Private Confidence Sets

Karan Chadha
Stanford University
knchadha@stanford.edu

John Duchi
Stanford University
jduchi@stanford.edu

Rohith Kuditipudi
Stanford University
rohithk@stanford.edu

Abstract

We consider statistical inference under privacy constraints. In particular, we give differentially private algorithms for estimating coverage probabilities and computing valid confidence sets, and prove upper bounds on the error of our estimates and the length of our confidence sets. Our bounds apply to broad classes of data distributions and statistics of interest, and for fixed ε we match the higher-order asymptotic accuracy of the standard (non-private) non-parametric bootstrap.

1 Introduction

The goal of statistical machine learning is to make inferences on parameters of the underlying population from which observed data has been drawn. Obtaining confidence sets for a statistic of interest is among the most basic of inferential tasks. Particularly in safety-critical applications such as healthcare, it is essential to be able to rigorously and accurately quantify uncertainties around point estimates or predictions.

To the extent that such applications often involve sensitive data, it is also crucial that any inferential procedure not compromise the privacy of the individuals from whom the data has been collected. Over the last decade plus, differential privacy [3] has emerged as the *de facto* framework for specifying such privacy constraints.

In this work, we aim to compute confidence sets for statistics of general (e.g., non-Gaussian) distributions while maintaining (pure) differential privacy. Building on the simplicity, generality and accuracy—at least in the absence of privacy considerations—of simulation-based inferential procedures such as the bootstrap, we propose private mechanisms for estimating the coverage of user-specified confidence sets, as well as computing valid confidence sets given a user-specified target coverage probability.

Notably, for fixed privacy parameter $\varepsilon = \Omega(1)$, our estimated coverage probabilities match the accuracy of the standard non-private nonparametric percentile-t bootstrap, which itself is asymptotically optimal, i.e. achieves the Cramér-Rao bound.

1.1 Preliminaries

Let $\mathcal{P}$ be a collection of distributions on a set $\mathcal{X} \subseteq \mathbb{R}$, and consider a statistic $\theta : \mathcal{P} \to \mathbb{R}$ of interest. We have an estimator, typically the plug-in estimator, $\widehat{\theta}_n : \mathcal{X}^n \to \mathbb{R}$ of $\theta(P)$, which converges at a rate r_n to $\theta(P)$, that is,

$$r_n(\theta(P) - \widehat{\theta}_n) \xrightarrow{d} T \tag{1}$$

for some random variable T.

We aim to estimate the cumulative distribution function (CDF) of the estimator, that is, for $t \in \mathbb{R}$ to estimate

$$J_n(t, P) := \mathbb{P}(r_n(\theta(P) - \widehat{\theta}_n) \leq t), \tag{2}$$

[1] For a set $\mathcal{S} \subseteq \mathbb{R}$, we use the notation $J_n(\mathcal{S}, P) := \mathbb{P}(r_n(\theta(P) - \widehat{\theta}_n) \in \mathcal{S})$.

35th Conference on Neural Information Processing Systems (NeurIPS 2021), Sydney, Australia.

so that we may ultimately compute valid confidence sets, defined as follows.

Definition 1.1. *(adapted from Definition 2 in [9]). Let $\alpha \in (0,1)$ and $X_{1:n} \overset{i.i.d.}{\sim} P$. A valid $(1-\alpha)$-level confidence set*[2] *for a statistic $\theta(P)$ is a (possibly randomized) function $C_\alpha : \mathcal{X}^n \to \mathcal{B}(\mathbb{R})$ such that for all $P \in \mathcal{P}$ we have*

$$\mathbb{P}\left(\theta(P) \in C_\alpha(X)\right) \geq 1 - \alpha,$$

where the probability is taken over the randomness of both C and P.

We consider C_α that satisfy differential privacy.

Definition 1.2. *(Differential privacy) For $\varepsilon, \delta \geq 0$, a mechanism $M : \mathcal{X}^n \to \mathcal{O}$ is (ε, δ) differentially private if for any measurable $O \in \mathcal{O}$ and $X, X' \in \mathcal{X}^n$ that differ in at most one element we have*

$$\mathbb{P}\left(M(X) \in O\right) \leq e^\varepsilon \cdot \mathbb{P}\left(M(X') \in O\right) + \delta.$$[3]

The results given in Section 2 are contingent on various (mild) regularity conditions on the data distribution and estimator under consideration, which are standard in the classical bootstrap literature. See Hall [7] for a canonical reference. Owing to space constraints, we defer stating these assumptions explicitly and instead present the results non-rigorously.

