

CS459/689

Privacy, Cryptography,
Network and Data Security

Differential Privacy

Intended Learning Outcomes

By the end of this lecture, you should be able to:

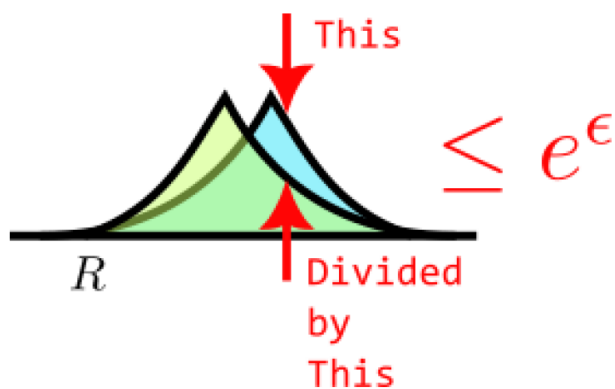
- **Describe the different properties of differential privacy such as composition and post processing.**
- **Apply the Laplace mechanism to simple statistics problems.**
- **Use randomized response to collect and analyze binary data.**
- **Discuss how the exponential mechanism can be applied to discrete problems.**

Recall: Differential privacy

Differential Privacy

A mechanism $M: \mathcal{D} \rightarrow \mathcal{R}$ is ϵ -differentially private (ϵ -DP) if the following holds for all possible sets of outputs $R \subset \mathcal{R}$ and all pairs of neighboring datasets $D, D' \in \mathcal{D}$:

$$\Pr(M(D) \in R) \leq \Pr(M(D') \in R) e^\epsilon$$




Differential Privacy Settings


Central DP vs. Local DP

- Depending on who runs the mechanism, there are two broad models for differential privacy.

Central Differential Privacy: there is a centralized (trusted) aggregator



Local Differential Privacy: each user runs the mechanism themselves and reports the result to the adversary/analyst




Central DP vs. Local DP

(Central) Differential Privacy

A mechanism $M: \mathcal{D} \rightarrow \mathcal{R}$ is ϵ -differentially private (ϵ -DP) if the following holds for all possible sets of outputs $R \subset \mathcal{R}$ and all pairs of neighboring datasets $D, D' \in \mathcal{D}$:


$$\Pr(M(D) \in R) \leq \Pr(M(D') \in R) e^\epsilon$$



(Local) Differential Privacy

A mechanism $M: \mathcal{D} \rightarrow \mathcal{R}$ is ϵ -differentially private (ϵ -DP) if the following holds for all possible sets of outputs $R \subset \mathcal{R}$ and all pairs of neighboring inputs $x, x' \in \mathcal{D}$:

$$\Pr(M(x) \in R) \leq \Pr(M(x') \in R) e^\epsilon$$



- They are “the same definition”, it’s just that the inputs to the mechanism and what we define as “neighbouring” inputs/datasets is usually different.

Central DP vs. Local DP

- **Central DP**

- Best accuracy, aggregation allows to hide in the crowd before we add noise.
- Need to trust the data collector.
- Hard to verify if noise was added.

- **Local DP**

- Accuracy not as good. Each user adds noise which can compound in the final result.
- User doesn't need to trust anybody and knows they added noise.

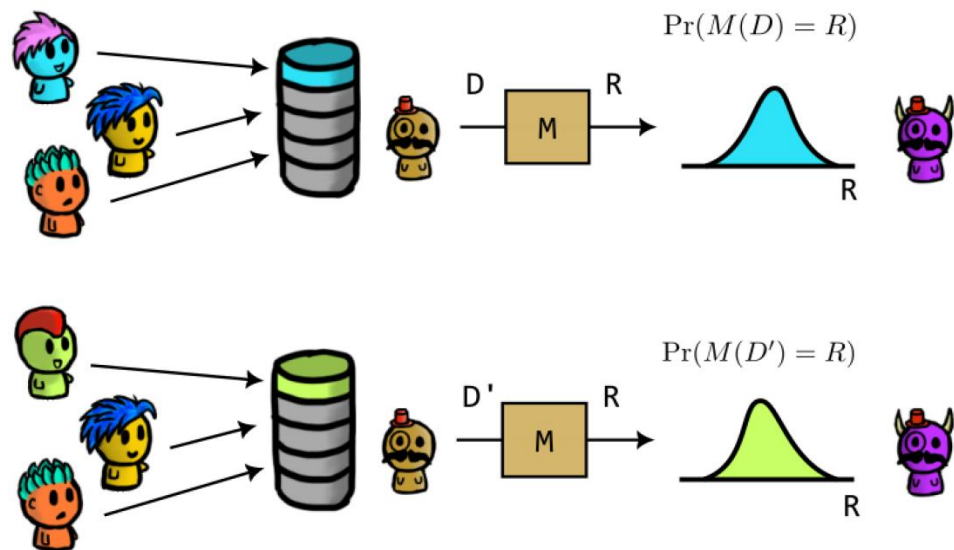
- **Shuffle Model of DP**

- Hybrid where users add less noise on the understanding a semi-trusted party aggregates and shuffles the results before they are made public.

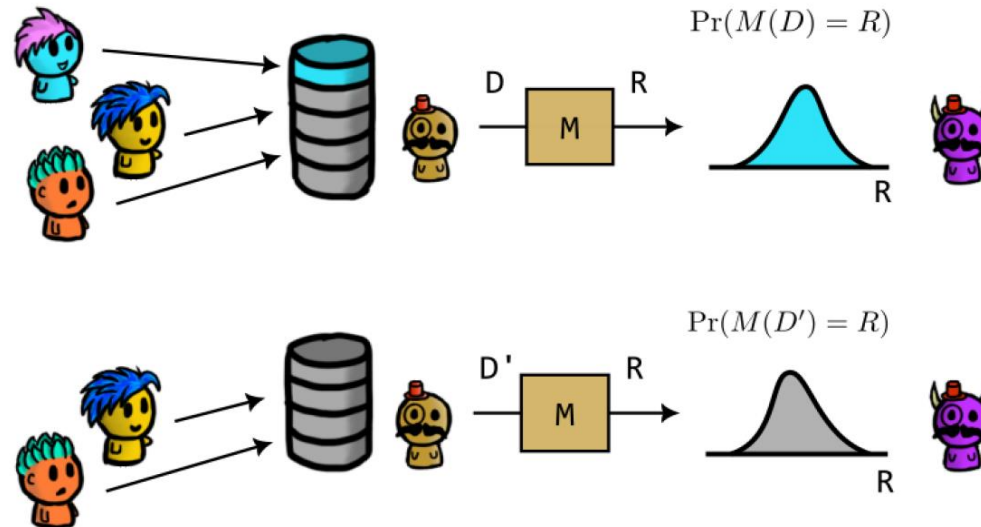
Bounded DP vs. Unbounded DP

- There are two “main” definitions for how we define neighboring datasets in the central model.

