy. The private histogram learner instead adds carefully calibrated noise to the empirical bin counts or frequencies. This makes the released histogram differentially private while still preserving enough accuracy to identify a bin with large true probability mass.

In (Karwa & Vadhan, 2018), this tool is mainly used as a localization procedure. The algorithm does not need to estimate the entire distribution accurately. Rather, it only needs to find a bin that is likely to be close to the center of the Gaussian distribution. The key point is that if a bin has very small true probability, then after privatization it is still unlikely to be selected as the heaviest bin. Thus, the private histogram learner allows the algorithm to locate a meaningful region of the real line without revealing too much information about any individual sample.

As studied in (Dwork et al., 2006), (Bun et al., 2019), and (Vadhan, 2017), there exist differentially private algorithms for learning histograms. The following lemma gives the version used in (Karwa & Vadhan, 2018). The first guarantee says that each privatized bin probability is close to the true bin probability. The second guarantee is especially important for the later Gaussian algorithms: it bounds the probability that a low-mass bin is incorrectly selected as the maximum.

Lemma 3.1 (Histogram Learner)**.** *For every $K \in \mathbb{N}\cup\{\infty\}$, domain Ω, for every collection of disjoint bins $B_1, \dots, B_K$ defined on Ω, $n \in \mathbb{N}$, $\epsilon, \delta \in (0, 1/n)$, $\beta > 0$ and $\alpha \in (0,1)$ there exists an (ϵ, δ)-differentially private algorithm $M : \Omega^n \to \mathbb{R}^K$ such that for every distribution $\mathcal{D}$ on Ω, if*

1. $X_1, \dots, X_n \overset{\text{iid}}{\sim} \mathcal{D}, \quad p_k = \mathbb{P}(X_i \in B_k)$,
2. $(\tilde{p}_1, \dots, \tilde{p}_K) \leftarrow M(X_1, \dots, X_n)$, *and*
3. $$n \geq \max\left\{ \min\left\{ \frac{8}{\epsilon\beta}\log\left(\frac{2K}{\alpha}\right), \frac{8}{\epsilon\beta}\log\left(\frac{4}{\alpha\delta}\right)\right\}, \frac{1}{2\beta^2}\log\left(\frac{4}{\alpha}\right)\right\} \tag{1}$$

then,

$$\mathbb{P}_{\substack{X\sim\mathcal{D}\\ M}}(|\tilde{p}_k - p_k| \leq \beta) \qquad \geq 1-\alpha \tag{2}$$

$$\mathbb{P}_{\substack{X\sim\mathcal{D}\\ M}}\left(\arg\max_k \tilde{p}_k = j\right) \quad \leq \begin{cases} np_j + 2e^{-(\epsilon n/8)\cdot(\max_k p_k)} & \text{if } K < 2/\delta, \\ np_j & \text{if } K \geq 2/\delta. \end{cases} \tag{3}$$

where the probability is taken over the randomness of M and the data $X_1, \dots, X_n$.

3.2. DP estimation of the range of a Gaussian R.V. with known variance

We first consider the range-estimation problem when the standard deviation is known, or more generally when a valid upper bound on the standard deviation is known. This is an intermediate subroutine. In the main known-variance confidence interval theorem, the variance σ^2 is assumed to be known exactly, but the range estimator itself is slightly more general: it works as long as the true standard deviation is at most the supplied value σ.

The goal of this algorithm is to output an interval $[X_{\min}, X_{\max}]$ that contains all the samples with high probability, while keeping the interval length as small as possible. The algorithm first partitions the possible range of the mean into bins of width roughly σ. It then runs the private histogram learner and selects the bin with the largest noisy count. Since most of the probability mass of a Gaussian distribution lies near its mean, the heaviest bin should be close to μ. After selecting this bin, the algorithm expands it by a term of order

$$\sigma\sqrt{\log(n/\alpha)}.$$

This expansion is large enough to contain all n Gaussian samples with probability at least $1-\alpha$.

