Accuracy of Gaussian mechanism:

Suppose $D = \{x^1, \ldots, x^n\}$ where each $x^i \in \mathbb{R}^k$ & $q_D = \frac{1}{n}\sum x^i \in \mathbb{R}^k$.

$\Delta_1^q = k/n$ & $\Delta_2^q = \frac{\sqrt{k}}{n}$

we output $z_D = q_D + N$

Laplace noise: $N \sim L(0, \frac{k}{\varepsilon n})$ to have ε-DP.

$$E[\| z_D - q_D \|_2] = O\left(\frac{k^{3/2}}{n\varepsilon}\right)$$

Gaussian noise: $N \sim \mathcal{N}(0, \frac{2k}{n^2\varepsilon^2}\log 1/\delta)$

$$E[\| z_D - q_D \|_2] = O\left(\frac{k}{\varepsilon n}\right)$$

* Gaussian mechanism leads to accuracy that is $O(k)$ better that of Laplace.

Properties of Approximate DP

1- Post-processing

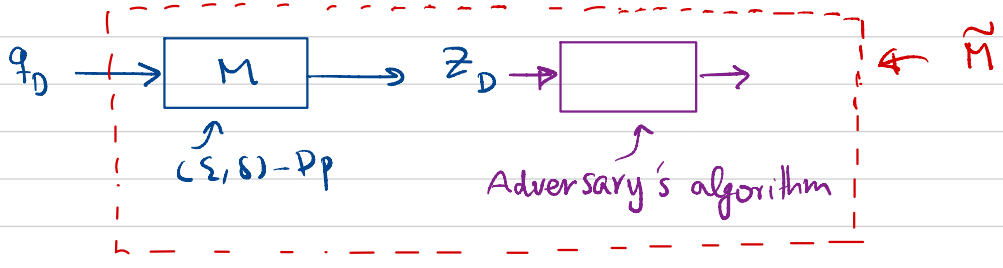

* Assumption: Adversary's algorithm doesn't have access to dataset.

If M is (ε, δ)-DP, then so is $\tilde{M}$.

In other words, approximate DP is closed under post-processing!

Proof. Similar as before, proof is a simple application of DP I.

2- Group privacy.

Theorem. If M is (ε, δ)-DP, then:

$$M_D(A) \leq e^{K\varepsilon} M_{D'}(A) + K e^{(K-1)\varepsilon} \cdot \delta \quad \forall \text{event } A$$

for any pair of datasets D & D' differing in K entries.

proof. Similar as before. construct $(K-1)$ pairwise neighboring dataset & apply the definition of approximate DP.

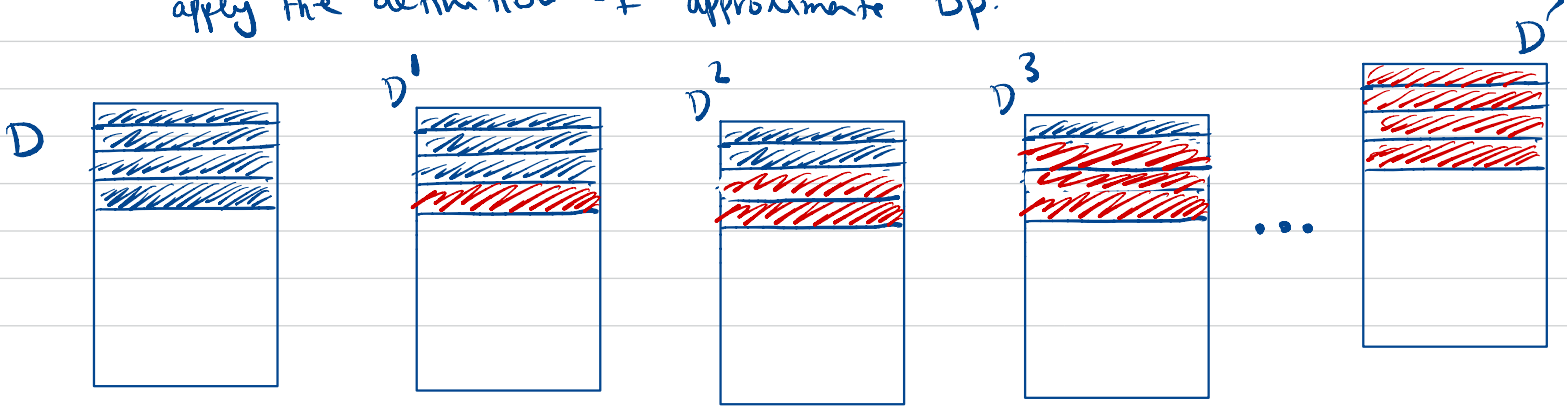

3- Basic composition:

* Adaptive or non-adaptive composition of k mechanism each of which is $(\varepsilon_i, \delta_i)$-DP is $\left(\sum_{i=1}^{k} \varepsilon_i, \sum_{i=1}^{k} \delta_i\right)$-DP.

