CSC2412: Properties of Differential Privacy & More Mechanisms

Sasho Nikolov

Review

Data model

Data set: (multi-)set X of n data points $X = \{x_1, \ldots, x_n\}$.

- each data point (or row) x_i is the data of *one person*
- each data point comes from a *universe* $\mathcal{X}$ e.g. $\mathcal{X} = \{0,1\}^d$

We call two data sets X and X' *neighbouring* if

1. (*variable* n) we can get X' from X by adding or removing an element
2. (*fixed* n) we can get X' from X by replacing an element with another

we will mostly use this

$X \sim X' \iff X, X'$ neighbouring

Differential Privacy

Definition
A mechanism $\mathcal{M}$ is ε-*differentially private* if, for any two neighbouring datasets X, X', and any set of outputs $S \subseteq \text{Range}(\mathcal{M})$

$$\mathbb{P}(\mathcal{M}(X) \in S) \leq e^{\varepsilon}\mathbb{P}(\mathcal{M}(X') \in S).$$

Basic Properties

Composition motivation

It would be nice if we can:

- Post-process outputs of DP algorithms without losing privacy.

E.g. average $\frac{(e^{\varepsilon}+1)y_i - 1}{e^{\varepsilon}-1}$ for the output $(y_1, \dots, y_n)$ of RR

- Build complex DP algorithms from simple ones.

E.g. use RR to answer many counting queries

- Allow an analyst to adaptively choose queries to ask

E.g. "smokers?" $\xrightarrow{\geq 25\%}$ "smokers are under 25 yrs old?"

$\downarrow$ $<25\%$

. . . .

Composition theorem

Suppose

- $\mathcal{M}_1(\cdot)$ is ε_1-DP
- $\mathcal{M}_2(\cdot, y)$ is ε_2-DP for any y in the range of $\mathcal{M}_1$

$\mathcal{M}_1$ takes X

$\mathcal{M}_2$ takes X and the output of $\mathcal{M}_1$

Then $\mathcal{M}(\cdot)$ given by $\mathcal{M}(X) = \mathcal{M}_2(X, \mathcal{M}_1(X))$ is $(\varepsilon_1 + \varepsilon_2)$-DP.

Epsilons add up

$\mathcal{M}_3(\cdot, z)$ ε_3-DP $\forall$ $z \in$ Range$(\mathcal{M}_2)$

$\mathcal{M}_3(X, \mathcal{M}_2(X, \mathcal{M}_1(X)))$ is $(\varepsilon_1+\varepsilon_2+\varepsilon_3)$-DP

and so on...

Post-processing

If $\mathcal{M}_2$ is 0-DP

i.e. $\mathcal{M}_2$ is only a function of the output of $\mathcal{M}_1$

then $\mathcal{M}_2(\mathcal{M}_1(X))$ is ε_1-DP

Proof of the composition theorem

Take some $X \sim X'$, $S \subseteq \text{Range}(\mathcal{M}_2)$

To prove: $\mathcal{M}(X) = \mathcal{M}_2(X, \mathcal{M}_1(X))$ is $(\varepsilon_1 + \varepsilon_2)$-DP

$$\mathbb{P}(\mathcal{M}(X) \in S) = \sum_{y \in \text{Range}(\mathcal{M}_1)} \mathbb{P}(\mathcal{M}_2(X, y) \in S) \cdot \mathbb{P}(\mathcal{M}_1(X) = y)$$

$$\leq \sum_{y \in \text{Range}(\mathcal{M}_1)} e^{\varepsilon_2}\, \mathbb{P}(\mathcal{M}_2(X', y) \in S) \cdot e^{\varepsilon_1}\, \mathbb{P}(\mathcal{M}_1(X') = y)$$

$$= e^{\varepsilon_1 + \varepsilon_2} \sum_{y \in \text{Range}(\mathcal{M}_1)} \mathbb{P}(\mathcal{M}_2(X', y) \in S) \cdot \mathbb{P}(\mathcal{M}_1(X') = y)$$

$$= e^{\varepsilon_1 + \varepsilon_2} \cdot \mathbb{P}(\mathcal{M}(X') \in S)$$

Group Privacy

What protection is offered to small groups rather than individuals?

- E.g., what can an adversary find out about my immediate family?

$X = \{x_1, x_2, \dots, x_i, \dots, x_j, \dots, x_n\}$

$X' = \{x_1, \dots, x_i', \dots, x_j', \dots, x_n\}$ $\Rightarrow$ 2-neighbouring

Definition
Two data sets X, X' are t-neighbours if they differ in the data of $\leq t$ individuals.

For any ε-DP mechanism $\mathcal{M}$, any t-neighbours X, X', and any set S of outputs

$$\mathbb{P}(\mathcal{M}(X) \in S) \leq e^{t\varepsilon}\mathbb{P}(\mathcal{M}(X') \in S).$$

Proof of group privacy property

X, X' t-neighbouring $\Rightarrow \exists\ X^0 = X, X^1, X^2, \ldots, X^t = X'$

$X^0 \sim X^1, X^1 \sim X^2, \ldots, X^{t-1} \sim X^t$

E.g. $= X^0$

$X = \{x_1, x_2, \ldots, x_i, \ldots, x_j, \ldots, x_n\}$

$X' = \{x_1, \ldots \quad, x_i', \ldots, x_j', \ldots, x_n\}$

$" X^2$

$X^1 = \{x_1, x_2, \ldots, x_i', \ldots, x_j, \ldots, x_n\}$

$\forall S$ set of outputs

$$\mathbb{P}(\mathcal{M}(X) \in S) \leq e^{\varepsilon}\, \mathbb{P}(\mathcal{M}(X^1) \in S) \leq e^{2\varepsilon}\, \mathbb{P}(\mathcal{M}(X^2) \in S)$$

$$\ldots \leq e^{t\varepsilon}\, \mathbb{P}(\mathcal{M}(X') \in S)$$