This range-estimation step is the key reason the final confidence interval does not depend on a worst-case bound on the data. Instead of requiring the analyst to provide a large deterministic bound B, the algorithm privately learns a data-dependent range whose length is comparable to the typical range of Gaussian samples.

Theorem 3.2. *For every $n \in \mathbb{N}$, $\sigma, \epsilon, \delta > 0$, $\alpha \in (0, 1/2)$, $R \in (0, \infty]$, there is a $w > 0$ and an (ϵ, δ)-differentially private algorithm $M : \mathbb{R}^n \to \mathbb{R} \times \mathbb{R}$ such that whenever $\mu \in (-R, R)$ and*

$$n \geq c \min\left\{\frac{1}{\epsilon}\log\left(\frac{R}{\sigma\alpha}\right), \frac{1}{\epsilon}\log\left(\frac{1}{\delta\alpha}\right)\right\},$$

(where c is a universal constant), if $X_1, \ldots, X_n$ are iid Gaussian random variables with mean μ and variance $\sigma_0^2 \leq \sigma^2$ (where σ^2 is known) and

$$(X_{\min}, X_{\max}) \leftarrow M(X_1, \ldots, X_n),$$

we have:

$$\mathop{\mathbb{P}}_{M}{}_{X\sim\mathcal{N}(\mu,\sigma_0^2)}\left(\forall i\ X_{\min} \leq X_i \leq X_{\max}\right) \geq 1-\alpha.$$

and with probability 1*,*

$$|X_{\max} - X_{\min}| = w = O\left(\sigma\sqrt{\log(n/\alpha)}\right).$$

3.3. DP estimation of variance

When the variance is unknown, the previous range estimator cannot be applied directly because it needs a scale parameter. Karwa and Vadhan therefore first construct a private constant-factor estimate of the standard deviation. This estimate does not need to be very precise. It only needs to be large enough so that the later range estimator does not underestimate the spread of the data.

The algorithm uses the fact that if

$$X_i \sim N(\mu, \sigma^2),$$

then differences of pairs remove the unknown mean:

$$X_{2i} - X_{2i-1} \sim N(0, 2\sigma^2).$$

Therefore, the absolute pairwise differences $|X_{2i} - X_{2i-1}|$ have typical size on the order of σ, independently of μ. The algorithm divides the positive real line into exponentially growing bins, such as

$$(2^j, 2^{j+1}].$$

It then runs the private histogram learner on these pairwise differences and selects the heaviest noisy bin. Because the bins grow geometrically, identifying the correct bin gives a constant-factor estimate of σ.

The output $\hat{\sigma}$ satisfies

$$\sigma \leq \hat{\sigma} \leq 8\sigma$$

with high probability. This one-sided behavior is important: overestimating σ only makes the final interval wider, but underestimating σ could cause the range estimator to miss some samples and break the coverage guarantee.

Theorem 3.3. *For every $n \in \mathbb{N}$, $\sigma_{\min} < \sigma_{\max} \in [0, \infty]$, $\epsilon, \delta > 0$, $\alpha \in (0, 1/2)$, there is an (ϵ, δ)-differentially private algorithm $M : \mathbb{R}^n \to [0, \infty)$ such that if $X_1, \ldots, X_n$ are iid Gaussian random variables with mean μ and with variance $\sigma^2 \in (\sigma_{\min}^2, \sigma_{\max}^2)$, and*

$$\hat{\sigma} \leftarrow M(X_1, \ldots, X_n),$$

we have:

1. *High probability bound: If*

$$n \geq c \min\left\{\frac{1}{\epsilon}\log\left(\frac{\log(\sigma_{\max}/\sigma_{\min})}{\alpha}\right), \frac{1}{\epsilon}\log\left(\frac{1}{\delta\alpha}\right)\right\},$$