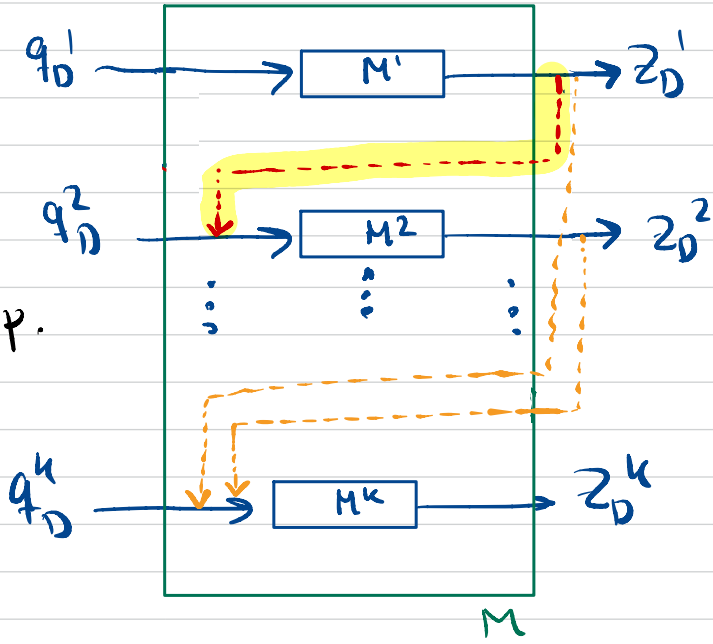

Proof. Dwork & Roth, Appendix B.

Basic composition looks rather identical to the basic composition for the pure DP. However, there is a fundamental difference.

* While basic composition is optimal for pure DP, it is indeed very far from being optimal for approximate DP.

There are several known attempts on improving (& even optimizing) composition results. Here, we introduce the oldest one.

Theorem (Advanced composition). Let M be ε_0-DP. The k-fold composition of M is (ε, δ)-DP for any $\delta \in (0,1)$ & $\varepsilon = \frac{k\varepsilon_0^2}{2} + \varepsilon_0\sqrt{2k \log \frac{1}{\delta}}$

Note that this theorem implies that composition of pure DP mechanism satisfies approximate DP; a phenomenon that can't be obtained from basic composition.

Proof. we need the following two lemmas:

1) Azuma Inequality: Let $Y_1, \ldots, Y_k$ be real-valued random variables such that $|Y_i| \leq \alpha$ & $E[Y_i \mid Y_1 = y_1, \ldots, Y_{i-1} = y_{i-1}] \leq \beta$.

Then, $$\Pr\left(\sum_{i=1}^{k} Y_i > k\beta + t\sqrt{k}\,\alpha\right) \le e^{-t^2/2}$$

* this type of results are typically called "concentration inequality" that aim to bound the probability of tail events

2) For any ε-DP mechanism:

$$D(M_D \| M_{D'}) \le (e^{\varepsilon} - 1)\,\varepsilon$$

In other words,

$$\left.\begin{aligned} E_{e^\varepsilon}(M_D \| M_{D'}) = 0 \\ E_{e^\varepsilon}(M_{D'} \| M_D) = 0 \end{aligned}\right\} \Rightarrow D(M_D \| M_{D'}) \le (e^{\varepsilon} - 1)\,\varepsilon$$

Proof of this result:

$$\begin{aligned} D(M_D \| M_{D'}) &\le D(M_D \| M_{D'}) + D(M_{D'} \| M_D) \\ &= \sum (M_{D'}(x) - M_D(x)) \log \frac{M_{D'}(x)}{M_D(x)} \\ &= \sum M_D(x) \cdot \underbrace{\left[\frac{M_{D'}(x) - M_D(x)}{M_D(x)}\right]}_{\le e^{\varepsilon} - 1} \cdot \underbrace{\log \frac{M_{D'}(x)}{M_D(x)}}_{\le \varepsilon} \\ &\le (e^{\varepsilon} - 1)\,\varepsilon \end{aligned}$$

With these two results at hand, we proceed with the proof of the advanced composition:

Consider two neighboring dataset $D \sim D'$ & let $Z_D^i \sim M_D^i$ & $Z_{D'}^i \sim M_{D'}^i$.

Define Privacy loss random variable for any given realization $z_1, \dots, z_{i-1}$,

$$Y_i := \log \frac{f_D^i(z_i \mid z_1, \dots, z_{i-1})}{f_{D'}^i(z_i \mid z_1, \dots, z_{i-1})} \quad \text{where } z_i \sim M_D^i$$

(f_D^i: pdf of M_D^i; $f_{D'}^i$: pdf of $M_{D'}^i$)

Since M^i is ε_0-Dp $\Rightarrow$ $|Y_i| \le \varepsilon_0$ w.p.1.