Bounded DP: D and D' have the same number of entries but differ in the value of one.




Unbounded DP: D and D' are such that you get one by deleting an entry from the other one.



Other notions of DP

- Many possible neighbouring definitions.
- For example, in location privacy:



Depending on how we define neighboring datasets D and D' , we get a different DP guarantee:

- User-level DP: we replace a user trajectory for another user's trajectory
 - Event-level DP: we replace the location of a user for another location
 - w -event DP: we replace a window of w consecutive locations of a user for another
- These are all DP and have their uses. It is important to understand, for each system/application, which notion of DP it provides.

DP Mechanisms

DP Mechanisms

- We are going to see different mechanisms that provide Differential Privacy and that can be applied to various systems.
- You need to understand why they provide DP, when you can use them, how to compute the ϵ level they provide, etc.
- We will see:
 1. The Laplace Mechanism (DP, continuous outputs)
 2. The Randomized Response Mechanism (DP, binary inputs/outputs)
 3. The Exponential Mechanism (DP, discrete outputs)
 4. The Gaussian Mechanism (approximate DP, continuous)
 5. General Discrete Mechanisms

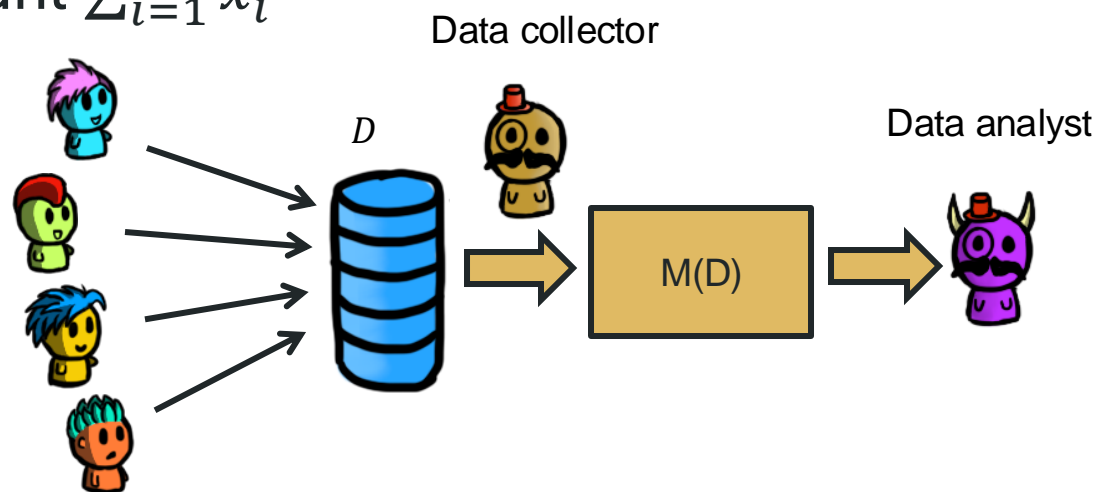
Why do we need DP mechanisms?

- Many cases where sensitive information can be of great benefit to society
 - Analysis of healthcare records
 - Statistics computed from the census
- Without proper protection we have learnt there are many inference attacks.
- DP mechanisms allow us to tune the privacy utility trade-off to still benefit from the sensitive data while providing privacy guarantees.

Example DP mechanism


- The dataset contains health data from n users, and the data analyst wants to know how many patients have tested positive for a virus
- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- Consider an analyst wants to report the count $\sum_{i=1}^n x_i$

Q: How could we make this private?



Example: the Laplacian mechanism

- Let $Y \sim \text{Lap}(b, \mu)$
 - A Laplace distribution!
- With PDF: $p_Y(y) = \frac{1}{2b} e^{-\frac{|y-\mu|}{b}}$



- Consider the mechanism that reports the true count of positive results plus Laplacian noise, i.e.,
 - $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .

Example: the Laplacian mechanism

- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .
- You can write $c = \sum_{i=2}^n x_i$.

Q: What do the worst-case distributions of $M(D)$ vs $M(D')$ look like?


Example: the Laplacian mechanism

- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .
- You can write $c = \sum_{i=2}^n x_i$.

Q: What do the worst-case distributions of $M(D)$ vs $M(D')$ look like?

Q: What is the maximum ratio between the distributions?

A:



Example: the Laplacian mechanism

- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .
- You can write $c = \sum_{i=2}^n x_i$.


Q: What do the worst-case distributions of $M(D)$ vs $M(D')$ look like?

Q: What is the maximum ratio between the distributions?

A: $\exp(1/b)$...




Let $b = 1/\epsilon$ and we have DP!

A:



Approximate DP

- Differential privacy is **very strict**. In the slide before, if we replace the Laplacian noise with a Laplace $y \sim Lap(1)$ truncated at $y > 1000$, the mechanism is basically “the same”:
 - $\Pr(y > 1000 | y \sim Lap(1)) = \frac{1}{2} \exp(-1000) \approx 10^{-435}$.
- However, if we truncate the Laplacian noise, the mechanism goes from $\epsilon = 1$ (good privacy) to $\epsilon = \infty$ (no privacy).



No matter where we do zoom, we'll always see this!