(where c is a universal constant),

$$\mathop{\mathbb{P}}_{M}{}_{X\sim\mathcal{N}(\mu,\sigma^2)}\left(\sigma \leq \hat{\sigma} \leq 8\sigma\right) \geq 1-\alpha.$$

2. *Expectation bound: If*

$$n \geq c \min\left\{\frac{1}{\epsilon}\log\left(\frac{\sigma_{\max}}{\sigma_{\min}}\right), \frac{1}{\epsilon}\log\left(\frac{1}{\delta\alpha}\right)\right\},$$

then

$$\mathop{\mathbb{E}}_{M}{}_{X\sim\mathcal{N}(\mu,\sigma^2)}\left[\hat{\sigma}^2\right] \leq \sigma^2 \cdot \left(c_1 + c_2 \log^2(n) \cdot \alpha\right)$$

for some universal constants c_1 and c_2.

3.4. DP estimation of the range of a Gaussian R.V. (unknown variance)

The unknown-variance range estimator combines the previous two subroutines. First, the algorithm runs the private variance-estimation algorithm to obtain a constant-factor estimate $\hat{\sigma}$. Then it runs the known-variance range estimator using $\hat{\sigma}$ as the supplied scale parameter.

This works because, with high probability, $\hat{\sigma}$ is an upper bound on the true standard deviation up to a constant factor. Thus, the enlarged bin in the range-estimation step is wide enough to contain the Gaussian samples. At the same time, since $\hat{\sigma}$ is at most a constant multiple of σ, the final range is still of the correct order:

$$O\left(\sigma\sqrt{\log(n/\alpha)}\right).$$

The output of this algorithm is

$$(\hat{\sigma}, X_{\min}, X_{\max}).$$

Here $\hat{\sigma}$ is the private scale estimate, and $[X_{\min}, X_{\max}]$ is the private range used later for clipping. This algorithm is the main bridge from the unknown-variance setting to the known-variance clipping strategy.

Theorem 3.4. *For every $n \in \mathbb{N}$, $\sigma_{\max} > \sigma_{\min} \in [0, \infty]$, $\epsilon, \delta > 0$, $\alpha \in (0, 1/2)$, $R \in [0, \infty]$, there is an (ϵ, δ)-differentially private algorithm $M : \mathbb{R}^n \to (0, \infty) \times \mathbb{R} \times \mathbb{R}$ such that whenever*

$$n \geq c \cdot \min\left\{ \max\left\{ \frac{1}{\epsilon}\log\left(\frac{R}{\sigma_{\min}\alpha}\right), \frac{1}{\epsilon}\log\left(\frac{\log(\sigma_{\max}/\sigma_{\min})}{\alpha}\right)\right\}, \frac{1}{\epsilon}\log\left(\frac{1}{\delta\alpha}\right)\right\},$$

(where c is a universal constant), if $X_1, \ldots, X_n$ are iid Gaussian random variables with mean $\mu \in (-R, R)$ and with variance $\sigma^2 \in (\sigma^2_{\min}, \sigma^2_{\max})$, and

$$(\hat{\sigma}, X_{\min}, X_{\max}) \leftarrow M(X_1, \ldots, X_n),$$

we have:

1.
$$\mathop{\mathbb{P}}_{\substack{X\sim\mathcal{N}(\mu,\sigma^2)\\ M}}\left(\forall i\ X_{\min} \leq X_i \leq X_{\max}\right) \geq 1 - \alpha$$

2.
$$\mathop{\mathbb{P}}_{\substack{X\sim\mathcal{N}(\mu,\sigma^2)\\ M}}\left(|X_{\max} - X_{\min}| \leq O\Big(\sigma\sqrt{\log(n/\alpha)}\Big)\right) \geq 1-\alpha$$

3.
$$\mathop{\mathbb{P}}_{\substack{X\sim\mathcal{N}(\mu,\sigma^2)\\ M}}\left(\sigma \leq \hat{\sigma} \leq 8\sigma\right) \geq 1 - \alpha$$