1.2 Related work

Whereas the principal focus of much early work on differential privacy was on on the design and analysis of algorithms for privately querying data and computing sample statistics, in recent years a growing body of work has sought to marry differential privacy and statistical inference—that is, to design differentially private procedures for inferring population parameters from sample statistics, with both the usual distribution-free privacy guarantees and also statistical guarantees under distributional assumptions on the given data.

Notably, Smith [11] showed that a large class of (asymptotically normal) statistical estimators admit private counterparts that asymptotically converge to the same Gaussian distribution. Perhaps motivated by this phenomenon, much of the emphasis of existing work on differentially private confidence sets has been on private mechanisms for covariance estimation—e.g., for univariate Gaussian [9], multivariate Gaussian [8] and sub-Gaussian data [2]; for private empirical risk minimization [1, 12]; and under the local model [6]—so that the normal approximation may then be applied to estimate coverage probabilities. Karwa and Vadhan [9] give matching upper and lower bounds, up to logarithmic factors, on the length of valid confidence sets for Gaussian mean estimation.

For non-Gaussian statistics, however, the normal approximation is merely first-order accurate. Indeed, it is well known in the non-private setting that data-driven, simulation-based inferential procedures such as the bootstrap are generally asymptotically more accurate than the normal approximation, i.e. *higher-order accurate*, in estimating the coverage of confidence sets, under suitable regularity conditions. Ferrando et al. [5] propose mechanisms for privatizing the parametric bootstrap, but do not provide quantitative error bounds. Our proposed mechanisms are similar to the frameworks proposed by Evans et al. [4], though we again distinguish ourselves via our emphasis on quantitative upper bounds on the error of our procedures.

2 Main Results

As a point of departure, consider the setting of mean estimation, with $\widehat{\theta}_n(X_{1:n}) := \overline{X}$. Algorithm 1 prescribes a natural first attempt at privatizing the standard nonparametric bootstrap, wherein each

[2] We will sometimes also use the term "confidence set" to refer to the output of C_α.

[3] We sometimes will refer to a mechanism as ε *differentially private* for $\delta = 0$.

time we compute our statistic on resampled data we add suitable noise to ensure the privacy of the statistic, assuming the domain $\mathcal{X}$ is bounded by an interval of width R,

Algorithm 1: Private bootstrap

Input : data $X_{1:n}$, threshold t
Output : estimate of $J_n(t, P)$

$T \leftarrow \lfloor \sqrt{n\varepsilon} \rfloor$, $\tilde{\theta} \leftarrow \widehat{\theta}_n(X_{1:n}) + \xi$ for $\xi \sim \text{Lap}\left(0, \frac{R}{n\varepsilon}\right)$
for $i \in 1, ..., T$ **do**
  Sample $X'_1, ..., X'_n$ uniformly with replacement from $X_{1:n}$
  $\tilde{\theta}_i \leftarrow \widehat{\theta}_n(X'_{1:n}) + \xi_i$ for $\xi_i \sim \text{Lap}\left(0, \frac{R}{n^{3/4}\varepsilon}\right)$
end
$\widehat{J}_n(t, P) \leftarrow \sum_{i=1}^{T} 1\left\{\sqrt{n}\left(\tilde{\theta} - \tilde{\theta}_i\right) \le t\right\}$
return $\widehat{J}_n(t, P)$

2.1 Upper bounds for estimating CDFs

Together with its simplicity and general applicability, the non-private bootstrap is typically more accurate than the normal approximation in estimating $J_n(t, P)$, under suitable regularity conditions on the distribution P. However, the usual pecking order flips under privacy constraints; whereas the normal approximation can be privatized in a manner that yields $\tilde{O}(\varepsilon^{-1}n^{-1/2})$ error [9], Algorithm 1 achieves a corresponding rate of just $O(\varepsilon^{-1/2}n^{-1/4})$.

Proposition 1. *(informal) For any $\delta > 0$, Algorithm 2 is $(\varepsilon_\delta, \delta)$ differentially private for $\varepsilon_\delta := \tilde{O}(\varepsilon)$,[4] and estimates $\widehat{J}_n(t, P)$ with $O_P(\varepsilon^{-1/2}n^{-1/4})$ error.*

Motivated by the failure of Algorithm 1 to match even the normal approximation, we propose Algorithm 2, a straightforward privatized variant of the bag of little bootstraps procedure [10], as a private mechanism for aggregating the results of multiple bootstrap procedures run on partitions of the data to estimate $J_n(t, P)$ to within $O(\varepsilon^{-1/2}n^{-1/2})$ error.

