Lecture 11

Today

- Recap Hypothesis testing
- subgaussian & zCDP
- RDP

HW 2 due Tuesday

Concentrated DP

PLRV $\quad L = L_A^{x \to x'}(Y) = \log\left(\frac{\Pr[A(x)=Y]}{\Pr[A(x')=Y]}\right)$

under $Y \sim A(x)$.

For $A(x) = (A_1(x), \ldots, A_k(x))$, output $Y = (Y_1, \ldots, Y_k)$

$$L = \sum_{j=1}^{k} L_j = \sum_{j=1}^{k} L_{A_j}^{x \to x'}(Y_j).$$

CLT says: $L \approx$ Gaussian

We want: mathematical tools to make this formal. Say: Gaussian DP

Def 1 A is ρ-zero-concentrated DP (ρ-zCDP) if $\forall x \sim x'$ and PLRV $L = L_A^{x \to x'}(y)$ satisfies, $\forall \lambda \geq 0$, $\mathbb{E}[\exp(\lambda L)] \leq \exp(\lambda(\lambda+1)\rho)$

Interpretation: PLRV looks like a Gaussian.

Why?

Claim If $X \sim \mathcal{N}(0,1)$, then $\forall \lambda \geq 0$

$$\mathbb{E}[\exp(\lambda X)] = \exp(\lambda^2/2)$$

More generally, a standard notion of "looks like Gaussian" is called "subgaussian."

Distribution has tails that die off at least as quickly as a Gaussian.

Claim (subgaussian properties) Let X be random variable with $\mathbb{E}X = 0$. There exist K_1, K_2, K_3 differing by an absolute constant factor such that the following are equivalent:

i) Tails: $\forall t \geq 0$,

$$\Pr\left[|X| \geq t\right] \leq 2\exp\left(-t^2/K_1^2\right)$$

ii) Moments: $\forall p \geq 1$

$$\left(\mathbb{E}|X|^p\right)^{1/p} \leq K_2\sqrt{p}$$

iii) Moment-generating function:

$$\forall \lambda \in \mathbb{R}, \quad \mathbb{E}\left[\exp(\lambda X)\right] \leq \exp\left(K_3^2 \lambda^2\right)$$

Pf (iii) $\Rightarrow$ (i), take Markov & $\lambda = t/2$.

Claim If A is ε-DP, then it is

$\left(\frac{1}{2}\varepsilon^2\right)$-zCDP

Claim If A is ρ-zCDP, then for any $\delta > 0$ it is (ε', δ)-DP for

$$\varepsilon' = \rho + 2\sqrt{\rho \log(1/\delta)}$$

Skip

Claim 1 Let $f: \mathcal{X}^n \to \mathbb{R}^d$ have global sensitivity $\Delta = \max_{x \sim x'} \|f(x) - f(x')\|_2$.

Let $A(x) = f(x) + \mathcal{N}(0, \sigma^2 \mathbb{I})$ for $\sigma > 0$.

Then $A(x)$ is ρ-zCDP for $\rho = \frac{\Delta^2}{2\sigma^2}$.

Claim 2 If A is ε-DP, then A is ρ-zCDP for $\rho = \frac{1}{2}\varepsilon^2$.

Claim 3 If A is the composition of k $\left(\frac{1}{2}\varepsilon^2\right)$-zCDP algorithms, then A is $\left(\frac{1}{2}\varepsilon^2 k\right)$-zCDP

Claim 4 If A is $\left(\frac{1}{2}\varepsilon^2 k\right)$-zCDP, then $\forall \delta > 0$ A is (ε', δ)-DP with

$$\varepsilon' = \varepsilon\sqrt{2k\log(1/\delta)} + \frac{1}{2}\varepsilon^2 k$$

Say:
For many mechanisms (esp. Gaussian) simplifies & sharpens discussions of composition

Strength of DP vision is our ability to quantify privacy leakage
Powerful tool to that end.

Rényi Differential Privacy

Say: slightly earlier (?), researchers did a bunch of complicated math about moments in order to track privacy loss while training neural networks privately.
Looked back at their math & realized it was built on this

Def Algorithm A is (α,ε)-Rényi differentially private (RDP) if $\forall x \sim x'$,

$L = L_A^{x \to x'}(Y)$ satisfies

$$\mathbb{E}\left[\exp((\alpha-1)L)\right] \le \exp((\alpha-1)\varepsilon)$$

Say: zCDP, like the subgaussian def, controls concentration at all moments. RDP gives control over specific moments.

Claim If A is $(\alpha, \alpha\rho)$-RDP for all $\alpha \in (1,\infty)$, then A is ρ-zCDP.

We use RDP when:

1. Need to reason about subsampling
2. Mechanisms don't satisfy zCDP

Ex $A(x) \sim \mathcal{N}(\mu_1, \sigma_1^2)$
$A(x') \sim \mathcal{N}(\mu_2, \sigma_2^2)$ for $\sigma_1^2 \approx \sigma_2^2$ (but not equal)

Usual definition

Def For distributions P, Q, the Rényi divergence of order $\alpha > 1$ is

$$D_\alpha(P \| Q) \triangleq \frac{1}{\alpha-1} \mathbb{E}_{x \sim Q}\left(\frac{P(x)}{Q(x)}\right)^\alpha$$

Def A is (α, ε)-RDP if $\forall x, x'$,

$$D_\alpha(A(x) \| A(x')) \leq \varepsilon.$$

observe: zCDP is bd on Rényi divergence $\forall \alpha > 1$

Claim If A is (α, ε)-RDP, then $\forall \delta > 0$ it is (ε', δ)-DP for $\varepsilon' = \varepsilon + \frac{\log 1/\delta}{\alpha - 1}$.