Note that we have

$$E[Y_i \mid Z_D^1 = z_1, \dots, Z_D^i = z_{i-1}] = D\left(P_{Z_D^i \mid z_1 \dots z_{i-1}} \,\|\, P_{Z_{D'}^i \mid z_1 \dots z_{i-1}}\right)$$

$$\le \varepsilon_0(e^{\varepsilon_0} - 1)$$

From the second lemma in previous page.

Note that $\sum Y_i$ is the privacy loss random variable for k-fold composition:

$$\log \frac{f_{Z_D^1 Z_D^2 \cdots Z_D^k}(z_1 \cdots z_k)}{f_{Z_{D'}^1 Z_{D'}^2 \cdots Z_{D'}^k}(z_1 \cdots z_k)} = \sum_{i=1}^{k} Y_i$$

where $(z_1 \cdots z_k) \sim$ joint density $f_{Z_D^1 \cdots Z_D^k}$

Recall that to prove that the k-fold composition is (ε, δ)-Dp we need this:

Any mechanism M is (ε, δ)-DP if $\Pr(Z_{D,D'} > \varepsilon) \leq \delta$

↑ PLRV

Thus, we need to say something like:

$$\Pr(\sum Y_i > \text{something}) \leq \delta$$

This is where Azuma Inequality kicks in.

Apply Azuma for $\alpha = \varepsilon_0$ & $\beta = \varepsilon_0(e^{\varepsilon_0} - 1)$

$$\Pr\Big[\sum Y_i > \underbrace{K\varepsilon_0(e^{\varepsilon_0}-1) + \varepsilon_0\overbrace{\sqrt{\log 1/\delta}\sqrt{k}}^{t}}_{\varepsilon}\Big] \le \delta$$

$\Rightarrow$ k-fold composition is (ε, δ)-DP with

$$\varepsilon = \underbrace{K\varepsilon_0(e^{\varepsilon_0}-1)}_{\varepsilon'} + \varepsilon_0\sqrt{k\log 1/\delta}$$

ε' can be approximated by $\varepsilon_0^2/2$

□

Comparison with basic composition:

Suppose we want to answer k queries of sensitivity one, with a given privacy guarantee (ε, δ)-DP.

Basic composition: k Laplace mechanism each of which $\frac{\varepsilon}{k}$-DP

so each adds $N \sim \mathcal{L}(0, k/\varepsilon)$.

Advanced composition: k Laplace mechanism each of which is $\left(\frac{\varepsilon}{\sqrt{2k \log 1/\delta}}\right)$-DP so each adds $N \sim \mathcal{L}\left(0, \sqrt{\frac{2k \log 1/\delta}{\varepsilon^2}}\right)$

* why?

Let M be $\frac{\varepsilon}{\sqrt{2k \log 1/\delta}}$-DP.

* The required variance is $O(\sqrt{k})$ smaller!

According to advanced composition, DP parameter of k-fold composition of M is

$$\frac{k}{2}\cdot\left(\frac{\varepsilon}{\sqrt{2k\log 1/\delta}}\right)^2+\left(\frac{\varepsilon}{\sqrt{2k\log 1/\delta}}\right)\sqrt{2k\log 1/\delta}=\frac{\varepsilon^2}{4\log 1/\delta}+\varepsilon\approx\varepsilon$$

Typical $\delta\approx 10^{-5}\Rightarrow \log 1/\delta\approx 15$

so $\frac{\varepsilon^2}{4\log 1/\delta}\approx\frac{\varepsilon^2}{60}$

In the advanced composition:

If $k\varepsilon_0^2$ is small, then we can ignore the first term.

M is ε_0-DP $\Rightarrow$ k-fold M is (ε,δ)-DP with $\delta\in(0,1)$

& $\varepsilon=\varepsilon_0\sqrt{2k\log 1/\delta}$

The previous advanced composition shows that k-fold composition of an ε_0-DP becomes an (ε, δ)-DP mechanism with ε increasing sub-linearly in k.

How about k-fold composition of $(\varepsilon_0, \delta_0)$-DP mechanism?

Here is a more general advanced composition:

Theorem (Advanced composition). Let M be $(\varepsilon_0, \delta_0)$-DP. The k-fold composition of M is (ε, δ)-DP where

$$\varepsilon = \frac{k\varepsilon_0^2}{2} + \varepsilon_0\sqrt{2k\log 1/\delta'}$$

$$\delta = k\delta_0 + \delta'$$

for any $\delta' \in (0,1)$.

Proof. The proof of this result relies on <u>some</u> neat reduction technique: Reduction of approximate DP to pure DP!