Approximate DP

- The following is a relaxation of the DP definition, that allows some tolerance:


(Approximate) Differential Privacy

A mechanism $M: \mathcal{D} \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private (ϵ, δ) -DP if the following holds for all sets of possible outputs $S \subset \mathcal{R}$ and all pairs of neighboring datasets $D, D' \in \mathcal{D}$:


$$\Pr(M(D) \in S) \leq \Pr(M(D') \in S) e^\epsilon + \delta$$

- When $\delta = 0$, this is the same as ϵ -DP (called pure DP).
- What does this mean?

We have two distributions
 $f(R|D)$ vs $f(R|D')$



We multiply one
(e.g., blue) by e^ϵ



The area of the green one not covered by
the blue one now will be $\leq \delta$



Approximate DP: interpretation

(Approximate) Differential Privacy

A mechanism $M: \mathcal{D} \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private $((\epsilon, \delta)$ -DP) if the following holds for all sets of possible outputs $S \subset \mathcal{R}$ and all pairs of neighboring datasets $D, D' \in \mathcal{D}$:

$$\Pr(M(D) \in S) \leq \Pr(M(D') \in S) e^\epsilon + \delta$$

- A mechanism $M: \mathcal{D} \rightarrow \mathcal{R}$ that provides ϵ -DP except for certain "bad" outcomes $B \subset \mathcal{R}$, where $\Pr(M(D) \in B) \leq \delta$ (for any $D \in \mathcal{D}$) also provides (ϵ, δ) -DP.
- Proof is not as simple as it seems, but it can be proven

The Laplace Mechanism – Sensitivity

- We already saw an example of this. Now, we will make it more formal.
- First, we need to bound the maximum change in the non-private function we want to compute.
- Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$, and two neighboring datasets $D \in \mathcal{D}$ and $D' \in \mathcal{D}$, the ℓ_1 -**sensitivity** of f is the maximum change that replacing D for D' can cause in the output:

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

- Can generalize to other norms (such as ℓ_2 which we will see later)

The Laplace Mechanism

- Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$, and two neighboring datasets $D \in \mathcal{D}$ and $D' \in \mathcal{D}$, the ℓ_1 -sensitivity of f is the maximum change that replacing D for D' can cause in the output:

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

- Given any function f and its ℓ_1 sensitivity, we can turn it into a DP mechanism if we add Laplacian noise to its output:

Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$ with ℓ_1 -sensitivity Δ_1 , the **Laplace mechanism** is defined as $M(D) = f(D) + (Y_1, Y_2, \dots, Y_k)$ where each Y_i is independently distributed following $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$.

The Laplace Mechanism

- We already saw an example of this. Now, we will make it more formal.
- Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$, and two neighboring datasets $D \in \mathcal{D}$ and $D' \in \mathcal{D}$, the ℓ_1 -sensitivity of f is the maximum change that replacing D for D' can cause in the output:

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

- Given any function f and its ℓ_1 sensitivity, we can turn it into a DP mechanism if we add Laplacian noise to its output:

Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$ with ℓ_1 -sensitivity Δ_1 , the **Laplace mechanism** is defined as $M(D) = f(D) + (Y_1, Y_2, \dots, Y_k)$, where Y_i is independently distributed following $Y \sim \text{Lap}(b)$ with

The Laplace mechanism provides ϵ -DP

Recall, our example

- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .
- You can write $c = \sum_{i=2}^n x_i$.

Q: What is the sensitivity?

Recall, our example

- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .
- You can write $c = \sum_{i=2}^n x_i$.

Q: What is the sensitivity?

A: 1

Recall, our example

- Let x_i be the test result for user i ($x_i = 0$ for negative, $x_i = 1$ for positive)
- Let D be the dataset where $x_1 = x_A$ is Alice, and D' is the dataset where $x_1 = x_B$ is Bob. Assume that $x_A = 1$ and $x_B = 0$.
- $M(D) = \sum_{i=1}^n x_i + Y$, where Y is noise from a Laplace distribution with mean 0 and scale b .
- You can write $c = \sum_{i=2}^n x_i$.

Q: What is the sensitivity?

A: 1

Q: What is the maximum ratio between the distributions?

Remember this?


A: $\exp(1/b)$...

Let $b = 1/\epsilon$ and we have DP!

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!




Q: what does smaller ϵ mean?

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!




Q: what does smaller ϵ mean?

A: more privacy

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!



Q: if we want more privacy, would we need to add more or less noise?

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!


Q: if we want more privacy, would we need to add more or less noise?

A: more noise. That's why $b \propto \frac{1}{\epsilon}$.

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!



Q: if changing D for D' can cause a huge change in $f(\cdot)$, is that a large or small sensitivity?

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!


Q: if changing D for D' can cause a huge change in $f(\cdot)$, is that a large or small sensitivity?

A: large sensitivity

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!



Q: if changing D for D' can have a huge impact in f , do we need a lot or a little noise to hide this impact?

The Laplace Mechanism – checkpoint!

The Laplace Mechanism: $M(D) = f(D) + Y$ where $Y \sim \text{Lap}(b)$ with $b = \frac{\Delta_1}{\epsilon}$ provides ϵ -DP

The variance is $2b^2$; higher b means more noise!

Q: if changing D for D' can have a huge impact in f , do we need a lot or a little noise to hide this impact?

A: a lot of noise.
That's why $b \propto \Delta_1$


Laplace Mechanism: examples

Example 1: D contains the test results for virus X of a set of users. We want to release the total number of users that tested positive. How do we make this ϵ -DP?

- Under unbounded DP
- Under bounded DP

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$



Laplace Mechanism: examples

Example 1: D contains the test results for virus X of a set of users. We want to release the total number of users that tested positive. How do we make this ϵ -DP?


- Under unbounded DP
- Under bounded DP

A: sensitivity is 1 in both cases

Add $Y \sim \text{Lap}\left(\frac{1}{\epsilon}\right)$

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$




Laplace Mechanism: examples

Example 2: D contains the salaries of a set of users. The salaries range from 20k to 200k. We want to release the **total** salary of the users. How do we make this ϵ -DP?

- Under unbounded DP
- Under bounded DP

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$



Laplace Mechanism: examples

Example 2: D contains the salaries of a set of users. The salaries range from 20k to 200k. We want to release the **total** salary of the users. How do we make this ϵ -DP?

