Lecture: Properties of ϵ-DP and Approximate DP

In this lecture, we shall discuss three key properties of ϵ-DP: (a) Composition, (b) Post-processing, and (c) Group privacy. We argue that DP degrades gracefully for all these operations.

A. Composition

Consider the setting where we cascade/compose one DP mechanism after another. Such a thought experiment sees practical realization often, since privacy-preserving mechanisms are designed in a modular fashion.

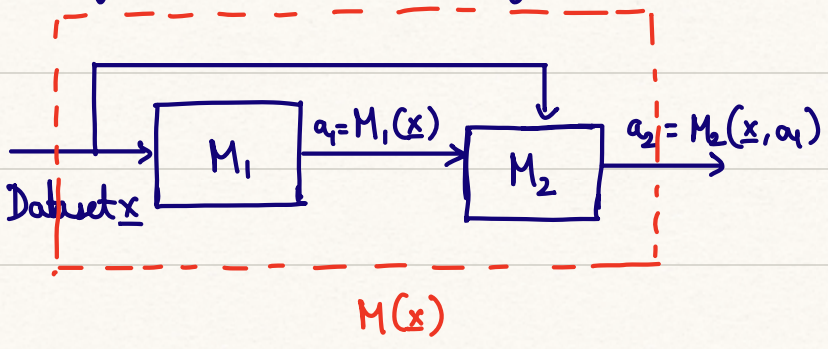

Claim: If M_1 is ϵ_1-DP and $M_2(\cdot, a)$ is ϵ_2-DP, for every a, then M is $(\epsilon_1+\epsilon_2)$-DP.

Proof: Fix neighbouring datasets $\underline{x}, \underline{x}'$. Now, note that (assuming the output spaces are discrete)

$$\begin{aligned}
&\mathbb{P}[M(\underline{x}) = (a_1, a_2)] \\
&\quad = \mathbb{P}[M_1(\underline{x}) = a_1] \cdot \mathbb{P}[M_2(\underline{x}, a_1) = a_2] \\
&\quad \leq e^{\epsilon_1} \cdot \mathbb{P}[M_1(\underline{x}') = a_1] \cdot e^{\epsilon_2} \cdot \mathbb{P}[M_2(\underline{x}', a_1) = a_2] \\
&\quad = e^{\epsilon_1+\epsilon_2} \cdot \mathbb{P}[M(\underline{x}') = (a_1, a_2)] \cdot \qquad \boxtimes
\end{aligned}$$

The above claim can be generalized further.

Theorem 1 (Basic Composition Theorem): Let $M_1, \dots, M_K$ be a sequence of mechanisms with $M_1 : \mathbb{D}_N \to \mathcal{Y}_1$, $M_2 : \mathbb{D}_N \times \mathcal{Y}_1 \to \mathcal{Y}_2, \dots, M_K : \mathbb{D}_N \times \mathcal{Y}_1^{K-1} \to \mathcal{Y}_K$ such that for every $i \in [K]$ and $(a_1, \dots, a_{i-1}) \in \mathcal{Y}_1 \times \dots \times \mathcal{Y}_{i-1}$, we have $M_i(\cdot, (a_1, \dots, a_{i-1}))$ is ϵ_i-DP. Then, the mechanism $M = M_K \circ M_{K-1} \circ \dots \circ M_1$ is $\left(\sum_{i=1}^{K} \epsilon_i\right)$-DP.

Remark: Loosely speaking, the composition of K Laplace mechanisms scales the amount noise by a factor of K.

We shall see a generalization of ϵ-DP, called (ϵ, δ)-DP, which achieves a scaling of $\approx \sqrt{K}$ upon composition.

B. Post-processing

The post-processing of the output of a DP mechanism is simply the task of acting on this output by any deterministic function which is <u>independent</u> of the dataset.

Theorem 2 (Post-processing): Let $M : \mathbb{D}_N \to \mathcal{Y}$ be an ϵ-DP mechanism

and $g: \mathcal{Y} \to \mathcal{Z}$ be any deterministic function. Then, we have that $g \circ M$ is ϵ-DP.

