CS839: Differential Privacy and Learning February 26, 2026

12: Zero-Concentrated Differential Privacy, Part 2

Instructor: Gavin Brown *Scribe:* Ronak Chauhan

Disclaimer: This document is intended as an informal supplement to in-class note-taking. It has not been given the level of scrutiny expected in polished lecture notes, let alone that reserved for peer-reviewed publications.

1 Motivating Gaussian DP

Consider algorithm $A = (A_1, \ldots, A_T)$. The total Privacy Loss Random Variable (PLRV) for A is the sum of the individual PLRVs for each step.

Thus, $L = L_A^{x \to x'}(Y) = \sum_{j=1}^{k} L_{A_j}^{x \to x'}(Y_j)$ under $Y \sim A(x)$.

Each term serves as evidence for an adversary trying to determine whether the data came from input x or x'.

As T tends to infinity, the Central Limit Theorem states that $L \approx$ Gaussian, i.e. the privacy loss random variable "looks like a Gaussian".

To formalize the idea that these PLRVs "look like Gaussians," we use mathematical tools like Gaussian DP and Zero-Concentrated DP.

2 Gaussian DP

Last class we talked about the hypothesis testing interpretation of differential privacy. It leads to the following definition: an algorithm A is *μ-Gaussian DP* (μ-GDP) if $\forall x \sim x'$, the difficulty of distinguishing $A(x)$ from $A(x')$ is no harder than distinguishing between a standard normal distribution $\mathcal{N}(0, 1)$ and a shifted normal distribution $\mathcal{N}(\mu, 1)$.

This captures the idea that the PLRV "looks like" a Gaussian.

3 Subgaussian Random Variables

Another standard notion of a distribution that "looks like" a Gaussian is called being *subgaussian*. Subgaussian distributions have tails that die off at least as quickly as a Gaussian distribution. We can also phrase this in terms of the moment-generating function:

$$\text{If } x \sim \mathcal{N}(0, 1), \text{ then } \forall \lambda \geq 0, \mathbb{E}[\exp(\lambda x)] = \exp(\frac{\lambda^2}{2})$$

Random variables which satisfy the following conditions are called *subgaussian*.

Claim 3.1. *Let X be a 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 three properties are equivalent:*

(i) ***Tail bound:***

$$\forall t \geq 0,\ \Pr[|X| \geq t] \leq 2\exp\left(\frac{-t^2}{K_1^2}\right)$$

(ii) ***Moment bound:***

$$\forall p > 1,\ (\mathbb{E}[|X|^p])^{1/p} \leq K_2\sqrt{p}$$

Higher moments correspond to skewness. The moments are not too big and they do not blow up as p increases.

(iii) ***Moment-Generating Function (MGF) bound:***

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

Subgaussian random variables have very light tails, their higher moments are not too big, and we have a bound on the moment-generating function.

We will prove just (iii) $\Rightarrow$ (i), which is most relevant for us. For the rest of the proof and much much more, see Vershynin [2018].

Proof. Assume (iii) holds with $K_3 = 1$ (rescaling K_3 does not change the proof).

Assume $K_3 = 1$
For any $\lambda > 0$, we can rewrite the probability using exponentiation:

$$\begin{aligned}
\Pr[X > t] &= \Pr[\lambda X > \lambda t] \\
&= \Pr[e^{\lambda X} > e^{\lambda t}] \\
&\leq \mathbb{E}[e^{\lambda X}] && \text{by Markov's inequality} \\
&\leq \exp(\lambda^2 - \lambda t)
\end{aligned}$$

This holds for any $\lambda \geq 0$. To obtain the bound we want, we pick $\lambda = t/2$:

$$\Pr[X > t] \leq \exp\left(\frac{t^2}{4} - \frac{t^2}{2}\right) = \exp\left(-\frac{t^2}{4}\right)$$

By symmetry, the same holds for $\Pr[X < -t]$, and applying a union bound yields the final tail property. □

4 ρ-Zero-Concentrated DP (ρ-CDP)

The definition of zCDP uses the MGF to control privacy loss.

Definition 4.1. An algorithm A is *ρ-Zero-Concentrated DP* (ρ-CDP) if $\forall x \sim x'$ PLRV $L = L_A^{x \to x'}(Y)$ satisfies the following:

$$\forall \lambda \geq 0, \mathbb{E}[\exp(\lambda L)] \leq \exp(\lambda(\lambda + 1)\rho)$$

The expectation above is the moment generating function. This definition ensures that the PLRV behaves like a Gaussian random variable

4.1 zCDP to Approximate DP Translation

Claim 4.2. *If algorithm A is ρ-ZCDP, then $\forall \delta > 0$, it is (ε', δ)-DP with $\varepsilon' = \rho + 2\sqrt{\rho \log(1/\delta)}$.*

As in the previous proof, to show that $\Pr[L > \varepsilon'] \leq \delta \implies A$ is (ε', δ)-DP, we use Markov's inequality on a non-negative random variable Z, where $\Pr[Z \geq t] \leq \frac{E[Z]}{t}$.

