Lecture 15

Private Mean Estimation

Binary Mean Estimation

$X_1, \dots, X_n$, $X_i \in \{0, 1\}$

$\tilde{p} = \frac{1}{n} \sum X_i$

$\hat{p} = \frac{1}{n} \sum X_i + \mathrm{Lap}\left(\frac{1}{\varepsilon n}\right)$

Dataset mean

$|\hat{p} - \tilde{p}| \leq O\left(\frac{1}{\varepsilon n}\right)$ (w.h.p.)

Acc α? $\frac{1}{\varepsilon n} \leq \alpha \Rightarrow n \geq \frac{1}{\alpha \varepsilon}$.

$X_1, \dots, X_n \overset{iid}{\sim} \mathrm{Bernoulli}(p)$ $\quad 0 \leq p \leq 1$

$E[X_i] = p$, $E\left[\frac{1}{n} \sum X_i\right] = p$.

$\mathrm{Var}[X_i] = p(1-p)$. $\mathrm{Var}\left[\frac{1}{n} \sum X_i\right] = \frac{1}{n^2} \mathrm{Var}\left[\sum X_i\right] = \frac{1}{n^2} \cdot n \mathrm{Var}[X_i]$

$= \frac{p(1-p)}{n} \leq \frac{1}{4n}$

$|p - \tilde{p}| \leq O\left(\frac{1}{\sqrt{n}}\right)$

$|p - \hat{p}| \leq |p - \tilde{p}| + |\tilde{p} + \hat{p}| \leq O\left(\frac{1}{\sqrt{n}} + \frac{1}{\varepsilon n}\right)$

Goal: α-error, $n \geq$?

Non-private $\quad$ Privacy

$\frac{1}{\sqrt{n}} + \frac{1}{\varepsilon n} \leq \alpha \Rightarrow n \geq O\left(\frac{1}{\alpha^2} + \frac{1}{\alpha \varepsilon}\right)$

$$\left|p - \left(\frac{1}{n}\sum X_i + N\right)\right| \leq \underbrace{\left|p - \frac{1}{n}\sum X_i\right|}_{\text{sampling error}} + \underbrace{\left|\frac{1}{n}\sum X_i - \left(\frac{1}{n}\sum X_i + N\right)\right|}_{\text{noise error}}$$

Unbounded Data

$X_1, \dots, X_n$, $X_i \in \mathbb{R}$. heights. Avg height?

$\tilde{\mu} = \frac{1}{n} \sum X_i$.

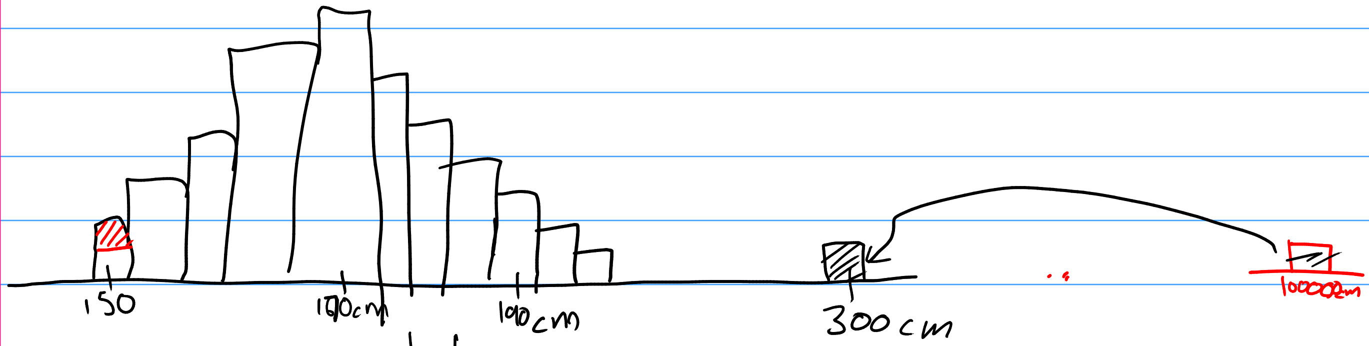

$300\text{cm} \geq$ Humans height ≥ 0 cm

1. Clip all data to given range (let f be fn which does this)
2. Compute $\frac{1}{n} \sum f(X_i) + \text{Lap}\left(\frac{300}{\varepsilon n}\right)$

$$f(X_i) = \begin{cases} 0 & \text{if } X_i < 0 \\ 300 & \text{if } X_i > 300 \\ X_i & \text{o.w.} \end{cases}$$

Private Parameter Estimation of a Dist

Goals

$X = X_1, \ldots, X_n$. Want $M(X)$ to satisfy:

1. Privacy: M is DP
2. Accurate: If $X \underset{iid}{\sim} P$ (w. approp properties), algo is accurate w.h.p.

Univariate Gaussian Estimation

$X_1, \dots, X_n$

Privacy: ε-DP.