Proof: Fix $\underline{x} \sim \underline{x}'$ and $S \subseteq \mathcal{Y}$. Let $T \triangleq \{y \in \mathcal{Y} : g(y) \in S\} = g^{-1}(S)$. Then,

$$\begin{aligned} \mathbb{P}\left[f(M(\underline{x})) \in S\right] &= \mathbb{P}\left[M(\underline{x}) \in T\right] \\ &\leq e^{\epsilon} \mathbb{P}\left[M(\underline{x}') \in T\right] \\ &= e^{\epsilon} \mathbb{P}\left[f(M(\underline{x}')) \in S\right]. \end{aligned}$$

$\blacksquare$

C. Group privacy

The following simple theorem holds:

Theorem 3 (Group privacy): Let $\underline{x}, \underline{x}'$ be datasets that differ in at most k samples. Then, for all $S \subseteq \mathcal{Y}$, we have that for an ϵ-DP mechanism M,

$$\mathbb{P}\left[M(\underline{x}) \in S\right] \leq e^{k\epsilon} \cdot \mathbb{P}\left[M(\underline{x}') \in S\right].$$

Proof: Simple exercise via "daisy".

Aside: Parallel composition

Consider the setting where a dataset $\mathcal{D} \in \mathbb{D}_N$ is chopped up/partitioned into smaller datasets. In particular, let $\mathcal{D} = (x_1, \ldots, x_N)$ and let

$\mathcal{D}_1 = x_{S_1}$, $\mathcal{D}_2 = x_{S_2}$, ..., $\mathcal{D}_K = x_{S_K}$, where $x_S \triangleq \bigcup_{i \in S} [x_i]$ and $\bigcup_{j=1}^{K} S_j = [N]$. Furthermore, we assume that the partitioning above is dataset-independent.

The following theorem then holds.

Theorem 4 (Parallel Composition Theorem): Given ϵ_i-DP mechanisms M_i, $i \in [K]$, acting on the samples in the subsets S_i, respectively, we have that the mechanism $M = (M_1, \ldots, M_K)$ satisfies $(\max_{i=1}^{K} \epsilon_i)$-DP.

Proof: Changing a sample in a dataset affects at most one subset in $S_1, \ldots, S_K$. ▨

Aside: "Add/Remove" model of DP

The definition of DP we have seen until now requires the indistinguishability of mechanism outputs on "neighbouring datasets," which are obtained by swapping a sample value with a different value, while keeping the total # samples fixed.

Def 1: Two datasets $\underline{x}$ and $\underline{x}'$ are said to be neighbouring in the "add/remove" model if $\underline{x}' = \underline{x}_{-i}$, for some $i \in [N]$, or $\underline{x} = \underline{x}'_{-i}$, for some $i \in [N]$.

Lemma 1: A mechanism M satisfying ϵ-DP in the "add/remove" model of DP obeys 2ϵ-DP in the "swap" model.

Proof: $\underline{x} \xrightarrow{\epsilon\text{-privacy loss}} \underline{x}_{\sim i} \longrightarrow \underline{x}_{\sim i} \cup \{x'_i\} = \underline{x}'$. ▨

($\underline{x}_{\sim i} \longrightarrow \underline{x}'$: ϵ-privacy loss.)

Approximate DP

We now discuss a relaxation of ϵ-DP, via the notion of "approximate DP". Such a notion obeys properties very similar to those of ϵ-DP, with significant strengthening in some.

Def 2 ((ϵ,δ)-DP): A mechanism $M : \mathbb{D}_N \to \mathcal{Y}$ is (ϵ,δ)-DP, for $\delta \in (0,1)$, if for all $S \subseteq \mathcal{Y}$ and all $\underline{x} \sim \underline{x}'$, we have

$$\mathbb{P}\left[M(\underline{x}) \in S\right] \leq e^{\epsilon} \cdot \mathbb{P}\left[M(\underline{x}') \in S\right] + \delta .$$

Remark: (i) We hence have that

$$e^{-\epsilon}\left(\mathbb{P}\left[M(\underline{x}') \in S\right] - \delta\right) \leq \mathbb{P}\left[M(\underline{x}) \in S\right] \leq e^{\epsilon} \cdot \mathbb{P}\left[M(\underline{x}') \in S\right] + \delta$$

(ii) In practice, one puts in the restriction that $\delta \leq 1/N$, to disallow certain "semantically non-private" mechanisms.