Proof. Take 2 adj inputs x, x'

$$x \sim x'$$

$$Y \sim A(x)$$

Denote PLRV as $L = L_A^{x \to x'}(Y)$ (Y drawn from $A(x)$)

Now, by definition of zCDP,

$$\forall \lambda \geq 0, \mathbb{E}[\exp(\lambda L)] \leq \exp(\lambda(\lambda + 1)\rho)$$

i.e

$$\forall \lambda \geq 0, E[\exp(\lambda L)] \leq \exp(\lambda^2 \rho) \exp(\lambda \rho)$$

$$\Pr[L > \varepsilon'] \overset{\lambda \geq 0}{=} \Pr[\lambda L > \lambda \varepsilon'] = \Pr[e^{\lambda L} > e^{\lambda \varepsilon'}]$$

$$\leq e^{-\lambda \varepsilon'} E[e^{\lambda L}]$$

$$\leq \exp(-\lambda \varepsilon' + \lambda(\lambda + 1)\rho)$$

$$= \exp\left(-\lambda\left(\rho + 2\sqrt{\rho \log\left(\frac{1}{\delta}\right)}\right) + \lambda^2 \rho + \lambda \rho\right) \quad \text{(plugging in value of } \varepsilon')$$

$$= \exp\left(-2\lambda\sqrt{\rho \log\left(\frac{1}{\delta}\right)} + \lambda^2 \rho\right) \leq \delta$$

i.e.

$$-2\lambda\sqrt{\rho \log\left(\frac{1}{\delta}\right)} + \lambda^2 \rho \leq \log(\delta)$$

i.e.

$$2\lambda\sqrt{\rho \log\left(\frac{1}{\delta}\right)} - \lambda^2 \rho = \log\left(\frac{1}{\delta}\right)$$

i.e

$$\lambda^2 \rho - 2\sqrt{\lambda^2 \rho \log\left(\frac{1}{\delta}\right)} + \log\left(\frac{1}{\delta}\right) = 0$$

Rewriting the left side as it is a perfect square:

$$(\lambda\sqrt{\rho})^2 - 2\left(\lambda\sqrt{\rho}\right)\sqrt{\log(\frac{1}{\delta})} + \left(\sqrt{\log(\frac{1}{\delta})}\right)^2 = 0$$

i.e.

$$\left(\lambda\sqrt{\rho} - \sqrt{\log(\frac{1}{\delta})}\right)^2 = 0$$

i.e.

$$\lambda\sqrt{\rho} - \sqrt{\log(\frac{1}{\delta})} = 0$$

i.e.

$$\lambda\sqrt{\rho} = \sqrt{\log(\frac{1}{\delta})}$$

Solving for λ, we get:

$$\lambda = \sqrt{\frac{\log(1/\delta)}{\rho}}$$

Since $\delta \in (0, 1)$, we know that $\log(1/\delta) > 0$. Because $\rho > 0$, it follows that λ is a real, strictly positive number ($\lambda > 0$).

By finding a valid $\lambda > 0$ that makes the probability bound exactly equal to δ, we have shown that $Pr[L > \varepsilon'] \leq \delta$. Thus the algorithm is (ε', δ)-DP. □

5 Rényi Differential Privacy (RDP)

Rényi Differential Privacy (RDP) emerged from research into tracking privacy loss while training neural networks privately.

An algorithm A is (α, ε)-RDP if for all adjacent datasets $x \sim x'$, the Privacy Loss Random Variable (PLRV) $L = L_A^{x \to x'}(Y)$ satisfies the following bound:

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

For distributions P and Q, the Rényi divergence of order $\alpha > 1$ is defined as:

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

An algorithm is considered (α, ε)-RDP if $D_\alpha(A(x)||A(x')) \leq \varepsilon$ for all adjacent x, x'.

5.1 Comparison with zCDP

While zCDP controls concentration at all moments (similar to sub-Gaussian definitions), RDP provides specific control over individual moments.

- **Connection:** If an algorithm A is $(\alpha, \alpha\rho)$-RDP for all $\alpha \in [1, \infty)$, then A is ρ-zCDP.

5.2 Conversion to (ε', δ)-DP

If an algorithm A is (α, ε)-RDP, then for any $\delta > 0$, it is (ε', δ)-DP where:

$$\varepsilon' = \varepsilon + \frac{\log(1/\delta)}{\alpha - 1}$$

References

Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.