- Under unbounded DP
- Under bounded DP


A: sensitivity is bounded by 180k in the bounded and 200k in the unbounded

Add $Y \sim \text{Lap}\left(\frac{180k}{\epsilon}\right)$ or

$$Y \sim \text{Lap}\left(\frac{200k}{\epsilon}\right)$$

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$




Laplace Mechanism: examples

Example 3: D contains the salaries of n users (n is public knowledge). The salaries range from 20k to 200k. We want to release the **average** salary of users. How do we make this ϵ -DP?

- Under bounded DP

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$



Laplace Mechanism: examples

Example 3: D contains the salaries of n users (n is public knowledge). The salaries range from 20k to 200k. We want to release the **average** salary of users. How do we make this ϵ -DP?


- Under bounded DP

A: sensitivity is bounded by $180k/n$

Add $Y \sim \text{Lap}\left(\frac{180k}{n\epsilon}\right)$

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$




Laplace Mechanism: examples

Example 4: D contains the age of a set of users. We want to release the histogram of ages $[0-10)$, $[10-20)$... $[100,110)$. How do we make this ϵ -DP?

- Under unbounded DP
- Under bounded DP

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$



Laplace Mechanism: examples

Example 4: D contains the age of a set of users. We want to release the histogram of ages $[0-10)$, $[10-20)$... $[100,110)$. How do we make this ϵ -DP?


- Under unbounded DP
- Under bounded DP

A: sensitivity is 1 in unbounded 2 in bounded

Add $Y \sim \text{Lap}\left(\frac{1}{\epsilon}\right)$ or $Y \sim \text{Lap}\left(\frac{2}{\epsilon}\right)$ to each bucket in the histogram (drawn fresh for each bucket)

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$




Laplace Mechanism: examples

Example 5: Alice wishes to report her annual salary x_A in a differentially private way. The salaries at her company range from 20k to 200k (and this is public information). What mechanism can she follow so that she gets ϵ -DP?

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$




Laplace Mechanism: examples

Example 5: Alice wishes to report her annual salary x_A in a differentially private way. The salaries at her company range from 20k to 200k (and this is public information). What mechanism can she follow so that she gets ϵ -DP?

A: sensitivity is bounded by 180k
Add $Y \sim \text{Lap}\left(\frac{180k}{\epsilon}\right)$

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$




Laplace Mechanism: examples

Example 6: Alice wishes to report her age x_A in a differentially private way. It is public information that she is between 18 and 100 years old. She adds Laplacian noise with $b = 3$ to her age, and reports the resulting value. What is the level of DP that she gets?

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$



Laplace Mechanism: examples

Example 6: Alice wishes to report her age x_A in a differentially private way. It is public information that she is between 18 and 100 years old. She adds Laplacian noise with $b = 3$ to her age, and reports the resulting value. What is the level of DP that she gets?


A: sensitivity is bounded by 82

$$b = \frac{82}{\epsilon} = 3$$

$$\epsilon = 82/3$$

$$\Delta_1 \doteq \max_{D, D'} \|f(D) - f(D')\|_1$$

$$f(D) + Y \text{ is } \epsilon\text{-DP if } Y \sim \text{Lap}\left(\frac{\Delta_1}{\epsilon}\right)$$

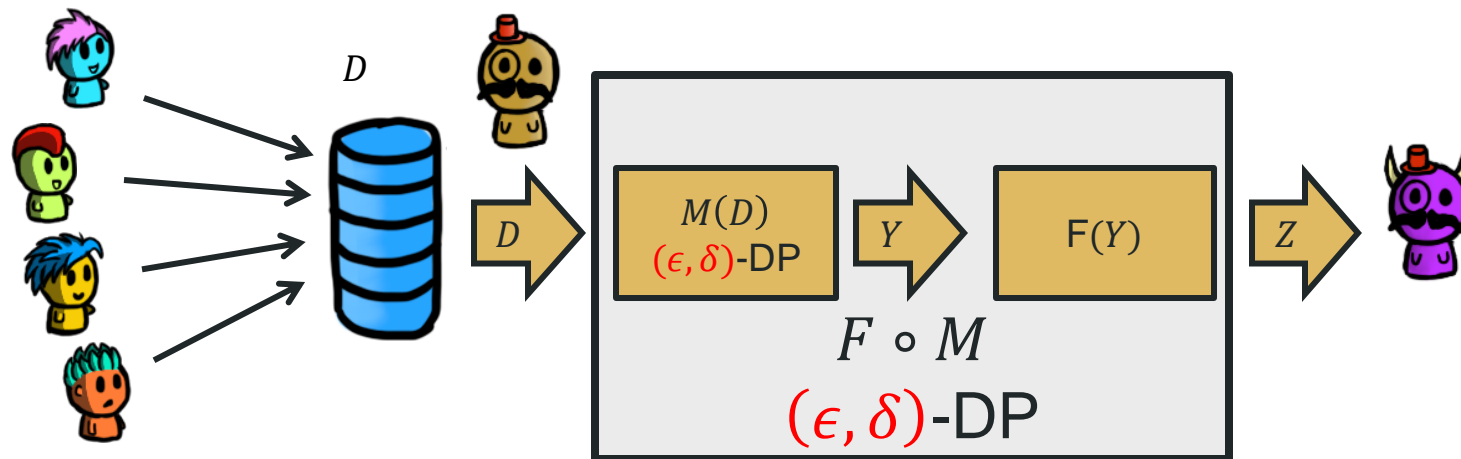


Properties of DP

Post-processing

Robustness to post-processing: Let $M: \mathcal{D} \rightarrow \mathcal{Y}$ be an (ϵ, δ) -DP mechanism, and let $F: \mathcal{Y} \rightarrow \mathcal{Z}$ be a (possibly randomized) mapping. Then, $F \circ M$ is (ϵ, δ) -DP.

- In layman terms, once you get a “privatized output” (Y) you cannot “unprivatize it” by running another mechanism.
- This makes a lot of sense: otherwise, the adversary could simply design an F that could “unprivatize” M !!