If $E_{e^{\varepsilon_0}}(P \| Q) \leq \delta_0$ and $E_{e^{\varepsilon_0}}(Q \| P) \leq \delta_0$

$\iff \exists\, P'$ & Q' such that

$E_{e^{\varepsilon_0}}(P' \| Q') = 0$ & $E_{e^{\varepsilon_0}}(Q' \| P') = 0$

&

$TV(P, P') \leq \frac{\delta_0}{e^{\varepsilon_0}+1}$ & $TV(Q, Q') \leq \frac{\delta_0}{e^{\varepsilon_0}+1}$

with this result, we can prove the advanced composition result using the previous composition result (for pure DP mechanisms).

□

Example: we wish to answer k adaptive queries each of sensitivity one.

Privacy budget $= (\varepsilon, \delta)$ & we want to use Gaussian mechanism:

According to advanced composition, k fold composition of an $(\varepsilon_0, \delta_0)$-DP Gaussian mechanism is $\left(\underbrace{\varepsilon_0\sqrt{2k\log 1/\delta'}}_{=\varepsilon}, \underbrace{k\delta_0 + \delta'}_{=\delta}\right)$-DP

* Note that here I ignore the first term in advanced composition!

$$k\delta_0 + \delta' = \delta \\ \delta' = \delta/2 \Rightarrow k\delta_0 = \delta/2 \Rightarrow \boxed{\delta_0 = \frac{\delta}{2k}}$$

$$\varepsilon_0\sqrt{2k\log\underbrace{1/\delta'}_{=2/\delta}} = \varepsilon \Rightarrow \varepsilon_0 = \frac{\varepsilon}{\sqrt{2k\log 2/\delta}}$$

$\Rightarrow$ Variance of Gaussian noise: $\frac{4k\log 2/\delta \log \frac{2k}{\delta}}{\varepsilon^2}$

As you can see, δ' is a design parameter. We need to pick δ' in a way that the resulting noise variance is minimal.

Another option for δ' in the above example is

$$\delta' = \delta_0$$

this choice results in variance $\frac{4K}{\varepsilon^2}\left(\log\frac{K+1}{\delta}\right)^2$ which is strictly worse than what we had before.

Rényi DP & composition results

Recall that differential privacy was intuitively defined by ensuring M_D & $M_{D'}$ are "close", according to HS divergence.

Why not other "distance" measures?

We already know that TV doesn't work.

It turns out that Rényi divergence is a good candidate.

Definition. (Mironov 2017) We say that a randomized mechanism M is (α, ε)-Renyi DP (or (α, ε)-RDP) for $\alpha > 1$ & $\varepsilon \geq 0$ if

$$D_\alpha(M_D \| M_{D'}) \leq \varepsilon$$

the main motivation [which made RDP extremely popular in AI] is the following result:

Theorem. For any $\alpha > 1$, we have

$$D_\alpha(P_X \| Q_X) + \min_x D_\alpha(P_{Y|X=x} \| Q_{Y|X=x}) \leq D_\alpha(P_{XY} \| Q_{XY}) \leq D_\alpha(P_X \| Q_X) + \max_x D_\alpha(P_{Y|X=x} \| Q_{Y|X=x})$$

In particular, if $P_{XY} = P_X \cdot P_Y$ & $Q_{XY} = Q_X \cdot Q_Y$, then:

$$D_\alpha(P_{XY} \| Q_{XY}) = D_\alpha(P_X \| Q_X) + D_\alpha(P_Y \| Q_Y).$$

Proof. See Lemma 2.2 of "Concentrated Differential Privacy" by Bun & Steinke.

this simple result leads to the following simple composition result that resembles basic composition.

Theorem. Let M^1 be (α, ε_1)-RDP & M^2 be (α, ε_2)-RDP. Then their (adaptive) composition is $(\alpha, \varepsilon_1 + \varepsilon_2)$-RDP.

Proof. Simple application the previous theorem.

According to this result, k-fold (adaptive) composition of any (α, ε)-RDP is $(\alpha, k\varepsilon)$-RDP.

Limitation of RDP. (α, ε)-RDP guarantee doesn't enjoy any operational interpretation, in terms of hypothesis testing.

In fact, it was shown* that D_α with $\alpha > 1$ has nothing to do with binary hypothesis testing performance.

* Read "Hypothesis testing interpretation & RDP" by Balle et al., AISTAT 2020.

How to fix it? RDP privacy guarantee needs to be converted back to approximate DP.

Theorem. If M is (α, ξ)-RDP with $\alpha > 1$, then it is (ε, δ)-DP with $\delta \in (0,1)$ & $\varepsilon = \xi + \frac{1}{\alpha - 1} \log \frac{1}{\delta}$.

Mironov

Moments accountant: was the state-of-the-art technique for studying composition of iterative algorithms used for training AI models (such as SGD). We discuss it further later.