Dimension-free Private Mean Estimation for Anisotropic Gaussians

Yuval Dagan, Michael I. Jordan, Xuelin Yang, Lydia Zakynthinou, Nikita Zhivotovskiy

Differential Privacy: Algorithm $A: \mathcal{X}^n \to \mathcal{W}$ is (ε, δ)-differentially private if, for any two data sets X, X' that differ in exactly one data point, for any measurable subset, $W \in \mathcal{W}$,

$$\Pr_A[A(X) \in W] \le e^{\varepsilon} \Pr_A[A(X') \in W] + \delta$$

Challenge: Differentially private inference often suffers from a *curse of dimensionality:* $n = \Omega(d)$ for number of sample n and data dimension d.

Our focus: High-dimensional mean estimation:
Given dataset $x \in \mathbb{R}^{n \times d}$ sampled from distribution $\mathcal{P}$ with mean μ, find $\hat{\mu}$: $\|\hat{\mu} - \mu\|_2 \le \alpha$. This is a fundamental task and a subroutine in many private algorithms.

Dependence on d: Sometimes it is unavoidable:
For isotropic Gaussians $\mathcal{N}(\mu, \sigma^2 \mathbb{I}_{d \times d})$,

$$n \gtrsim \frac{d\sigma^2}{\alpha^2} + \frac{d\sigma}{\alpha\varepsilon} + \frac{\log(1/\delta)}{\varepsilon}$$

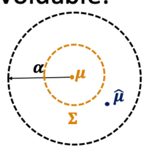

samples are necessary and sufficient for (ε, δ)-DP mean estimation [KLSU'19, BGSUZ'21].

But real-world data are often anisotropic (far from isotropic). The signal can be concentrated in few directions – the rest are noise.
Non-privately, the sample complexity scales with the *effective rank:*

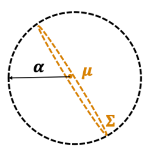

$$n \gtrsim \frac{\mathrm{Tr}(\Sigma)}{\alpha^2} = \frac{\sum_{i=1}^d \sigma_i^2}{\alpha^2} \Longrightarrow \| \mu_X - \mu \|_2 \le \alpha$$

Can we achieve similar dimension-independent sample complexity under privacy?

For pure DP, it's impossible via *packing [HT10]*.
For (ε, δ)-DP, we present:

- An (ε, δ)-DP (sub)Gaussian mean estimator with **dimension-free** sample complexity, for known Σ.
- The error achieved by this estimator is **optimal**.
- An (ε, δ)-DP (sub)Gaussian mean estimator with **milder dependence on the dimension** for unknown Σ.

Prior Approaches

1. Learning the mean in Mahalanobis norm
   - ✓ Implies bound on error in Euclidean norm: $\|\hat{\mu} - \mu\|_\Sigma \le \alpha \Rightarrow \|\hat{\mu} - \mu\|_2 \le \alpha\sigma_1$.
   - ✓ Can be done with unknown Σ with the same sample size [BGSUZ'21].
   - ✗ Requires $n = \Omega(d)$.
2. "Folklore": clip + spherical noise [KV'18]
   Find private coarse mean $\tilde{\mu}$ to clip data
   Add noise $\mathcal{N}(0, \mathrm{Tr}(\Sigma)\mathbb{I}_{d \times d}/\varepsilon^2 n^2)$

   $$n \gtrsim \frac{\mathrm{Tr}(\Sigma)}{\alpha^2} + \frac{\sqrt{d\ \mathrm{Tr}(\Sigma)}\sqrt{\log 1/\delta}}{\alpha\varepsilon} + \frac{\sqrt{d}}{\varepsilon}$$

   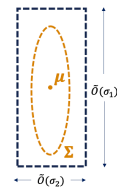

   - ✗ Spherical noise incurs $\sqrt{d\ \mathrm{Tr}(\Sigma)}/\varepsilon n$ error.
3. [PLAN'23]: clip + **rescaled** noise
   Find private coarse mean $\tilde{\mu}$ to clip data
   Add noise $\mathcal{N}\left(0, \mathrm{Tr}\left(\Sigma^{1/2}\right)\Sigma^{1/2}/\varepsilon^2 n^2\right)$

   $$n \gtrsim \frac{\mathrm{Tr}(\Sigma)}{\alpha^2} + \frac{\mathrm{Tr}\left(\Sigma^{1/2}\right)\sqrt{\log 1/\delta}}{\alpha\varepsilon} + \frac{\sqrt{d}}{\varepsilon}$$

   - ✓ Allows for more error in directions of small variance.
   - ✗ Coarse estimation requires $\sqrt{d}/\varepsilon$

All prior approaches require $n \gtrsim \sqrt{d}$. Ours do not!