It is **very important** that F does not depend on D (other than through Y) at all! Otherwise, this will not hold!

Sequential Composition

Naïve composition: Let $M = (M_1, M_2, \dots, M_k)$ be a sequence of mechanisms, where M_i is (ϵ_i, δ_i) -DP. Then M is $(\sum_{i=1}^k \epsilon_i, \sum_{i=1}^k \delta_i)$ -DP

- This means that running k mechanisms on the same sensitive dataset, and publishing all k results, the ϵ s and δ s add up (privacy decrease as we publish more results).
- Recall, the attacks we saw in lecture 16...
 - More queries meant more leakage... this captures that.

Sequential Composition

- However, if we allow the overall δ to be slightly larger, we can get a much smaller ϵ :

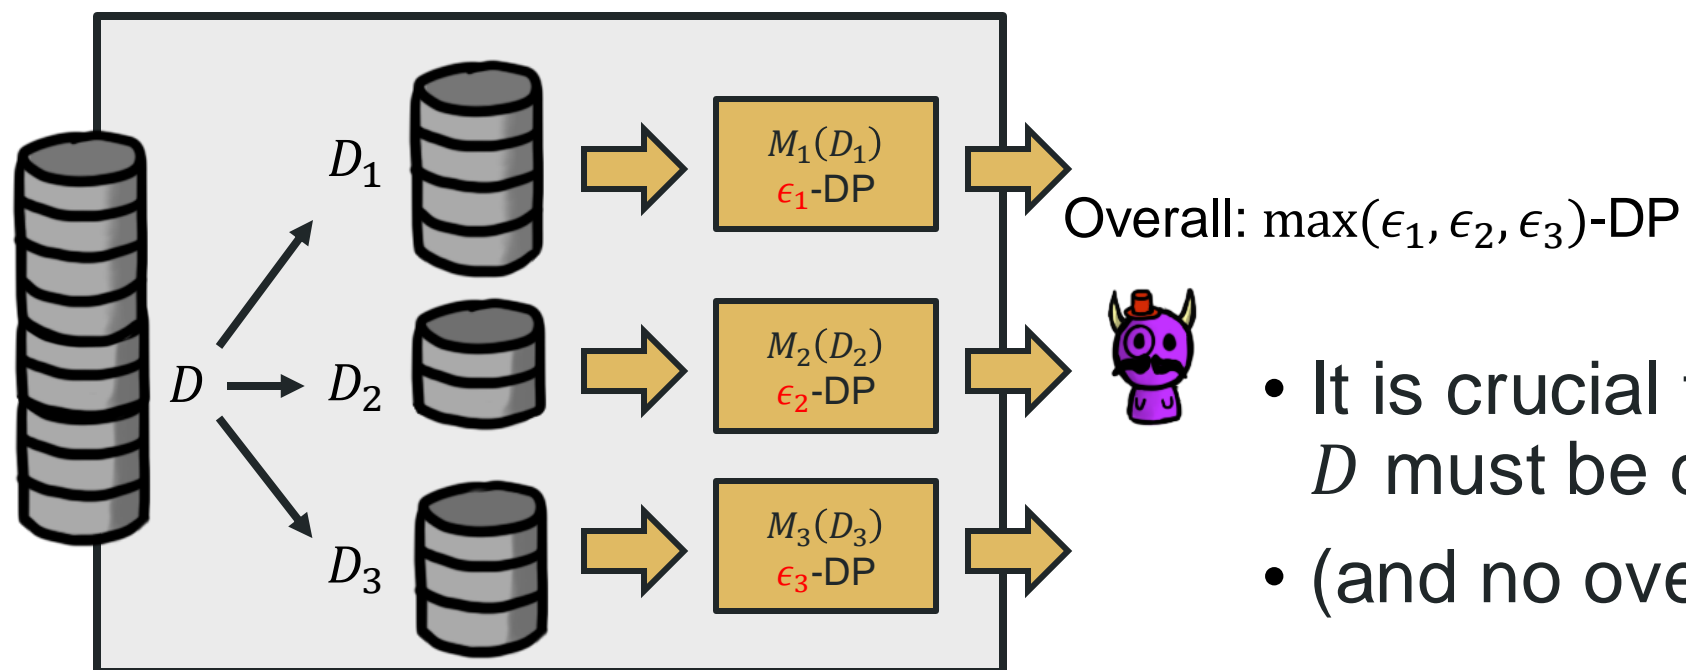
Advanced composition: Let $M = (M_1, M_2, \dots, M_k)$ be a sequence of mechanisms, where M_i is (ϵ, δ) -DP.

Then M is $\left(\epsilon \sqrt{2k \cdot \ln \left(\frac{1}{\delta'} \right)} + \frac{k\epsilon(e^\epsilon - 1)}{e^{\epsilon+1}}, k\delta + \delta' \right)$ -DP

- Note that the overall ϵ only grows on the order of \sqrt{k} now (loosely speaking), and that if we allow higher δ' then we can get a smaller overall ϵ .

Parallel Composition


Parallel Composition: Let $M = (M_1, M_2, \dots, M_k)$ be sequence of mechanisms, where M_i is ϵ_i -DP. Let D_1, D_2, \dots, D_k let a deterministic partition of D . Publishing $M_1(D_1), M_2(D_2), \dots, M_k(D_k)$ satisfies $(\max_{i \in [1, \dots, k]} \epsilon_i)$ -DP.




