Unbiased Differentially Private Mechanism

Lower Bounds on Error via Dimension Reduction

Haohua Tang

Differential Privacy

An algorithm A is ε-differentially private (ε-DP) if for every two neighboring datasets X, X', and every measurable subset S of the range of A, A satisfies

$$\mathbb{P}[A(X) \in S] \leq e^{\varepsilon}\mathbb{P}[A(X') \in S]$$

Neighboring = different in only 1 data point.

E.g. {(1,2), (2,3), (3,4)} and {(1,2), (4,5), (3,4)}

We assume ε is small enough $\mathrm{s.t.}\ \varepsilon \cong e^{\varepsilon} - 1$.

Property of DP: Post-processing (doing anything not looking at datapoints) preserves DP.

Unbiased Mechanisms

We say a mechanism M for answering query f is unbiased if for every dataset X, M satisfies

$$\mathbb{E}[M(X)] = f(X)$$

E.g. adding any noise with mean 0.

Error

The l_2 error of a mechanism M for answering query f is

$$\sqrt{\mathbb{E}[(M(X) - f(X))^2]}$$

This is $\sqrt{\boldsymbol{tr}(\Sigma)}$ for unbiased M, where Σ is the covariance of $M(X)$.

Mean Point Problem

$X = \{x_1, x_2, \dots, x_n\},\ x_i \in K \subseteq \mathbb{R}^d,\ f(X) = \frac{1}{n}\sum_{i=1}^{n} x_i$

Reason of studying this relatively simple query: for other linear f, we can shift the space and solve mean point problem there, and shifting back is post-processing.

We will assume $K = UB^d$ for some $\mathbf{U} \succcurlyeq \mathbf{0}$.

Support function

The support function of non-empty closed convex set $K \subseteq \mathbb{R}^d$ is defined to be:

$$h_K(\theta) = \sup_{x \in K}\{\theta^T x\}$$

Where $\theta \in \mathbb{R}^d$.

We also define width function to be:

$$w_K(\theta) = h_K(\theta) + h_K(-\theta)$$

When $0 \in K$ we have:

$$w_K(\theta) \geq h_K(\theta)$$

Reduction to 1-dimension

Idea: For all direction (1 dim), we show the variance on that direction is large.

Formally, for all $\theta \in \mathbb{R}^d$:

$$\sqrt{Var(\theta^T M(X))} \gtrsim \frac{w_K(\theta)}{\varepsilon n}$$

We will show this later.

We can also show that

$$\sqrt{Var(\theta^T M(X))} = h_{\Sigma^{0.5} B_2^d}(\theta)$$

Reduction to 1-dimension

We use the following fact (can be proved by Hyperplane Separation Theorem):

$$A \subseteq B \Leftrightarrow \forall\theta, h_A(\theta) \leq h_B(\theta)$$

This gives us $K \subseteq C\varepsilon n\Sigma^{0.5}B_2^d$ for some absolute constant C.

Lower bound on error:

$$\sqrt{\boldsymbol{tr}(\Sigma)} \gtrsim \frac{min\left\{\sqrt{\boldsymbol{tr}(A)}: V \succcurlyeq 0 \wedge K \subseteq V^{0.5}B_2^d\right\}}{\varepsilon n}$$

1 dimensional problem

Now we can focus on 1 dimensional setup and show

$$\sqrt{Var(\theta^T M(X))} \gtrsim \frac{w_K(\theta)}{\varepsilon n}$$

We use HCR bound to obtain this.

HCR Bound

Lemma. Hammersley–Chapman–Robbins (HCR) lower bound:

For distributions P and Q,

$$\chi^2(P\|Q) \geq \frac{\left(\mathbb{E}_P[Y]-\mathbb{E}_Q[Y]\right)^2}{Var_Q(Y)}$$

The Chi-square divergence is defined to be

$$\chi^2(P\|Q) = \mathbb{E}_q[\left(\frac{p(y)}{q(y)} - 1\right)^2]$$

Selecting P and Q

Idea: Obtain the χ^2 divergence from DP and obtain the $\left(\mathbb{E}_P[Y] - \mathbb{E}_Q[Y]\right)^2$ term using M is unbiased.

Let P and Q be the distribution of $\theta^T M(X_1)$ and $\theta^T M(X_2)$. Let $p(y)$ and $q(y)$ denote the PDF of P and Q.

We let X_2 be arbitrary from K^n. We select X_1 $s.t. \left|\theta^T\left(f(X_1) - f(X_2)\right)\right| \geq \frac{w_K(\theta)}{2n}$ and X_1 and X_2 are neighbouring datasets. By this we have

$$\left(\mathbb{E}_P[Y] - \mathbb{E}_Q[Y]\right)^2 \geq \left(\frac{w_K(\theta)}{2n}\right)^2$$

Such X_1 always exists! (by linearity of f)

$\chi^2(P\|Q)$

Let $r(y) = \frac{p(y)}{q(y)}$.

Observations:

1. $\mathbb{E}_q[r(y) - 1] = 0$.
2. $r(y) - 1 \in [e^{-\varepsilon} - 1, e^{\varepsilon} - 1]$, by definition of DP.

