Lecture 11

- Selection Problem
- Approximate DP
  - Gaussian Mechanism.

Announcements:

→ Recitation (in-person) Friday.

→ Guest Lecture (Jon Ullman)
Will post Zoom / Youtube Link.
Next Weds.

Selection Problem

$\mathcal{Y}$: possible outcomes

$q: \mathcal{Y} \times \mathcal{X}^n \to \mathbb{R}$ "score" function

measures how good y is on dataset x.

q is Δ-sensitive if $\forall\, y \in \mathcal{Y}$

$q(y; \cdot)$ has $GS_q \le \Delta$.

Exponential Mechanism. $\mathcal{A}_{EM}(x, q, \varepsilon, \Delta)$

Output an outcome y with probability proportional to

$$\exp\left(\frac{\varepsilon}{2\Delta} q(y; x)\right).$$

For this class, assume outcome space $\mathcal{Y}$ is finite.

$$\mathbb{P}[\mathcal{A}_{EM}(x, q, \varepsilon, \Delta) = y] = \frac{1}{C_x} \cdot \exp\left(\frac{\varepsilon}{2\Delta} q(y; x)\right)$$

"Normalization factor" $C_x = \sum_{y' \in \mathcal{Y}} \exp\left(\frac{\varepsilon}{2\Delta} q(y'; x)\right).$

Privacy Proof.

Theorem. For every Δ-sensitive q, $A_{EM}(\cdot, q, \varepsilon, \Delta)$ is ε-DP. } Privacy Guarantee.

Utility Guarantee.

$$\forall t > 0, \quad \underset{\hat{y} \sim A_{EM}}{\mathbb{P}}\Big[\underbrace{q_{max}(x) - q(\hat{y}; x)}_{\text{error}} \geq \underbrace{\frac{2\Delta}{\varepsilon}\left(\ln(d) + t\right)}_{\text{threshold / error Bound}} \Big] < \underbrace{e^{-t}}_{\text{"failure probability"}} \quad 1\%$$

Tail Bound on the error.

Selection Problem Example

Heavy Hitter

Example. A set of websites $\{1, \dots, d\}$

Each user submits $X_i \subseteq \{1, \dots, d\}$

Winner: website with the highest score: $\forall j \in \{1, \dots, d\}$

$$q(j; x) = \left| \{ i \mid j \in X_i \} \right|.$$

" # users visiting "

$$\text{Error} = \underbrace{\max_j q(j; x)}_{q_{max}} - q(A(x); x).$$

Error Bounds:

① Exp Mech.

W. p. 99%, error $< \dfrac{2 \ln(100d)}{\varepsilon}$

② Laplace Mech.

Skip? $\longrightarrow$ Release $q(y; x) + \text{Lap}\left(\dfrac{GS_q}{\varepsilon}\right) \quad \forall y.$

$$GS_q = \max_{x, x'} \sum_{j=1}^{d} \left| q(j; x) - q(j; x') \right| = d.$$

$$\text{Error} \geq \left(\frac{d}{\varepsilon}\right)$$

Report Noisy Max

$$A_{RNM}(x,\ q,\ \Delta,\ \varepsilon):$$

1) For every outcome $y = 1, 2, \dots, d$

"Fresh"

$\tilde{q}(y) = q(y; x) + \text{Noise}$ ← Not Releasing

2) Return $\hat{y} = \arg\max_{y \in \{1,\dots,d\}} \tilde{q}(y)$

"Noisy max"

$\hat{y}$ maximizes $\tilde{q}(y)$

What Noise

→ Laplace Noise

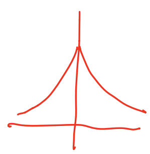

Lap.

RNM + Noise ≥

→ Exponential distribution

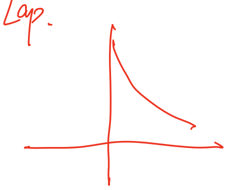

Exp

→ Gumbel distribution

↳ (Recover Exp mech.)

RNM (noise ~ Gumbel) = exp mech.

density $h(z) = e^{-(z + e^{-z})}$

Other Perturbations

Theorem. A_{RNM} is ε-DP. ⟵ Privacy

$\forall x$, $y = A_{RNM}(x, q, \Delta, \varepsilon)$

Expectation: $\mathbb{E}[q(y;x)] \geq q_{max} - \frac{2\Delta}{\varepsilon}[\ln(d)+1]$

Tail:

$\forall t>0$, $\mathbb{P}[q(y;x) < q_{max} - \frac{2\Delta}{\varepsilon}(\ln(d)+t)] < \exp(-t)$

↑ Utility

Approximate Differential Privacy

- (ε, δ) - DP
- Gaussian Mechanism.

Definition. (Differential Privacy)

A is ε-differentially private if

for all neighbors x and x'

for all subsets E of outputs

$$\mathbb{P}[A(x) \in E] \leq e^{\varepsilon}\, \mathbb{P}[A(x') \in E]$$

Is this too stringent?

Suppose there is some $E \subseteq Y$ such that

$$\mathbb{P}[A(x) \in E] \leq \frac{1}{2^{100}}$$

and $\mathbb{P}[A(x') \in E] = 0$.