- It is crucial that the partition of D must be deterministic!
- (and no overlap)

More Mechanisms


Randomized Response (RR)




- Now we consider a mechanism with binary inputs and outputs, i.e., $M: \{0,1\} \rightarrow \{0,1\}$. This makes more sense in the local setting, where $x \in \{0,1\}$ and the outputs is $y \in \{0,1\}$.
- For example, x can be the answer to a yes/no question:
 - Have you voted for party X?
 - Have you tested positive for virus Y?
 - Have cheated in any assignment this term?
- Instead of reporting x , Alice follows the following process:



RR - Question



- Instead of reporting x , Alice follows the following process:



Q: compute these probabilities with an unbiased coin:


$$\Pr(y = 0|x = 0)$$

$$\Pr(y = 1|x = 0)$$


$$\Pr(y = 0|x = 1)$$

$$\Pr(y = 1|x = 1)$$

RR - Question



- Instead of reporting x , Alice follows the following process:



Q: compute these probabilities with an unbiased coin:

$$\Pr(y = 0|x = 0)$$

$$\Pr(y = 1|x = 0)$$

$$\Pr(y = 0|x = 1)$$

$$\Pr(y = 1|x = 1)$$

A:

$$\Pr(y = 0|x = 0) = 0.75$$

$$\Pr(y = 1|x = 0) = 0.25$$

$$\Pr(y = 0|x = 1) = 0.25$$

$$\Pr(y = 1|x = 1) = 0.75$$


Randomized Response (RR)

Differential Privacy (local model, discrete outputs)

A mechanism $M: \mathcal{X} \rightarrow \mathcal{Y}$ is ϵ -differentially private (ϵ -DP) if the following holds for all possible outputs $y \in \mathcal{Y}$ and all pairs of neighboring datasets $x, x' \in \mathcal{X}$:

$$\Pr(M(x) = y) \leq \Pr(M(x') = y) e^\epsilon$$

Q: what is the level of DP that RR provides?




Randomized Response (RR)

Differential Privacy (local model, discrete outputs)

A mechanism $M: \mathcal{X} \rightarrow \mathcal{Y}$ is ϵ -differentially private (ϵ -DP) if the following holds for all possible outputs $y \in \mathcal{Y}$ and all pairs of neighboring datasets $x, x' \in \mathcal{X}$:

$$\Pr(M(x) = y) \leq \Pr(M(x') = y) e^\epsilon$$

Q: what is the level of DP that RR provides?



A:


$$\frac{\Pr(y = 0|x = 0)}{\Pr(y = 0|x = 1)} = 3$$

$$\frac{\Pr(y = 0|x = 1)}{\Pr(y = 0|x = 0)} = \frac{1}{3}$$

The maximum ratio is 3. So $\epsilon = \log 3 \approx 1.10$.

Randomized Response (RR): Statistical Analyses


- More generally, we can have any probabilities p and $1 - p$.



Q: what is the ϵ in this case?

Randomized Response (RR): Statistical Analyses

- More generally, we can have any probabilities p and $1 - p$.



Q: what is the ϵ in this case?

Q: When $p \rightarrow 0.5$, $\epsilon \rightarrow 0$, does this make sense?


A:

$$\epsilon = \log\left(\max\left\{\frac{p}{1-p}, \frac{1-p}{p}\right\}\right)$$

Randomized Response (RR): Statistical Analyses

- Even though it is hard to guess the x given y (unless $p \rightarrow 1$ or 0), when multiple users report outputs we can get an estimate of the percentage of users that had $x = 1$.
- Assume there are n users reporting values, and a fraction p_0 have $x = 0$, while a fraction $p_1 = 1 - p_0$ have $x = 1$.

Q: How many answers $y = 1$ should we get, on average?




Randomized Response (RR): Statistical Analyses

- Even though it is hard to guess the x given y (unless $p \rightarrow 1$ or 0), when multiple users report outputs we can get an estimate of the percentage of users that had $x = 1$.
- Assume there are n users reporting values, and a fraction p_0 have $x = 0$, while a fraction $p_1 = 1 - p_0$ have $x = 1$.


Q: How many answers $y = 1$ should we get, on average?

A: $E\{y\} = p_0 \cdot (1 - p) + (1 - p_0) \cdot p$



Randomized Response (RR): Statistical Analyses

$$\mathbf{A: } E\{y\} = p_0 \cdot (1 - p) + (1 - p_0) \cdot p$$




- You can also see this using the law of total probability:
$$E\{y\} = \Pr(y = 1) = \Pr(y = 1|x = 0) \Pr(x = 0) + \Pr(y = 1|x = 1) \Pr(x = 1)$$
- Therefore, the analyst can estimate $E\{y\}$ empirically using the reported values (let this be \bar{y}), and then compute p_0 by solving $\bar{y} = p_0 \cdot (1 - p) + (1 - p_0) \cdot p$.
- This gives us an estimator for p_0 :

$$\hat{p}_0 = \frac{\bar{y} - p}{1 - 2p}$$

Q: Can this gives us a negative estimate? Why?

Randomized Response (RR): Statistical Analyses

$$\mathbf{A: } E\{y\} = p_0 \cdot (1 - p) + (1 - p_0) \cdot p$$



- You can also see this using the law of total probability:
$$E\{y\} = \Pr(y = 1) = \Pr(y = 1|x = 0) \Pr(x = 0) + \Pr(y = 1|x = 1) \Pr(x = 1)$$
- Therefore, the analyst can estimate $E\{y\}$ empirically using the reported values (let this be \bar{y}), and then compute p_0 by solving $\bar{y} = p_0 \cdot (1 - p) + (1 - p_0) \cdot p$.
- This gives us an estimator for p_0 :

$$\hat{p}_0 = \frac{\bar{y} - p}{1 - 2p}$$

Q: Can this gives us a negative estimate? Why?


A: It can happen, this will only approach the true percentage as $n \rightarrow \infty$.

Statistical analysis with RR: exercise

- **Disclaimer:** you have $\epsilon = 1.1$ (high-ish privacy); no matter what you report in this exercise, you can always claim it was not your true answer (**plausible deniability**).
- Let's learn how many of you cheated in an exam/assignment before/after covid times.

Statistical analysis with RR: exercise

- $x = 1$ means “I have cheated”. Flip two coins, run randomized response:






	During covid	After covid
Number of participants		
Number of $y = 1$		
Empirical avg: \bar{y}		
Estimate of non-cheaters: $\hat{p}_0 = 1.5 - 2\bar{y}$		
Estimate of cheaters: $\hat{p}_1 = 2\bar{y} - 0.5$		

Exponential Mechanism




- Sometimes, adding Laplacian noise could destroy the utility of a mechanism.
 - What if we want noise that is not symmetrical?
- Sometimes, we do not want to make numerical answers private, but we want to be able to report objects/classes/categories.
 - How do we do this privately?
- The exponential mechanism can be used to provide DP in many settings.
- The idea is that we will report an output privately, but with a probability *proportional to its utility*.