An optimal estimator under known Σ

There exists an (ε, δ)-DP estimator which given n samples from $\mathcal{N}(\mu, \Sigma)$ with known Σ, returns $\hat{\mu}$ s.t. $\| \hat{\mu} - \mu \|_2 \le \alpha$ if

$$n \gtrsim \frac{\mathrm{Tr}(\Sigma)}{\alpha^2} + \frac{\mathrm{Tr}(\Sigma^{1/2})\sqrt{\log 1/\delta}}{\alpha\varepsilon} + \frac{\log 1/\delta}{\varepsilon}$$

Idea: replace coarse estimation with a pre-processing check
which only uses $\frac{\log 1/\delta}{\varepsilon}$ samples and ensures

$$\forall \ell, j : \left\|\Sigma^{-1/4}\left(x_\ell - x_j\right)\right\|_2^2 \le \mathrm{Tr}\left(\Sigma^{1/2}\right)$$

using the FriendlyCore filter of [TCKMS'21] with this predicate.

1. Filter data by the predicate and abort if too few remain.
2. Release empirical mean of modified dataset with Gaussian noise $\mathcal{N}\left(0, \frac{\log(1/\delta)\mathrm{Tr}(\Sigma^{1/2})}{\varepsilon^2 n^2}\right)\Sigma^{1/4}$

Optimality via *fingerprinting [BUV14]*:
Any (ε, δ)-DP algorithm for the Gaussian mean with accuracy α requires $n = \Omega\left(\frac{\sum_{i \in [d]} \sigma_i}{\alpha\varepsilon \log^2 d}\right)$.

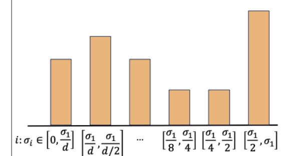

Must $\exists$ bucket of coordinates $m \in [\log d]$ with same variance $\sigma_{(m)}$ of size $d_{(m)} \propto \sum \sigma_i / \sigma_{(m)}$
$\Rightarrow$ bucket m is isotropic Gaussian
$\Rightarrow$ need $d_{(m)}\sigma_{(m)}/\alpha\varepsilon$ samples to learn its mean [KLSU2019]

A $d^{1/4}$ependence for unknown Σ

- Privately learning Σ in spectral norm requires $n \gtrsim d^{1.5}$ [KMS'22].
- Techniques from [BGSUZ'21] require $n \gtrsim d$.

Consider diagonal covariance.

1. Approach 1: Privately learn all σ_i, then run known-covariance algorithm $\Rightarrow$ Requires $n \gtrsim \sqrt{d}/\varepsilon$.
2. Approach 2: Only privately learn $\mathrm{Tr}(\Sigma)$ and add spherical Gaussian noise $\Rightarrow$ Error scales with $\sqrt{d\mathrm{Tr}(\Sigma)}/\varepsilon n$.

Idea: learn as many large σ_i as the sample size allows, add spherical noise to the remaining coordinates.
We learn the top $k \approx \varepsilon^2 n^2$ variances. Error of remaining coordinates is $\frac{\sqrt{d}}{\sqrt{k}}$ times larger than the optimal.

There exists an (ε, δ)-DP estimator which given n samples from $\mathcal{N}(\mu, \Sigma)$ with unknown diagonal Σ, returns $\hat{\mu}$ s.t. $\| \hat{\mu} - \mu \|_2 \le \alpha$ if

$$n \gtrsim \frac{\mathrm{Tr}(\Sigma)}{\alpha^2} + \frac{\mathrm{Tr}(\Sigma^{1/2})\sqrt{\log 1/\delta}}{\alpha\varepsilon} + \frac{d^{1/4}\sqrt{\mathrm{Tr}(\Sigma^{1/2})}\log 1/\delta}{\sqrt{\alpha}\varepsilon}$$

Optimal error under unknown covariance? Our algorithm matches the known-covariance error for special cases: e.g. exponentially decaying variances.

Aumüller, Lebeda, Nelson, Pagh (2023). PLAN:Variance-aware Private Mean Estimation.
Tsfadia Cohen Kaplan Mansour Stemmer (2021). FriendlyCore: Practical Differentially Private Aggregation
Brown, Gaboardi, Smith, Ullman, Zakynthinou (2021). Covariance-aware Private Mean Estimation without Private Covariance Estimation.

Paper:

Partially funded by the European Union (ERC-2022-SYG-OCEAN-101071601).