Accuracy: If $X_1, \dots, X_n \sim N(\mu, 1)$, $|\mu| \le R$, estimate μ whp.

Claim: Whp. $X_1, \dots, X_n \in [\mu - O(\sqrt{\log n}), \mu + O(\sqrt{\log n})]$

pf:

$$\Pr_{X \sim N(\mu,1)}[|X - \mu| \ge t] \le 2\exp\left(-\frac{t^2}{2}\right) .$$

$t = \sqrt{20 \log n}$, $\le \frac{2}{n^{10}}$.

Union bound: $\le \frac{2}{n^9}$. $\square$

Corr: $X_1, \dots, X_n \in [-R - O(\sqrt{\log n}), R + O(\sqrt{\log n})]$.

Naïve Approach

(Naive) Thm: ε-DP algo, est mean of $N(\mu,1)$ $(|\mu|\le R)$ to acc α,

$$n=\tilde{O}\left(\frac{1}{\alpha^2}+\frac{R}{\alpha\varepsilon}\right) \text{ samples.}$$

Proof:

1. Clip dataset to $[-R-O(\sqrt{\log n}), R+O(\sqrt{\log n})]$.
2. $\hat{\mu}=\frac{1}{n}\sum \underbrace{f(X_i)}_{\text{clipped pts}}+\mathrm{Lap}\left(\frac{2R+O(\sqrt{\log n})}{n\varepsilon}\right)$

ε-DP.

Clipping won't move pts.

$|\mu-\tilde{\mu}|\le O\left(\frac{1}{\sqrt{n}}\right)$

$|\tilde{\mu}-\hat{\mu}|=\left|\mathrm{Lap}\left(\frac{2R+O(\sqrt{\log n})}{n\varepsilon}\right)\right|\le\tilde{O}\left(\frac{R}{n\varepsilon}\right)$

$|\mu-\hat{\mu}|\le\tilde{O}\left(\frac{1}{\sqrt{n}}+\frac{R}{n\varepsilon}\right)$.

Acc $\alpha\Rightarrow \underline{n\ge\tilde{O}\left(\frac{1}{\alpha^2}+\frac{R}{\alpha\varepsilon}\right)}$ $\blacksquare$

$$\left|\mu-\left(\frac{1}{n}\sum f(X_i)+N\right)\right|\le\underbrace{\left|\mu-\frac{1}{n}\sum X_i\right|}_{\substack{\text{sampling}\\ O\left(\frac{1}{\sqrt{n}}\right)}}+\underbrace{\left|\frac{1}{n}\sum X_i-\frac{1}{n}\sum f(X_i)\right|}_{\substack{\text{bias}\\ 0\ (\text{w.h.p.})}}+\underbrace{\left|\frac{1}{n}\sum f(X_i)-\left(\frac{1}{n}\sum f(X_i)+N\right)\right|}_{\substack{\text{noise}\\ \tilde{O}\left(\frac{R}{n\varepsilon}\right)}}$$

Thm: ε-DP Algo. Est μ of $N(\mu,1)$ $(|\mu|\le R)$ to acc α

$$n=\tilde{O}\left(\frac{1}{\alpha^2}+\frac{1}{\alpha\varepsilon}+\frac{\log R}{\varepsilon}\right) \text{ samples}$$

Histogram-based Approach

Karwa-Vadhan '18

1. Coarse estimate: $\tilde{O}\left(\frac{\log R}{\varepsilon}\right)$ $\quad \frac{\varepsilon}{2}$
2. Fine estimate $\tilde{O}\left(\frac{1}{\alpha^2}+\frac{1}{\alpha\varepsilon}\right)$ $\quad \frac{\varepsilon}{2}$

1. $[-R - O(\sqrt{\log n}),\ R + O(\sqrt{\log n})]$

Divide into $\underbrace{O\left(\frac{R}{\sqrt{\log n}}\right)}_{\le O(R)}$ intervals of width $O(\sqrt{\log n})$

Laplace Histograms $\qquad$ $\mathrm{Lap}(\frac{1}{\varepsilon})$ noise to each

Max error $\le \frac{\log \#\text{of bins}}{\varepsilon} \le \frac{\log R}{\varepsilon}$

$$\frac{n}{2} - \frac{\log R}{\varepsilon} > 0 + \frac{\log R}{\varepsilon} \implies n \gtrsim \frac{\log R}{\varepsilon}$$

Coarse Lem: ε-DP algo, finds I s.t. $\mu \in I$. $\quad N(\mu,1),\ |\mu| \le R$

$$n \ge \tilde{O}\left(\frac{\log R}{\varepsilon}\right) \implies |I| \le O(\sqrt{\log n})$$

Now: $R \le O(\sqrt{\log n})$

$$n = \tilde{O}\left(\frac{1}{\alpha^2} + \frac{\sqrt{\log n}}{\alpha\varepsilon}\right) = \tilde{O}\left(\frac{1}{\alpha^2}+\frac{1}{\alpha\varepsilon}\right).$$

Shrinking Confidence Intervals

Biswas, Dong, K, Ullman

Laplace Mech.