Private Auction: noise is not great for DP!



- A set of users wants to buy an item, and each has a private amount they are willing to pay: v_i .
- The retailer sees the v_i 's and could choose the largest price p that maximizes the revenue (number of clients with $v_i \geq p$, times p).
- However, the p chosen this way would reveal information about the users' valuations v_i , which can be privacy-sensitive.

Private Auction: noise is not great for DP!



Issue here: the revenue (utility) is very sensitive to the choice of p :

- If $p = 1$, then the revenue is \$3
- If $p = 1.01$, then the revenue drops to \$1.01
- If $p = 3.01$, then the revenue is \$3.01
- But at $p = 3.02$, the revenue drops to \$0

Adding noise to p before making it public can destroy the utility (revenue)

The Exponential Mechanism

Given a database $D \in \mathcal{D}$, a set of outputs \mathcal{H} and a score function $s: \mathcal{D} \times \mathcal{H} \rightarrow \mathbb{R}$, the **exponential mechanism** M_E chooses an output $h \in \mathcal{H}$ with probability proportional to:

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

Here, Δ is the sensitivity of the score function, defined as

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$




The Exponential Mechanism

Given a database $D \in \mathcal{D}$, a set of outputs \mathcal{H} and a score function $s: \mathcal{D} \times \mathcal{H} \rightarrow \mathbb{R}$, the **exponential mechanism** M_E chooses an output $h \in \mathcal{H}$ with probability proportional to:

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

- In order to compute the actual probability $\Pr(M_E(D) = h)$, we need to compute the values of the score function for every $h \in \mathcal{H}$. This can sometimes be very expensive.
- The exponential mechanism chooses items proportional to the score function
- The epsilon smooths this distribution
- The set of outputs is public knowledge, the choice is sensitive

The Exponential Mechanism – an example






- **Q:** how can we use the exponential mechanism in this scenario?

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

The Exponential Mechanism – an example






- **Q:** how can we use the exponential mechanism in this scenario?

A: we can discretize the set of possible outputs, e.g., $\mathcal{H} = \{0.1, 0.2, \dots, 10\}$ (assuming the maximum price of the item is \$10). This is the set of possible values p . Compute the probability of each and sample with that probability.

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

The Exponential Mechanism – an example






- Then, the retailer computes $s(D, h)$ for each possible output h . Note that D is simply $\{v_1, v_2, v_3\}$ in this case.

Q: what will be the sensitivity?

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

The Exponential Mechanism – an example



$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

- Then, the retailer computes $s(D, h)$ for each possible output h . Note that D is simply $\{v_1, v_2, v_3\}$ in this case.




Q: what will be the sensitivity?

A: the maximum effect that an item can have in the revenue is \$10, assuming the maximum price of the item is \$10).

The Exponential Mechanism – an example






- **Q:** Assume $\mathcal{H} = \{1, 2, 3, 4\}$ compute the probability of selecting each output, when $\epsilon = 1$.



$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

The Exponential Mechanism – an example



$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

- **Q:** Assume $\mathcal{H} = \{1, 2, 3, 4\}$ compute the probability of selecting each output, when $\epsilon = 1$.


A: sensitivity would be 4

- Scores would be $\{3, 2, 3, 0\}$
- $\Pr(M_E(D) = 1) = \exp\left(\frac{3}{8}\right) / \sum_h \exp\left(\frac{s(D, h)}{8}\right)$
- $\Pr(M_E(D) = 2) = \exp\left(\frac{2}{8}\right) / \sum_h \exp\left(\frac{s(D, h)}{8}\right)$
- $\Pr(M_E(D) = 3) = \exp\left(\frac{3}{8}\right) / \sum_h \exp\left(\frac{s(D, h)}{8}\right)$
- $\Pr(M_E(D) = 4) = 1 / \sum_h \exp\left(\frac{s(D, h)}{8}\right)$
- $\sum_h \exp\left(\frac{s(D, h)}{8}\right) = 2\exp\left(\frac{3}{8}\right) + \exp\left(\frac{2}{8}\right) + 1$

The Exponential Mechanism – an example

- Assume we want to make a small decision tree for classifying heart attacks based on cholesterol
- Given the following dataset we want to choose a threshold h that maximizes accuracy of the classifier $f(c)$:


Cholesterol (c)	Heart Attack (y)
216	0
501	1
100	0
535	1
214	1



- Let $s(D, h) = \frac{1}{n} \sum_i (f_h(c_i) == y_i)$

The Exponential Mechanism – an example

Cholesterol (c)	Heart Attack (y)
216	0
501	1
100	0
535	1
214	1



$$s(D, h) = \frac{1}{n} \sum_i (f_h(c_i) == y_i)$$

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

$$\Delta = \max_h \max_{D, D'} |s(D, h) - s(D', h)|$$

- **Q:** Assume $\mathcal{H} = \{100, 200, 300, 400, 500\}$ compute the probability of selecting each output, when $\epsilon = 1.25$.

Just checking...

Given a database $D \in \mathcal{D}$, a set of outputs \mathcal{H} and a score function $s: \mathcal{D} \times \mathcal{H} \rightarrow \mathbb{R}$, the **exponential mechanism** M_E chooses an output $h \in \mathcal{H}$ with probability proportional to:

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

Q: What is the runtime complexity of the exponential mechanism in relation to \mathcal{H}

Just checking...

Given a database $D \in \mathcal{D}$, a set of outputs \mathcal{H} and a score function $s: \mathcal{D} \times \mathcal{H} \rightarrow \mathbb{R}$, the **exponential mechanism** M_E chooses an output $h \in \mathcal{H}$ with probability proportional to:

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

Q: What is the runtime complexity of the exponential mechanism in relation to \mathcal{H}

A: $O(|\mathcal{H}|)$

Just checking...

Given a database $D \in \mathcal{D}$, a set of outputs \mathcal{H} and a score function $s: \mathcal{D} \times \mathcal{H} \rightarrow \mathbb{R}$, the **exponential mechanism** M_E chooses an output $h \in \mathcal{H}$ with probability proportional to:

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

Q: What is the effect of reducing epsilon on the probability of each item?