Moreover, if

$$n \geq c \cdot \min\left\{\frac{1}{\epsilon}\log\left(\frac{\sigma_{\max}}{\sigma_{\min}}\right), \frac{1}{\epsilon}\log\left(\frac{1}{\delta\alpha}\right)\right\},$$

then

$$\mathop{\mathbb{E}}_{\substack{X\sim\mathcal{N}(\mu,\sigma^2)\\ M}}\left[\hat{\sigma}^2\right] \leq \sigma^2\left(c_1 + c_2\log^2(n)\cdot\alpha\right).$$

3.5. Main theorems (Upper Bound)

The main theorems of mean estimation can be split into two cases:

1. variance is known
2. variance is unknown

3.5.1. Variance is known

We now move from range estimation to confidence interval construction. In the known-variance setting, the variance σ^2 is known exactly, just as in the classical Gaussian confidence interval

$$\bar{X} \pm \frac{\sigma}{\sqrt{n}} z_{1-\alpha/2}.$$

The private algorithm follows the same basic structure, but it must first control sensitivity.

The algorithm has three main steps. First, it privately estimates a range $[X_{\min}, X_{\max}]$ containing the data with high probability. Second, it clips each sample to this interval:

$$Y_i = \begin{cases} X_i, & X_i \in [X_{\min}, X_{\max}], \\ X_{\max}, & X_i > X_{\max}, \\ X_{\min}, & X_i < X_{\min}. \end{cases}$$

After clipping, changing one data point can change the clipped empirical mean by at most

$$\frac{X_{\max} - X_{\min}}{n}.$$

Therefore, the Laplace mechanism can be applied to the clipped mean.

Finally, the algorithm outputs an interval centered at the noisy clipped mean. The radius of the interval contains two parts: the usual Gaussian sampling error and an additional term for the Laplace noise. The resulting interval has fixed width, and its length is bounded by

$$\max\left\{\frac{\sigma}{\sqrt{n}}O\left(\sqrt{\log(1/\alpha)}\right), \frac{\sigma}{\epsilon n}\operatorname{polylog}\left(\frac{n}{\alpha}\right)\right\}.$$

The first term matches the classical non-private rate, while the second term is the extra cost of privacy.

Theorem 3.5. *Let $\mathcal{I}$ be the set of all possible intervals in $\mathbb{R}$. For every $n \in \mathbb{N}$, $\sigma, \epsilon, \delta > 0$, $\alpha \in (0, 1/2)$, $R \in (0, \infty]$, there is $\beta > 0$ and an (ϵ, δ)-differentially private algorithm $M : \mathbb{R}^n \to \mathcal{I}$ such that if $X_1, \ldots, X_n$ are iid random variables from $\mathcal{N}(\mu, \sigma^2)$, with $\mu \in (-R, R)$ and $I \leftarrow M(X_1, \ldots, X_n)$, then*

$$\mathop{\mathbb{P}}_{\substack{X \sim \mathcal{N}(\mu, \sigma^2) \\ M}} \big(I(X_1, \ldots, X_n) \ni \mu\big) \geq 1 - \alpha.$$

Moreover, $|I| = \beta$ and if

$$n \geq c \min\left\{ \frac{1}{\epsilon} \log\left(\frac{R}{\sigma\alpha}\right), \frac{1}{\epsilon} \log\left(\frac{1}{\delta\alpha}\right) \right\},$$

(where c is a universal constant), then

$$\beta \leq \max\left\{ \frac{\sigma}{\sqrt{n}} O\left(\sqrt{\log\left(\frac{1}{\alpha}\right)}\right), \frac{\sigma}{\epsilon n} \operatorname{polylog}\left(\frac{n}{\alpha}\right) \right\}. \tag{4}$$

3.5.2. Variance is unknown

In the unknown-variance setting, the algorithm must also account for uncertainty in the variance. The non-private confidence interval uses the sample variance and a t-quantile rather than the known value of σ. The private algorithm follows the same idea, but it must estimate both the mean and the variance in a way that preserves differential privacy.