Naive method: Clips data, noises empirical mean

BDKU: Naive + confidence interval

1. Clip dataset to $[-R-O(\sqrt{\log n}), R+O(\sqrt{\log n})]$.
2. $Z = \frac{1}{n}\sum f(X_i) + \mathrm{Lap}\left(\frac{2R+O(\sqrt{\log n})}{n(\varepsilon/\log R)}\right)$
3. Return interval centered at Z, of width $O\left(\frac{1}{\sqrt{n}} + \frac{R+\sqrt{\log n}}{\varepsilon n}\right)$

Claims: - If $n \geq O\left(\frac{\log R}{\varepsilon}\right)$ and $R \geq C\sqrt{\log n}$, then returned interval is const factor smaller

- If $X_1, \ldots, X_n \sim N(\mu, 1)$, interval $\ni \mu$ w.h.p.

Proof: $\mu \in [-R, R]$. Final: $O\left(\frac{1}{\sqrt{n}} + \frac{R+\sqrt{\log n}}{\varepsilon n}\right)$

$$\boxed{n \geq \frac{100 \log R}{\varepsilon}} \qquad \leq \cancel{\frac{\sqrt{\varepsilon}}{10}} \boxed{+\frac{R}{100}} + \cancel{\frac{\sqrt{\log n}}{100}}$$

$$|Z-\mu| \leq \left|\frac{1}{n}\sum X_i - \mu\right| + \left|\mathrm{Lap}\left(\frac{2R+O(\sqrt{\log n})}{n\varepsilon}\right)\right|$$

$$\leq O\left(\frac{1}{\sqrt{n}}\right) + O\left(\frac{R+\sqrt{\log n}}{n\varepsilon}\right) \text{ whp } \square$$

$R \to O(\sqrt{\log n})$

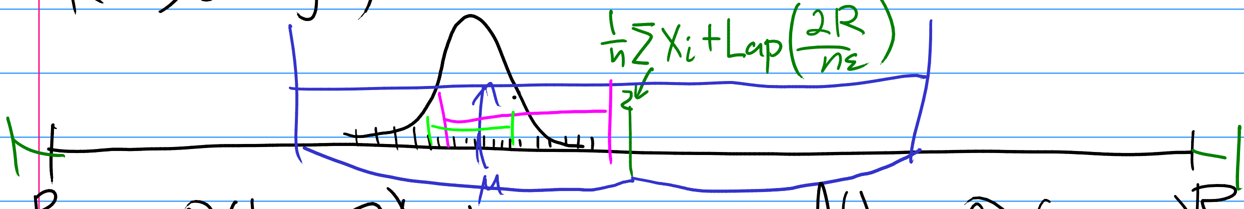

$t = O(\log R)$ iters $\Rightarrow$ width $= O(\sqrt{\log n})$

$+\varepsilon$-DP. $\varepsilon' = \frac{\varepsilon}{T}$ $\boxed{n \gtrsim \frac{\log R}{\varepsilon}}$ $R = O(\sqrt{\log n})$

$n = \tilde{O}\left(\frac{1}{\alpha^2} + \frac{\sqrt{\log n}}{\alpha\varepsilon}\right) = \boxed{\tilde{O}\left(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon}\right)}.$

Beyond Univariate Gaussians

CoinPress $\mu \in [-R, R]$

1. Clip the data. Do this based on the confidence interval containing the parameter, combined with the tail bounds of the distribution class.

2. Compute the empirical estimator for the quantity, and add noise proportional to the sensitivity (which should be bounded due to clipping).

3. Define a new confidence interval centered around this estimate, with a width based on the sampling error and the noise added.

4. Repeat.

Multivariate Mean Estimation

Theorem 5. *There exists an (ε, δ)-differentially private algorithm which estimates the mean of $N(\mu, I)$ (where $\|\mu\|_2 \le R$) to ℓ_2-accuracy α, given*

d-dimensional

$$n = \tilde{O}\left(\frac{d}{\alpha^2} + \frac{d\sqrt{\log(1/\delta)}}{\alpha\varepsilon} + \frac{\sqrt{d \log R \log(1/\delta)}}{\varepsilon}\right)$$

samples.

