CS 355 Lecture 9 : Differential Privacy

Last week:

- MPC
- PIR

Today: "MPC is not enough"

- Brief recap on MPC and PIR
  - ↳ Does MPC "imply" PIR? (no)
  - ↳ Does MPC protect "privacy"? (no)
- Intro to differential privacy
  - Defining privacy
  - The Laplace mechanism
  - Composition & post-processing

HW 3 is out!

PIR:

index into X

Client $(i, 1^\lambda)$ $\xrightarrow{q}$ Server $(X, 1^\lambda)$

$\xleftarrow{a}$

$X_i \leftarrow \text{Recon}(a)$

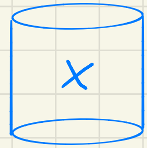

Two flavors:

k-server
- non-collusion assumption
- information theoretic security

Single-server
- cryptographic assumption
- computational security

Question: Does MPC "imply" PIR?

e.g., given a maliciously-secure 2PC protocol for arbitrary functions f, can you construct single-server PIR?

Answer: Not necessarily!

Why: MPC says we can compute the function $f(i, X) = x_i$ privately and efficiently.

client's input (i); server's input (X)

efficiently:
- ↳ parties run in poly(λ) time
- ↳ Communication is poly(λ)

MPC protocol doesn't guarantee that communication is $o(|X|)$

PIR is an example of a cryptographic primitive with a stricter notion of "efficient" than "everything should be poly(λ)". Many interesting primitives in crypto have such requirements (e.g. Fully Homomorphic Encryption) and a lot of recent research tries to make MPC and ZK proofs that are efficient in a practical sense.

MPC and privacy:

We want to know if smoking causes cancer

Study: collect (s, c) from n participants and compute

$s \in \{0,1\}$ "smoker" $c \in \{0,1\}$ "has cancer"

# of smokers that have cancer

$$y = \frac{\sum_{i=1}^{n} c_i \wedge s_i}{\sum_{i=1}^{n} s_i}$$

# of smokers

	Smokes	has cancer
Alice	?	?
Bob	✓	✓
Charlie	×	✓
Daisy	✓	×

Adversary corrupted these parties (Bob, Charlie, Daisy)

MPC protocol secure against $n-1$ corruptions $\Longrightarrow$ $\boxed{y = 2/3}$

↳ does Alice smoke?
does Alice have cancer?

- Is this really "private"?
- Do we really need to know the exact value of y to understand the link between smoking and cancer?

↳ there's sampling noise anyway so no!

Private data analysis

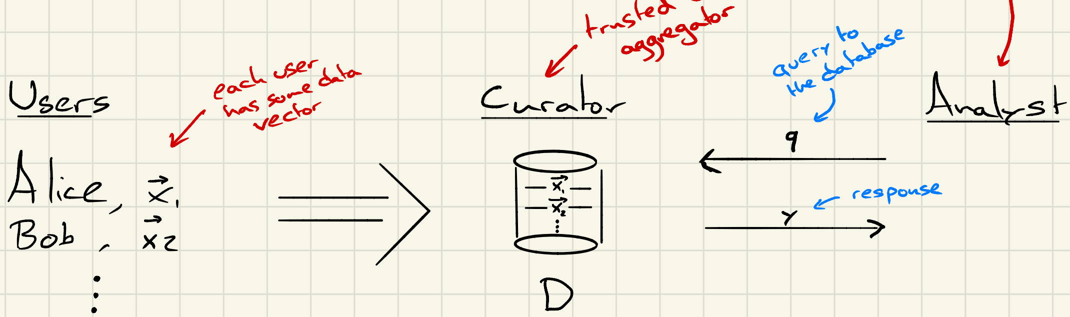

Strong privacy notion (Dalenius, 1977):

"The analyst learns nothing about Alice that it couldn't have learned without the database D"

This is too strong! E.g. if the Analyst knows that Alice smokes, and it learns from D that "smoking increases cancer risk", it has learned something about Alice (she is more likely to have cancer).

↳ this remains true even if Alice is not in the database!

↳ the only way to satisfy this notion is if the Analyst learns nothing at all!

Differential privacy

"The analyst learns nothing about Alice that it couldn't also have learned if Alice was not in the database"

Differential privacy is a promise [Dwork, Roth 2014]

Whether Alice agrees to give her data to the curator or not has no influence on what an adversary can learn!
(so she might as well give her data for the advancement of Science!)

Note: the def. prevents the analyst from learning individual facts about Alice but not from learning "population-level" facts such as "smoking causes cancer".