The algorithm first runs the unknown-variance range estimator to obtain a private interval $[X_{\min}, X_{\max}]$. It then clips the data to this interval and adds Laplace noise to the clipped empirical mean. In addition, it privately estimates the sample variance of the clipped data. Because the clipped data lie in a bounded interval, the clipped sample variance has bounded sensitivity, so Laplace noise can also be added to this quantity.

The final interval is centered at the noisy clipped mean and uses the private variance estimate together with a t-quantile. The algorithm deliberately makes the private variance estimate conservative, so that the interval remains valid even after adding privacy noise. Unlike the known-variance case, the interval length is random because it depends on the private variance estimate. Therefore, the theorem bounds the expected length of the interval.

The final expected length has the same general form as in the known-variance case:

$$\max\left\{ \frac{\sigma}{\sqrt{n}} O\left(\sqrt{\log(1/\alpha)}\right), \frac{\sigma}{\epsilon n} \operatorname{polylog}\left(\frac{n}{\alpha}\right) \right\}.$$

Again, the first term is the usual non-private statistical error, and the second term is the additional privacy cost.

Theorem 3.6. *Let $\mathcal{I}$ be the set of all possible intervals in $\mathbb{R}$. For every $n \in \mathbb{N}$, $\sigma_{\min} < \sigma_{\max} \in [0, \infty]$, $\epsilon, \delta > 0$, $\alpha \in (0, 1/2)$, $R \in [0, \infty)$, there is an (ϵ, δ)-differentially private algorithm $M : \mathbb{R}^n \to \mathcal{I}$ such that if $X_1, \ldots, X_n$ are iid random variables from $\mathcal{N}(\mu, \sigma^2)$, where $\mu \in (-R, R)$ and $\sigma \in (\sigma_{\min}, \sigma_{\max})$, and $I \leftarrow M(X_1, \ldots, X_n)$, then*

$$\mathop{\mathbb{P}}_{\substack{X \sim \mathcal{N}(\mu, \sigma^2) \\ M}} \big(I(X_1, \ldots, X_n) \ni \mu\big) \geq 1 - \alpha.$$

Moreover, if

$$n \geq \frac{c_1}{\epsilon} \min\Bigg\{ \max\left\{ \log\left(\frac{R}{\sigma_{\min}}\right), \log\left(\frac{\sigma_{\max}}{\sigma_{\min}}\right) \right\}, \log\left(\frac{1}{\delta}\right) \Bigg\} + \frac{c_2}{\epsilon} \log\left(\frac{\log(1/\epsilon)}{\alpha}\right).$$

(where c_1 and c_2 are universal constants), then

$$\begin{aligned} \beta &:= \mathop{\mathbb{E}}_{\substack{X \sim \mathcal{N}(\mu, \sigma^2) \\ M}} \big[\, |I(X_1, \ldots, X_n)| \,\big] \\ &\leq \max\left\{ \frac{\sigma}{\sqrt{n}} O\left(\sqrt{\log\left(\frac{1}{\alpha}\right)}\right), \frac{\sigma}{\epsilon n} \operatorname{polylog}\left(\frac{1}{\alpha}\right) \right\}. \end{aligned}$$

3.6. Main theorems (Lower Bound)

They first proved a theorem without privacy.