Algorithm 2: Private bootstrap ensemble

Input : data $X_{1:n}$, threshold t, ensemble size s
Output : estimate of $J_n(t, P)$

$b \leftarrow \lfloor n/s \rfloor$, $X_i \leftarrow X_{\pi(i)}$ for random permutation π of $[n]$ (i.e. shuffle the data)
for $k \in 1, ..., s$ **do**
  for $i \in 1, ..., T$ **do**
    Sample $X_1^i, ..., X_n^i$ uniformly with replacement from $X_{(k-1)b:kb}$
  end
  $\widehat{J}_n^k(t, P) \leftarrow \sum_{i=1}^{T} 1\left\{\sqrt{n}\left(\widehat{\theta}_n(X_{(k-1)b:kb}) - \widehat{\theta}_n(X^i_{1:n})\right) \le t\right\}$
end
$\widehat{J}_n(t, P) \leftarrow \frac{1}{s}\sum_{k=1}^{s} \widehat{J}_n^k(t, P) + \xi$ for $\xi \sim \text{Lap}(0, 1/s\varepsilon)$
return $\widehat{J}_n(t, P)$

Theorem 1. *(informal) For $\varepsilon > 1/\sqrt{n}$, setting $s = \sqrt{\frac{n}{\varepsilon}}$ we have that Algorithm 2 is ε differentially private and estimates $J_n(t, P)$ with $O_P(\varepsilon^{-1/2}n^{-1/2})$ error.*

Note that Algorithm 2 achieves lower error than Algorithm 1 *while also satisfying a more stringent privacy constraint*, i.e. pure ε differential privacy. But while Algorithm 2 does better than Algorithm 1 in estimating $J_n(t, P)$ for a single fixed query t, one apparent advantage retained by Algorithm 1 is that we can reuse the resampled statistics to release the full histogram of the plug-in estimator while incurring no further privacy loss.

[4] We have hidden the usual logarithmic factors in δ, n and ε that result from advanced composition

Fortunately however, drawing inspiration from standard techniques for private histogram estimation, we can obtain a similar uniform approximation result from Algorithm 2 while maintaining the same error rate up to logarithmic factors.

Theorem 2. *(informal) We can construct an ε differentially private estimate $\widehat{J}_n(\cdot, P)$ such that $||\widehat{J}_n(\cdot, P) - J_n(\cdot, P)||_\infty = \tilde{O}_P(\varepsilon^{-1/2} n^{-1/2})$.*

While we have thus far limited the scope of our discussion to mean estimation, analogous results for Algorithm 2 can be shown for general statistics under suitable regularity conditions, including statistics which converge to a non-Gaussian random variable T, in which case the error in estimating $J_n(t, P)$ is typically determined by the rate of convergence, where we may have $r_n \neq n^{-1/2}$.

Finally, turning our attention to the studentized sample mean, define

$$J_n^{\text{st}} :=:= \mathbb{P}\left(\frac{\sqrt{n}(\widehat{\theta}_n - \theta(P))}{\widehat{\sigma}} \leq t \right), \tag{3}$$

where $\widehat{\sigma}$ is the usual unbiased empirical estimate of the standard deviation of P.

Whereas the error of the (non-private) normal approximation in estimating J_n^{st} remains $\Omega(n^{-1/2})$, using a minor variation of Algorithm 2, wherein we truncate the s independent bootstrap estimates around their median before averaging we can achieve a significantly improved dependence on the sample size n, matching the rate of the standard non-private bootstrap.[5]

Theorem 3. *(informal) We can estimate $J_n^{\text{st}}(t, P)$ with ε differential privacy and $O_P(1/\varepsilon n)$ error.*

Studentization also allows us to improve upon the uniform bound in Theorem 2.

Theorem 4. *(informal) We can construct an ε differentially private estimate $\widehat{J}_n^{\text{st}}(\cdot, P)$ of $J_n^{\text{st}}(\cdot, P)$ such that $||\widehat{J}_n^{\text{st}}(\cdot, P) - J_n^{\text{st}}(\cdot, P)||_\infty = \tilde{O}(1/\varepsilon^{1/2} n^{3/4})$.*

2.2 From CDFs to confidence sets

In computing a confidence set, one generally aims to minimize the expected size of the set while maintaining the desired coverage probability, i.e. validity. Assuming certain regularity conditions, we can obtain such a confidence set given a uniformly good estimate of $J_n(\cdot, P)$ such as described in the statement of Theorem 2, thus further motivating the results in Section 2.