(ε, δ) - DP

A is (ε, δ)-differentially private if
for all neighbors x and x'
for all subsets E of outputs

$$\mathbb{P}[A(x) \in E] \leq \underbrace{e^{\varepsilon}}_{\text{Multiplicative Approximation}} \mathbb{P}[A(x') \in E] + \underbrace{\delta}_{\text{Additive Approximation}}$$

Naming Convention:

ε-DP, $(\varepsilon, 0)$-DP, "pure" DP

(ε, δ)-DP, "approximate" DP

Intepretation of δ:

probability of "Privacy Failure"

"Name & Shame" Algorithm

$NS_\delta(x_1, x_2, \dots, x_n)$

For each $i = 1, \dots, n$

Release $Y_i = \begin{cases} X_i & \text{w.p. } \delta \\ \text{"}\perp\text{"} & \text{w.p. } (1-\delta) \end{cases}$

NS_δ satisfies $(\overset{\varepsilon}{0}, \delta)$-DP.

If $\delta > \frac{1}{n}$, (e.g. $\frac{20}{n}$) release ≈ 20 in the clear.

$\rightarrow \delta \ll \frac{1}{n}$

$\frac{1}{2^{20}}$, $\frac{1}{2^{100}}$, (Paper writing: $\delta = \frac{1}{n^2}$)

Preserve Nice Properties.

① Post-Processing

(ε,δ)-DP $A: \mathcal{X}^n \to Y$, $f: Y \to Y'$

$$f(A(\cdot)) \text{ is } (\varepsilon,\delta)\text{-DP}$$

② Adaptive Composition

$A_1: \mathcal{X}^n \longrightarrow Y_1$ (ε_1,δ_1)-DP

$A_2: \mathcal{X}^n \times Y_1 \longrightarrow Y_2$ (ε_2,δ_2)-DP

"Basic" Composition: $\left(A_1(x),\ A_2(x, A_1(x))\right)$ is $(\varepsilon_1+\varepsilon_2,\ \delta_1+\delta_2)$-DP.

Example: Truncated Laplace.

For f with $GS_f = 1$.

$$A(x) = f(x) + Z,$$

Truncated Laplace

$$p(z) = \begin{cases} \frac{1}{C} \cdot \exp(-|z|), & |z| \le \tau \\ 0 & |z| > \tau \end{cases}$$

$$C = \int_{-\tau}^{\tau} e^{-|z|} \, dz.$$

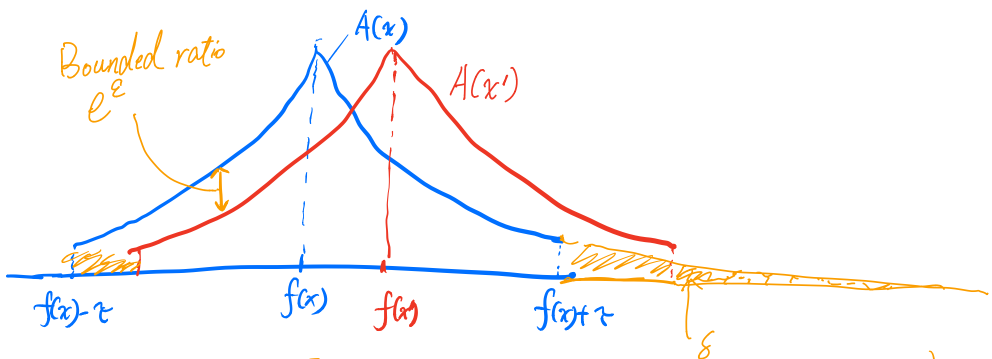

$$\mathbb{P}_{A(x)}\left[\boxed{} \right] \le \delta \implies \tau \approx O\left(\log\left(\tfrac{1}{\delta}\right)\right)$$

(ε, δ)-DP

Gaussian Mechanism

1-dim case $\quad f: \mathcal{X}^n \mapsto \mathbb{R}$, w/ $GS_f = \Delta$

$$A(x) = f(x) + N\left(0, \frac{2\Delta^2 \log(2/\delta)}{\varepsilon^2}\right)$$

Gaussian

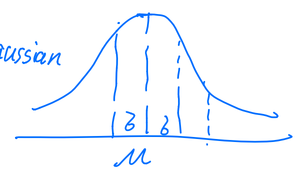

Gaussian dist $= N(\mu, \sigma^2)$

"sigma"

$$p_{\mu,\sigma^2}(y) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-(x-\mu)^2}{2\sigma^2}\right)$$

pi

Lap

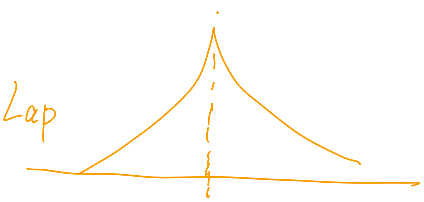

How to prove (ε, δ)-DP?

For ε-DP. suffices to prove

$$\forall y \in Y, \quad \mathbb{P}[A(x)=y] \le e^{\varepsilon} \mathbb{P}[A(x')=y]$$

$$\Leftrightarrow \forall E \subseteq Y, \quad \mathbb{P}[A(x) \in E] \le e^{\varepsilon} \mathbb{P}[A(x') \in E]$$

$\forall y \in$ "Good set w/ bounded ratio of e^{ε}"

$\mathbb{P}$["Bad set"] $\le \delta$.

Theorem. For any $\varepsilon \leq 1$ $\delta > 0$, the (1-dim) Gaussian Mechanism satisfies (ε, δ)-DP.

~~Proof.~~

1-dim case $f: \mathcal{X}^n \mapsto \mathbb{R}$, w/ $GS_f = \Delta$

$$A(x) = f(x) + N\left(0, \frac{2\Delta^2 \log(2/\delta)}{\varepsilon^2}\right)$$