Definition: Two databases D, D' are neighboring, denoted $D \sim D'$, if $|D| = |D'|$ and D, D' differ in a single row.

$$D: \begin{array}{|c|} \hline \text{—}\ \vec{x}_1\ \text{—} \\ \cdots \\ \text{—}\ \vec{x}_i\ \text{—} \\ \cdots \\ \text{—}\ \vec{x}_n\ \text{—} \\ \hline \end{array} \qquad\qquad D': \begin{array}{|c|} \hline \text{—}\ \vec{x}_1\ \text{—} \\ \cdots \\ \text{—}\ \vec{x}_i'\ \text{—} \\ \cdots \\ \text{—}\ \vec{x}_n\ \text{—} \\ \hline \end{array}$$

Let M be the curator's (randomized) algorithm for answering the analyst's queries, i.e.

$$M : \underbrace{\mathcal{X}^n}_{\text{space of n-row databases}} \times \underbrace{Q}_{\text{query space}} \to \underbrace{\mathcal{Y}}_{\text{output range}} \qquad\qquad M(\underbrace{D}_{\text{database}}, \underbrace{q}_{\text{query}}) = \underbrace{y}_{\text{answer}}$$

Definition [ε-differential privacy, Dwork-McSherry-Nissim-Smith 2006]:

An algorithm M is ε-DP if for every pair of neighboring databases D, D', every query $q \in Q$ and every event $S \subseteq \mathcal{Y}$:

$$\mathbb{P}\big[M(D,q) \in S\big] \le e^{\varepsilon} \cdot \mathbb{P}\big[M(D',q) \in S\big]$$

Remarks

- Any "bad" event when Alice is in the DB would have happened with similar probability if Alice was not in the DB
- Think of $\varepsilon > 0$ as a small constant
  ↳ Why can't ε be negligible (say in $|D|$)? [Homework]

Q: Can any query be answered with differential privacy (and some utility)?

A: No!

Example:

database of salaries

query

$$P\left[M\left(\boxed{\begin{array}{l}\text{Henry } 1\$ \\ \text{Florian } 1\$ \\ \text{Dima } 1\$\end{array}}, \text{"max salary"}\right) \in S\right] \approx P\left[M\left(\boxed{\begin{array}{l}\text{Bill Gates } 1M\$ \\ \text{Florian } 1\$ \\ \text{Dima } 1\$\end{array}}, \text{"max salary"}\right) \in S\right]$$

If M gives an accurate answer with high Probability for one of the two databases, it must give a very inaccurate answer with roughly the same probability for the other database.

The Laplace Mechanism: ε-DP for low-sensitivity queries

We will show how to get differential privacy for real-valued queries $q : X^n \to \mathbb{R}$. For example, "counting queries":

$D =$

name	smoker
Alice	1
Bob	1
Charlie	0
Daisy	1

$q =$ # of smokers

$$q(D) = 3$$

<u>Definition (sensitivity)</u>: For a query $q : \mathcal{X}^n \to \mathbb{R}$, the sensitivity of q is $\Delta q = \max_{D \sim D'} |q(D) - q(D')|$

Q: What is the sensitivity of a counting query?

A: $\boxed{1}$

Q: What is the sensitivity of the "maximum salary" query?

A: $\boxed{\text{unbounded}}$

Definition (Laplace distribution):

The centered Laplace distribution with parameter b, $Lap(b)$, has density $f_{Lap(b)}(z) = \frac{e^{-|z|/b}}{2b}$

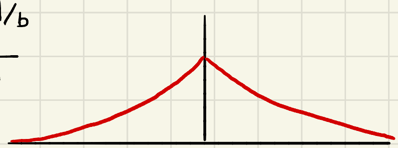

mean $= 0$
Variance $= 2b^2$

Laplace Mechanism $M_L(D, q)$
1. Compute $q(D)$
2. Sample $v \sim Lap\left(\frac{\Delta q}{\varepsilon}\right)$ — query sensitivity
3. Output $q(D) + v$

- M_L is ε-DP

Proof: For any $D \sim D'$, $y \in \mathbb{R}$ and query q, let $b = \frac{\Delta q}{\varepsilon}$. Then,

$$\frac{\mathbb{P}[M_L(D,q)=y]}{\mathbb{P}[M_L(D',q)=y]} = \frac{\mathbb{P}_{v \sim Lap(b)}[v = y - q(D)]}{\mathbb{P}_{v \sim Lap(b)}[v = y - q(D')]} = \frac{\frac{1}{2b} \cdot e^{-|y-q(D)|/b}}{\frac{1}{2b} \cdot e^{-|y-q(D')|/b}}$$

$$= e^{\frac{\varepsilon}{\Delta q} \cdot (|y - q(D')| - |y - q(D)|)}$$

reverse triangular inequality $\quad |x| - |y| \leq |x - y|$

$$\leq e^{\frac{\varepsilon}{\Delta q} \cdot |q(D') - q(D)|}$$

def. of sensitivity

$$\leq e^{\frac{\varepsilon}{\Delta q} \cdot \Delta q}$$

$$= e^{\varepsilon} \qquad \square$$

1000-words:

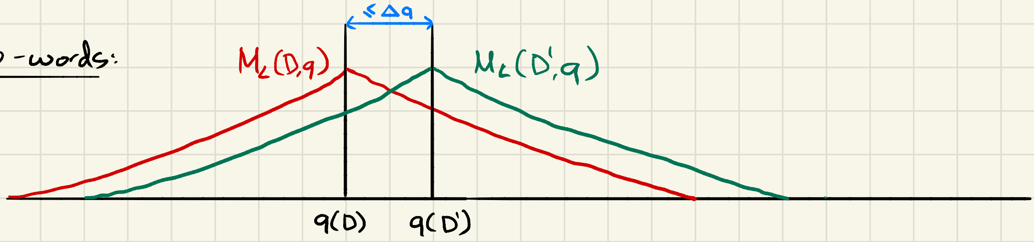

For any value returned by M_L on D, M_L would have returned the same value with approx. the same probability on D'.

- M_L is accurate :

$$\forall \beta > 0, \quad \mathbb{P}\left[\left| M_L(D,q) - q(x) \right| > \frac{\Delta q}{\varepsilon} \cdot \ln\left(\frac{1}{\beta}\right) \right] \leq \beta$$

Example: for a counting query ($\Delta q = 1$), if M_L is ε-DP for $\varepsilon = 0.1$, then with 99% probability, the error in the counting query will be less than $\frac{1}{0.1} \cdot \ln\left(\frac{1}{0.01}\right) \simeq 46$

Note that the noise, and thus error, are independent of the size of the database.

Proof: Follows from the following standard concentration inequality for the Laplace distribution:

$$\mathbb{P}_{V \sim Lap(b)}\left[|V| > c \cdot b \right] < e^{-c} \qquad \text{for any constant } c$$

Properties of differential privacy

- Post-processing: Let $\mathcal{M} : \mathcal{X}^n \times Q \to \mathcal{Y}$ be ε-DP and let $f : \mathcal{Y} \to Z$ be any (randomized) function. Then $(f \circ \mathcal{M}) : \mathcal{X}^n \to Z$ is ε-DP.

  function composition

Proof (for deterministic f):

Fix any neighboring databases D, D', query q and event $S \subseteq Z$.
Let $T = \{ y \in \mathcal{Y} : f(y) \in S \}$. Then:

$$
\begin{aligned}
\mathbb{P}[f(\mathcal{M}(D,q)) \in S] &= \mathbb{P}[\mathcal{M}(D,q) \in T] \\
&\leq e^{\varepsilon} \cdot \mathbb{P}[\mathcal{M}(D',q) \in T] \\
&= e^{\varepsilon} \cdot \mathbb{P}[f(\mathcal{M}(D',q)) \in S]
\end{aligned}
$$

$\square$

Remark: Whatever the analyst does with the answers to the queries, DP is guaranteed !

• Composition: Let $M_1, M_2, \dots, M_n$ be algorithms where $M_i : \mathcal{X}^n \times Q \to \mathcal{Y}_i$ is ε_i-DP.

Then $M(D,q) \mapsto (M_1(D,q), M_2(D,q), \dots, M_n(D,q))$ is ε-DP for $\varepsilon = \sum_{i=1}^n \varepsilon_i$.

Proof: Fix any $D \sim D'$, query Q and $(y_1, \dots, y_n) \in \underbrace{\mathcal{Y}_1 \times \dots \times \mathcal{Y}_n}_{\text{range of } M}$

$$
\begin{aligned}
\mathbb{P}\left[M(D,q) = (y_1, \dots, y_n)\right] &= \prod_{i=1}^n \mathbb{P}\left[M_i(D,q) = y_i\right] \\
&\le \prod_{i=1}^n e^{\varepsilon_i} \mathbb{P}\left[M_i(D',q) = y_i\right] \\
&= e^{\sum_{i=1}^n \varepsilon_i} \cdot \mathbb{P}\left[M(D',q) = (y_1, \dots, y_n)\right]
\end{aligned}
$$

$\square$

Remark: The more queries are answered, the less privacy remains!