Lemma. $\forall x \in \mathbb{R}, \forall \alpha \in \mathbb{R}^+$,

$$x \in [e^{-\alpha} - 1, e^{\alpha} - 1] \wedge \mathbb{E}[x] = 0 \Longrightarrow \mathbb{E}[x^2] \leq e^{-\alpha}(e^{\alpha} - 1)^2.$$

Applying this directly we have

$$\chi^2(P\|Q) = \mathbb{E}_q[(r(y) - 1)^2] \leq e^{-\varepsilon}(e^{\varepsilon} - 1)^2$$

1 Dimensional Lower Bound

By HCR Bound we have the following lower bound:

$$\sqrt{Var_Q(Y)} \geq \frac{w_K(\theta)}{2ne^{-0.5\varepsilon}(e^{\varepsilon}-1)}$$

Since X_2 is selected arbitrarily, when ε is small this is exactly what we want

$$\sqrt{Var(\theta^T M(X))} \gtrsim \frac{w_K(\theta)}{\varepsilon n}$$

This lower bound is asymptotically tight for 1 dimension.

Higher Dimensional Lower Bound

Now we have

$$\sqrt{\boldsymbol{tr}(\Sigma)} \gtrsim \frac{min\left\{\sqrt{\boldsymbol{tr}(A)}: V \succcurlyeq 0 \wedge K \subseteq V^{0.5} B_2^d\right\}}{\varepsilon n}$$

By $K = \mathrm{U} B_2^d$, $\; min\left\{\sqrt{\boldsymbol{tr}(A)}:\ V \succcurlyeq 0\ \wedge\ K \subseteq V^{0.5} B_2^d\right\} = \sqrt{\boldsymbol{tr}(U^T U)}$

So we have error of $\Omega(\frac{1}{\varepsilon n}\sqrt{\boldsymbol{tr}(U^T U)})$.

Unfortunately, this is not tight: Error of the Laplace Mechanism is $\mathrm{O}\left(\frac{\sqrt{d}}{\varepsilon n}\sqrt{\boldsymbol{tr}(U^T U)}\right)$.

There is a better approach (Packing Lower Bound) that yields a tight lower bound.

☹, but we can get asymptotically tight lower bound for zCDP!

zCDP

An algorithm A is ρ-zero-concentrated differentially private (ρ-zCDP) if for every two neighboring databases X, X', and every measurable subset S of the range of A, and for all $\alpha \in (1, \infty)$, A satisfies

$$D_\alpha(A(X) \| A(X')) \leq \rho\alpha$$

The α-Rényi divergence is defined to be

$$D_\alpha(P \| Q) = \frac{1}{\alpha - 1} \log\left(\mathbb{E}_{y \sim Q}\left[\left(\frac{P(y)}{Q(y)}\right)^\alpha\right]\right)$$

Setting $\alpha = 2$ this gives us lower bound on Chi-square divergence.

Lower Bound for zCDP

Plug in $\alpha = 2$ we have

$$\chi^2(P\|Q) = \mathbb{E}_q[(r(y)-1)^2] = \mathbb{E}_q[r(y)^2] - 1 \leq e^{2\rho} - 1 \cong 2\rho$$

This gives us

$$\sqrt{\boldsymbol{tr}(\Sigma)} \gtrsim \frac{min\left\{\sqrt{\boldsymbol{tr}(A)}: V \succcurlyeq 0 \wedge K \subseteq V^{0.5} B_2^d\right\}}{\sqrt{\rho}n}$$

For $K = \mathrm{U}B_2^d$, we have error $\Omega(\frac{1}{\sqrt{\rho}n}\sqrt{\boldsymbol{tr}(U^T U)})$.

Matches error of the Gaussian Mechanism $\mathrm{O}\left(\frac{1}{\sqrt{\rho}n}\sqrt{\boldsymbol{tr}(U^T U)}\right)$.

If your algorithm is unbiased, you cannot do asymptotically better on zCDP than just adding a Gaussian noise.

Future Works

1. Lower bounds for ADP(work in progress, we believe this is tight)
2. More general spaces.(general closed convex set)
3. More general queries.(non-linear ones?)

Thank you!

The Laplace Mechanism

l_1 Global Sensitivity is $GS_{f,l_1} = \sup_{X_1, X_2 neighbours} \| f(X_1) - f(X_2) \|_1$.

The Laplace Mechanism adds noise of $Z_i \sim Lap(\frac{GS_{f,l_1}}{\varepsilon})$ to each of the d dimensions to achieve ε-DP.

PDF of Laplace distruibution is $Lap(\lambda)$ is $h_\lambda(y) = \frac{\exp(-\frac{|y|}{\lambda})}{2\lambda}$.

Variance of Laplace distribution is $2\lambda^2$.

The Laplace Mechanism

Consider running the Laplace Mechanism when $K = B_2^d$. (Same query)

Sensitivity is $GS_{f,l_1} = \frac{2\sqrt{d}}{n}$, so adding noise of $\mathrm{Z}_i \sim Lap(\frac{2\sqrt{d}}{\varepsilon n})$ to each of the d dimensions.

Covariance would be $\frac{8d}{\varepsilon^2 n^2} I$.

The Laplace Mechanism

Back to $K = UB_2^d$ .We can add noise of $\mathbf{UZ}_i$ to each of the d dimensions.

Still DP since this is post-processing.

Covariance would be $\frac{8d}{\varepsilon^2 n^2} UU^T$.

Error of $\mathrm{O}\left(\frac{\sqrt{d}}{\varepsilon n}\sqrt{\boldsymbol{tr}(U^T U)}\right)$.

The Gaussian Mechanism

l_2 Global Sensitivity is $GS_{f,l_2} = \sup_{X_1, X_2 neighbours} \|f(X_1) - f(X_2)\|_2$.

The Gaussian Mechanism adds noise of $Z_i \sim \mathcal{N}(0, \frac{(GS_{f,l_2})^2}{2\rho})$ to each of the d dimensions to achieve ρ-zCDP.

By similar argument (shifting with U) we can get error is $O\left(\frac{1}{\sqrt{\rho}n}\sqrt{\boldsymbol{tr}(U^T U)}\right)$.