Theorem 5. *(informal) Given $\widehat{J}_n(\cdot, P)$ satisfying $||\widehat{J}_n(\cdot, P) - J_n(\cdot, P)||_\infty \leq \xi$ and the existence of $\mathcal{S} \subseteq \mathbb{R}$ with $J_n(\mathcal{S}, P) \geq 1 - \alpha$, we can return $\mathcal{S}' \subseteq \mathbb{R}$ such that $|\mathcal{S}'| \leq |\mathcal{S}| + \sigma\xi$ and $J_n^f(\mathcal{S}', P) \geq 1 - \alpha$.*

For $f : \mathbb{R} \to \mathbb{R}$—e.g., f could be a private mechanism for releasing $\widehat{\theta}_n$[6]—let

$$J_n^f(t, P) := \mathbb{P}(r_n(f(\widehat{\theta}_n - \theta(P)) \leq t). \tag{4}$$

We can convert any estimate $\widehat{J}_n$ of J_n into an estimate of J_n^f by simulating a random variable with CDF $\widehat{J}_n$ and applying f to repeated draws from the simulation. The following result bounds the error of such an estimate.

Proposition 2. *(informal) Given $\widehat{J}_n(\cdot, P)$, we can construct an estimate $\widehat{J}_n^f(\cdot, P)$ of $J_n^f(\cdot, P)$ such that $||\widehat{J}_n^f(\cdot, P) - J_n^f(\cdot, P)||_\infty \leq ||\widehat{J}_n(\cdot, P) - J_n(\cdot, P)||_\infty$.*

After normalizing by the $\sqrt{n}$ scaling, Proposition 2 and Theorem 5 together imply for mean estimation that if there exists $\mathcal{S}$ such that $J_n^f(\mathcal{S}, P) \geq 1 - \alpha$, then we can construct, from the output of Algorithm 2, an interval around $f(\widehat{\theta}_n)$ of width at most $|\mathcal{S}| + O(\sigma/\varepsilon^{1/2} n)$ such that this interval is guaranteed to be a valid $1 - \alpha$ confidence set.

Taking f as an ε differentially private mechanism for releasing $\widehat{\theta}_n$, we thus obtain a 2ε differentially private $1 - \alpha$ confidence set.

[5] In particular, the non-parametric bootstrap achieves $O(1/\sqrt{n})$ error without studentization and $O(1/n)$ error with studentization.

[6] Note that restricting the domain of f to be $\mathbb{R}$ is a strong assumption, as generally private mechanisms can be functions of the data itself, i.e. $X_{1:n}$. However, many common mechanisms such as the Laplace mechanism do satisfy this assumption.

References

[1] M. Avella-Medina, C. Bradshaw, and P.-L. Loh. Differentially private inference via noisy optimization. *arXiv:2103.11003 [math.ST]*, 2021.

[2] S. Biswas, Y. Dong, G. Kamath, and J. Ullman. Coinpress: Practical private mean and covariance estimation. In *Advances in Neural Information Processing Systems 20*, 2020.

[3] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Theory of Cryptography Conference*, pages 265–284, 2006.

[4] G. Evans, G. King, M. Schwenzfeier, and A. Thakurta. Statistically valid inferences from privacy protected data. 2021.

[5] C. Ferrando, S. Wang, and D. Sheldon. General-purpose differentially-private confidence intervals. *arXiv:2006.07749 [cs.LG]*, 2020.

[6] M. Gaboardi, R. Rogers, and O. Sheffet. Locally private mean estimation: Z-test and tight confidence intervals. *arXiv:1810.08054 [cs.DS]*, 2018.

[7] P. Hall. *The Bootstrap and Edgeworth Expansion*. Springer, 1992.

[8] G. Kamath, J. Li, V. Singhal, and J. R. Ullman. Privately learning high-dimensional distributions. *arXiv:1805.00216 [cs.DS]*, 2018.

[9] V. Karwa and S. Vadhan. Finite sample differentially private confidence intervals. *arXiv:1711.03908 [cs.CR]*, 2017.

[10] A. Kleiner, A. Talwalkar, P. Sarkar, and M. I. Jordan. A scalable bootstrap for massive data. *Journal of the Royal Statistical Society, Series B*, 76(4):795–816, 2014.

[11] A. Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Forty-Third Annual ACM Symposium on the Theory of Computing*, pages 813–822. ACM, 2011.

[12] Y. Wang, D. Kifer, and J. Lee. Differentially private confidence intervals for empirical risk minimization. *arXiv:1804.03794 [cs.LG]*, 2018.