$Y \sim N(0, I)$, $\|Y\|_2 \approx \sqrt{d}$

$\mu \in B_2(0, R)$

1. Clip the data. Do this based on the confidence interval containing the parameter, combined with the tail bounds of the distribution class. $O(\sqrt{d} + \sqrt{\log n})$ $\rightarrow$ Clip to ball $B_2(0, R + \sqrt{d} + \sqrt{\log n})$

2. Compute the empirical estimator for the quantity, and add noise proportional to the sensitivity (which should be bounded due to clipping). $\frac{1}{n}\sum X_i + N(0, \ldots)$ $\sigma = \left(\frac{R + \sqrt{d} + \sqrt{\log n}}{n\varepsilon}\right)$

3. Define a new confidence interval centered around this estimate, with a width based on the sampling error and the noise added. $\sqrt{\frac{d}{n}} + \sqrt{d}\left(\frac{R + \sqrt{d} + \sqrt{\log n}}{n\varepsilon}\right)$ err in each coord

4. Repeat. log R times

$\rightarrow$ Radius $R = \tilde{O}(\sqrt{d})$

$\hookrightarrow$ Use naive estimator.

$\frac{\sqrt{d}\,R}{n\varepsilon}$

$n \geq \frac{\log \sqrt{d}\sqrt{\log R}}{\varepsilon}$

Covariance

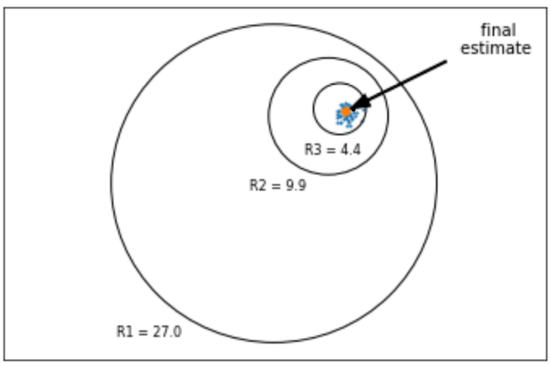

Heavy-Tailed Data

Sub-Gaussian

↳ all moments bounded

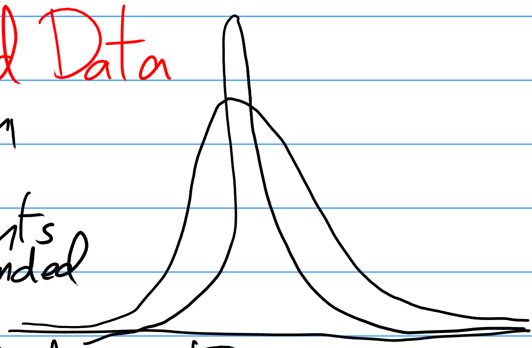

Heavy-tailed dist. P

Moments: $\mathop{E}_{X\sim P}[(X-\mu)^k]$ ← kth moment

Simple example:

$\mu = E[X] \in [-1, 1]$, $E[(X-\mu)^2] \le 1$

$|\mu - (\frac{1}{n}\sum f(X_i) + N)|$

$\le \underbrace{|\mu - \frac{1}{n}\sum X_i|}_{\text{sampling } O(\frac{1}{\sqrt{n}})} + \underbrace{|\frac{1}{n}\sum X_i - \frac{1}{n}\sum f(X_i)|}_{\text{bias from clipping } O(\frac{1}{\tau})} + \underbrace{|\frac{1}{n}\sum f(X_i) - (\frac{1}{n}\sum f(X_i) + N)|}_{\text{noise for privacy } O(\frac{\tau}{n\varepsilon})}$

$[-1-\tau, 1+\tau]$. $\frac{1}{\sqrt{n}} + \frac{1}{\tau} + \frac{\tau}{n\varepsilon} \le \alpha$. $\tau = \frac{\alpha}{3}$, $n \ge O(\frac{1}{\alpha^2 \varepsilon})$

$\Pr[|X-\mu| \ge 10\sqrt{n}] \le 1/100n$. Union bound: $O(\frac{1}{\alpha^2} + \frac{1}{\alpha\varepsilon})$

$X_1, \dots, X_n \sim P$, $\in [\mu - 10\sqrt{n}, \mu + 10\sqrt{n}]$

0 bias ⟺ $\tau = \Theta(\sqrt{n})$. $O(\frac{1}{\sqrt{n}})$, 0, $O(\frac{1}{\sqrt{n}\varepsilon})$

$N = Lap(\frac{\Theta(\sqrt{n})}{n\varepsilon})$

$\frac{1}{\sqrt{n}\varepsilon} \le \alpha \to n \ge \frac{1}{\alpha^2\varepsilon^2}$.

Claim: Clipping to $[\mu-\tau, \mu+\tau]$ gives

$\underbrace{|\frac{1}{n}\sum X_i - \frac{1}{n}\sum f(X_i)|}_{\text{bias}} \le O(1/\tau)$

K. Singhal Ullman '20

Sorry I said $\tau = \frac{\alpha}{3}$, it should be $\tau = \frac{3}{\alpha}$!