Other ways of measuring DP

- Renyi DP
- Zero-concentrated DP
- Gaussian DP.

Main idea: easier to deal with than (ε, δ)-DP, better lecture composition.

- Convert back to (ε, δ)-DP for final guarantee, but internally work with alternate measurement.

Starting point: divergences between distributions

Remember from advanced composition lecture

Max-divergence: $D_\infty(P \| Q) = \max_{S \subseteq \text{Supp}(Q)} \left[\log \frac{P(S)}{Q(S)} \right]$

M is ε-DP iff $D_\infty(M(D), M(D')) \leq \varepsilon \quad \forall D \sim D'$

$\left(\log \left(\frac{\Pr[M(D) \in S]}{\Pr[M(D') \in S]} \right) \leq \varepsilon \iff \Pr[M(D) \in S] \leq e^\varepsilon \Pr[M(D') \in S] \right)$

KL-Divergence: $D(P \| Q) = \mathop{\mathbb{E}}_{x \sim P} \left[\ln \frac{P(x)}{Q(x)} \right]$

Rényi Divergence: (of order $\alpha > 1$):

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

Claim: - Limit as $\alpha \to \infty$ gives max-divergence!

- Limit as $\alpha \to 1$ gives KL-Divergence

Def: A mechanism M is (α, ε)-RDP if

$$D_\alpha(M(D), M(D')) \le \varepsilon \quad \forall D \sim D'$$

Fact: Rényi divergence is monotone w.r.t. α:

If $\alpha \ge \alpha'$, then $D_\alpha(P\|Q) \ge D_{\alpha'}(P\|Q)$

$\Rightarrow$ Since ε-DP is (∞, ε)-RDP, ε-DP $\Rightarrow (\alpha, \varepsilon)$-DP $\forall \alpha$.

So weaker than ε-DP. But what about compared to (ε, δ)-DP?

Idea: simpler and better!

Simpler - easier to analyze Gaussian mechanism

- Easier to prove composition!

Better: - Can convert to (ε, δ)-DP
- often gives stronger bounds! Used in Google's RAPPOR
- "Better" privacy guarantee: doesn't just give up on privacy when δ!

Interpretation:

If $D_\alpha(M(D) \| M(D'))$ is small, then while unlikely events can have larger probability difference, not unbounded!

Composition:

Thm: Let $f: \mathcal{D} \to R_1$ be (α, ε_1)-RDP,
$g: R_1 \times \mathcal{D} \to R_2$ be (α, ε_2)-RDP, then
mechanism that returns (X, Y) on input D where $X = f(D)$,
$Y = g(X, D)$
satisfies $(\alpha, \varepsilon_1 + \varepsilon_2)$-DP.

Pf: Let $h: \mathcal{D} \to R_1 \times R_2$ be sequential mechanism.
- Let X distribution for $f(D)$ X'
Y dist. of $g(X, D)$ Y' for D'
Z dist of (X, Y) Z'

$$\exp\left((\alpha-1)\, D_\alpha(h(D)\,\|\,h(D'))\right) = \mathop{E}_{h(D')}\left[\left(\frac{Z}{Z'}\right)^\alpha\right]$$

$$= \int_{R_1 \times R_2} \left(\frac{Z(x,y)}{Z'(x,y)}\right)^\alpha Z'(x,y)\, dx\, dy$$

$$= \int_{R_1 \times R_2} Z(x,y)^\alpha Z'(x,y)^{1-\alpha}\, dx\, dy$$

$$= \int_{R_1} \int_{R_2} \left(X(x)\,Y(x,y)\right)^\alpha \left(X'(x)\,Y'(x,y)\right)^{1-\alpha} dy\, dx$$

$$= \int_{R_1} X(x)^\alpha X'(x)^{1-\alpha} \left(\int_{R_2} Y(x,y)^\alpha Y'(x,y)^{1-\alpha}\, dy \right) dx$$

$$= \int_{R_1} X(x)^\alpha X'(x)^{1-\alpha} \cdot \mathop{E}_{y \sim Y(x)}\left[\left(\frac{Y(x,y)}{Y'(x,y)}\right)^\alpha\right] dx$$

$$= \int_{R_1} X(x)^\alpha X'(x)^{1-\alpha} \exp\left((\alpha-1)\, D_\alpha(g(x,D)\,\|\,g(x,D'))\right) dx$$

$$\le \int_{R_1} X(x)^\alpha X'(x)^{1-\alpha} dx \;\; \exp\left((\alpha-1)\cdot\varepsilon_2\right)$$

$$\le \exp\left((\alpha-1)\varepsilon_1\right) \cdot \exp\left((\alpha-1)\varepsilon_2\right)$$

$$= \exp\left((\alpha-1)(\varepsilon_1+\varepsilon_2)\right)$$

$$\Rightarrow D_\alpha(h(D)\,\|\,h(D')) \le \varepsilon_1+\varepsilon_2 \quad \checkmark$$

Thm: If f is (α, ε)-RDP, then it is

$$\left(\varepsilon + \frac{\log 1/\delta}{\alpha - 1}, \delta\right)\text{-DP} \quad \forall\, 0 < \delta < 1$$

Thm: If f has sensitivity 1, then Gaussian mechanism adding $N(0, \sigma^2)$ noise is $\left(\alpha, \frac{\alpha}{2\sigma^2}\right)$-RDP

$\forall\, \alpha > 1$

$\Rightarrow$ Can add noise without "precommitting" to δ!

In practice,

- analyzing via RDP then converting to (ε, δ)-DP better than using (ε, δ)-DP and advanced composition
- Most mechanisms don't need to commit to α ahead of time: afterwards can try diff. α's and see which gives best analysis

zCDP:

RDP a little complex: 2 parameters!

Intuition from Gaussian: $(\alpha, \rho\alpha)$ $\forall \alpha$: ε a linear fn of α!

Def: M is ρ-zCDP if

$$D_\alpha(M(D) \| M(D')) \leq \rho\alpha \quad \forall \alpha > 1$$

⇒ Gaussian mechanism is $\frac{1}{2\sigma^2}$-zCDP

much simpler!

Can still convert from ρ-zCDP → RDP → (ε, δ)-DP

- In practice, people tend to prefer RDP, since more precise, can choose α

- In theory, zCDP easier to deal with.

Thm: Let $f: \mathcal{D} \to R_1$ ρ_1-zCDP,

$g: R_1 \times \mathcal{D} \to R_2$ ρ_2-zCDP

⇒ composition is $(\rho_1 + \rho_2)$-zCDP