Theorem 3.7. *Let $X_1, \ldots, X_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$ where $\mu \in (-R, R)$, σ^2 is known, and $R > c\sigma^2$ for a fixed (but arbitrarily small) constant $c > 0$. Let $I = I(X_1, \ldots, X_n)$ be any measurable set where $I \subseteq (-R, R)$ such that for every $\alpha \in (0, 1/2)$, and $\mu \in (-R, R)$, if*

$$\mathbb{P}_{\underline{X} \sim N(\mu, \sigma^2)} \big(I(X_1, \ldots, X_n) \ni \mu\big) \geq 1 - \alpha$$

,

$$\mathbb{E}_{\underline{X} \sim N(\mu, \sigma^2)} \big[|I(X_1, \ldots, X_n)|\big] \geq \frac{2\sigma}{\sqrt{n}} z_{1 - \frac{\alpha}{2}} - \tilde{O}\left(\frac{1}{n}\right).$$

Since $z_{1-\frac{\alpha}{2}} = \Theta(\sqrt{\log(1/\alpha)})$, the lower bound for statistical error is then $\Omega(\sigma\sqrt{\frac{\log(1/\alpha)}{n}})$, which matches the first term in the upper bound (4).

Then they showed the following lower bound for privacy error.

Theorem 3.8. *Let $M(X_1, \ldots, X_n)$ be any (ϵ, δ)-DP algorithm that outputs a random measurable set $S(X_1, \ldots, X_n) \subset (-R, R)$ such that, for every $\alpha \in (0, 1)$, and $\mu \in (-R, R)$, whenever $X_1, \ldots, X_n \overset{\text{iid}}{\sim} N(\mu, \sigma^2)$ where $\mu \in (-R, R)$ and σ^2 is known, if*

1.

$$\mathbb{P}^M_{\underline{X}\sim N(\mu,\sigma^2)}\big(S(X_1,\ldots,X_n)\ni\mu\big)\geq 1-\alpha.$$

2.

$$\mathbb{E}^M_{\underline{X}\sim N(\mu,\sigma^2)}\big[|S(X_1,\ldots,X_n)|\big]\leq\beta$$

then,

1. *For all $\alpha\in(0,\frac{1}{2})$, and $\delta<\alpha/2n$,*

$$\beta\geq c\cdot\min\left(\frac{\sigma}{\epsilon n}\log\left(\frac{1}{\alpha}\right),\,R\right). \tag{5}$$

2. *If $\beta<\sigma<R$, then,*

$$n\geq c_1\cdot\min\left(\frac{1}{\epsilon}\log\left(\frac{1}{\alpha}\right),\,\frac{1}{\epsilon}\log\left(\frac{1}{\delta}\right)\right)$$

and

$$n\geq c_2\cdot\min\left(\frac{1}{\epsilon}\log\left(\frac{R}{\sigma}\right),\,\frac{1}{\epsilon}\log\left(\frac{1}{\delta}\right)\right).$$

Remark: Suppose $\beta<R$, then this lower bound matches the second term in upper bound (4) up to some polylog terms.

4. Future Directions

Karwa and Vadhan's work suggests several natural directions for future research. The first direction is to close the remaining gap between the upper and lower bounds. In the known-variance case, the lower bound shows that any differentially private confidence interval must pay a privacy cost of order

$$\Omega\left(\frac{\sigma}{\epsilon n}\log\left(\frac{1}{\alpha}\right)\right),$$

while their upper bound has a privacy term of the form

$$\frac{\sigma}{\epsilon n}\mathrm{polylog}\left(\frac{n}{\alpha}\right).$$

Thus, the leading statistical term already matches the classical non-private rate, but there remains a polylogarithmic gap in the privacy-dependent term. A sharper algorithm or lower bound would clarify whether this gap is only an artifact of the analysis or is genuinely necessary.

A second direction is to improve the practical performance of the algorithm. Although the algorithm has strong finite-sample guarantees, it can be conservative in practice because it first privately estimates a range and then adds noise calibrated to the clipped sensitivity. Later work such as (Du et al., 2020) studies more practical differentially private confidence intervals for Gaussian means and compares several algorithms empirically. This line of work asks whether one can preserve the finite-sample validity of Karwa and Vadhan's approach while producing intervals closer to the optimal non-private interval length.

A third direction is to extend the central idea of the paper—private localization, clipping, and noisy estimation—to other statistical tasks. For example, (Kamath et al., 2020) studies private mean estimation for heavy-tailed distributions, where one must again control the influence of extreme observations before adding privacy noise. Similarly, (Biswas et al., 2020) develops CoinPress, a practical private estimation framework for multivariate sub-Gaussian mean and covariance estimation. These works suggest that the localization-and-clipping idea from (Karwa & Vadhan, 2018) is not only useful for one-dimensional Gaussian confidence intervals, but also forms a general strategy for private estimation with unbounded or weakly bounded data.

Finally, another broad direction is to develop general-purpose methods for valid inference after privatization. Instead of designing a separate confidence interval for each statistic, later work has studied bootstrap and simulation-based approaches for differentially private confidence intervals. For instance, (Ferrando et al., 2022) develops a parametric bootstrap approach, while (Chadha et al., 2024) studies private resampling methods for confidence intervals in settings such as mean estimation, median estimation, and logistic regression. These works move beyond the specific Gaussian setting and aim to build a more general theory of statistically valid inference under differential privacy.

References

Biswas, S., Dong, Y., Kamath, G., and Ullman, J. Coinpress: Practical private mean and covariance estimation. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 14475–14485. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/a684eceee76fc522773286a895bc8436-Paper.pdf.

Bun, M., Nissim, K., and Stemmer, U. Simultaneous private learning of multiple concepts. *Journal of Machine Learning Research*, 20(94):1–34, 2019. URL http://jmlr.org/papers/v20/18-549.html.

Cai, B., Daskalakis, C., and Kamath, G. Priv'IT: Private and sample efficient identity testing. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 635–644. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/cai17a.html.

Chadha, K., Duchi, J., and Kuditipudi, R. Resampling methods for private statistical inference. *arXiv preprint arXiv:2402.07131*, 2024.

Du, W., Foot, C., Moniot, M., Bray, A., and Groce, A. Differentially private confidence intervals. *arXiv preprint arXiv:2001.02285*, 2020.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pp. 265–284. Springer, 2006. doi: 10.1007/11681878_14.

Ferrando, C., Wang, S., and Sheldon, D. Parametric bootstrap for differentially private confidence intervals. In Camps-Valls, G., Ruiz, F. J. R., and Valera, I. (eds.), *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 1598–1618. PMLR, 28–30 Mar 2022. URL https://proceedings.mlr.press/v151/ferrando22a.html.

Gaboardi, M., Lim, H., Rogers, R., and Vadhan, S. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2111–2120, New York, New York, USA, 20–22 Jun 2016. PMLR. URL https://proceedings.mlr.press/v48/rogers16.html.

Kamath, G., Singhal, V., and Ullman, J. Private mean estimation of heavy-tailed distributions. In Abernethy, J. and Agarwal, S. (eds.), *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 2204–2235. PMLR, 09–12 Jul 2020. URL https://proceedings.mlr.press/v125/kamath20a.html.

Karwa, V. and Vadhan, S. Finite sample differentially private confidence intervals. pp. 44:1–44:9, Dagstuhl, Germany, 2018. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik., 2018. ITCS. URL http://drops.dagstuhl.de/opus/volltexte/2018/8344/.

Lehmann, E. L. and Romano, J. P. *Testing statistical hypotheses*. Springer Texts in Statistics. Springer, New York, third edition, 2005. ISBN 0-387-98864-5.

Sheffet, O. Differentially private least squares: Estimation, confidence and rejecting the null hypothesis. *CoRR*, abs/1507.02482, 2015. URL http://arxiv.org/abs/1507.02482.

Vadhan, S. *The Complexity of Differential Privacy*, pp. 347–450. Springer International Publishing, Cham, 2017. ISBN 978-3-319-57048-8. doi: 10.1007/978-3-319-57048-8_7. URL https://doi.org/10.1007/978-3-319-57048-8_7.