Just checking...

Given a database $D \in \mathcal{D}$, a set of outputs \mathcal{H} and a score function $s: \mathcal{D} \times \mathcal{H} \rightarrow \mathbb{R}$, the **exponential mechanism** M_E chooses an output $h \in \mathcal{H}$ with probability proportional to:

$$\Pr(M_E(D) = h) \propto \exp\left(\frac{\epsilon \cdot s(D, h)}{2\Delta}\right)$$

Q: What is the effect of reducing epsilon on the probability of each item?

A: The probabilities become more similar. As epsilon tends to 0, probabilities tend to $\frac{1}{|\mathcal{H}|}$

The Exponential Mechanism is Generic!

Q: What is the probability of selection when the score function is $s(D, h) = -|f(D) - h|$

The Exponential Mechanism is Generic!

Q: What is the probability of selection when the score function is $s(D, h) = -|f(D) - h|$

Q: What distribution is this?

A: $\propto \exp\left(-\frac{\epsilon|f(D) - h|}{2\Delta}\right)$

The Exponential Mechanism is Generic!

Q: What is the probability of selection when the score function is $s(D, h) = -|f(D) - h|$

A: $\propto \exp\left(-\frac{\epsilon|f(D) - h|}{2\Delta}\right)$

Q: What distribution is this?

A: Even the Laplace mechanism is an instantiation of the exponential mechanism!

Bonus Content

The Gaussian Mechanism

- So far, we have seen mechanisms for pure DP. Let's see one for approximate DP.
- First, given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$, we define the ℓ_2 -sensitivity as:

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$

The Gaussian Mechanism

- Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$, we define the ℓ_2 -sensitivity as:

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$

- The Gaussian mechanism simply adds Gaussian noise to the output of the function:

Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$ with ℓ_2 -sensitivity Δ_2 , the **Gaussian mechanism** is defined as $M(D) = f(D) + (Y_1, Y_2, \dots, Y_k)$ where each Y_i is independently distributed as $Y_i \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$.

The Gaussian Mechanism

- Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$, we define the ℓ_2 -sensitivity as:

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$

- The Gaussian mechanism simply adds Gaussian noise to the output of the function:

Given a function $f: \mathcal{D} \rightarrow \mathbb{R}^k$ with ℓ_2 -sensitivity Δ_2 , the **Gaussian mechanism** is defined as $M(D) = f(D) + (Y_1, Y_2, \dots, Y_k)$ where each Y_i is independent and distributed as $Y_i \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$

The Gaussian mechanism provides ϵ, δ -DP

Let's think about this

The Gaussian mechanism $M(D) = f(D) + Y$ where $Y \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$ provides (ϵ, δ) -DP.

Q: does the relationship between the privacy parameter ϵ and the noise variance σ^2 make sense?

Let's think about this

The Gaussian mechanism $M(D) = f(D) + Y$ where $Y \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$ provides (ϵ, δ) -DP.

Q: does the relationship between the privacy parameter ϵ and the noise variance σ^2 make sense?

A: yes, to provide more privacy (lower ϵ) we need more noise (higher σ^2).

Let's think about this

The Gaussian mechanism $M(D) = f(D) + Y$ where $Y \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$ provides (ϵ, δ) -DP.

Q: if we fix the noise level (σ), what is the relationship between ϵ and δ , and why?

Let's think about this

The Gaussian mechanism $M(D) = f(D) + Y$ where $Y \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$ provides (ϵ, δ) -DP.

Q: if we fix the noise level (σ), what is the relationship between ϵ and δ , and why?

A: for a fixed noise, ϵ and δ will be inversely proportional: if we want allow for a higher δ then that level of noise can provide lower ϵ 's.


Let's think about this

The Gaussian mechanism $M(D) = f(D) + Y$ where $Y \sim N(0, \sigma^2)$ with $\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$ provides (ϵ, δ) -DP.

Q: if we fix the noise level (σ), what is the relationship between ϵ and δ , and why?

A: for a fixed noise, ϵ and δ will be inversely proportional: if we want allow for a higher δ then that level of noise can provide lower ϵ 's.

This is not just for the Gaussian mechanism, but all ϵ, δ -DP mechanisms:




Gaussian Mechanism: examples

Example 1: D contains the salaries of a set of n users. The salaries range from 10k to 200k. We want to release the **total** salary of the users. What is the σ^2 of the gaussian mechanism under bounded DP assuming $\delta = 1/n^2$

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$

$f(D) + Y$ is (ϵ, δ) -DP if
 $Y \sim N(0, \sigma^2)$

$$\sigma^2 = 2 \ln \left(\frac{1.25}{\delta} \right) \Delta_2^2 / \epsilon^2$$



Gaussian Mechanism: examples

Example 1: D contains the salaries of a set of n users. The salaries range from 10k to 200k. We want to release the **total** salary of the users. What is the σ^2 of the gaussian mechanism under bounded DP assuming $\delta = 1/n^2$


A: sensitivity is 190k

$$\sigma^2 = 2 \ln(1.25 n^2) (190k)^2 / \epsilon^2$$

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$

$f(D) + Y$ is (ϵ, δ) -DP if
 $Y \sim N(0, \sigma^2)$

$$\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$$




Gaussian Mechanism: examples

Example 2: D contains the age of a set of users. We want to release the histogram of ages $[0-10)$, $[10-20)$... $[100,110)$. What is the σ^2 of the gaussian mechanism under bounded DP assuming $\delta = 1/n^2$

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$

$f(D) + Y$ is (ϵ, δ) -DP if
 $Y \sim N(0, \sigma^2)$

$$\sigma^2 = 2 \ln \left(\frac{1.25}{\delta} \right) \Delta_2^2 / \epsilon^2$$



Gaussian Mechanism: examples

Example 2: D contains the age of a set of users. We want to release the histogram of ages $[0-10)$, $[10-20)$... $[100,110)$. What is the σ^2 of the gaussian mechanism under bounded DP assuming $\delta = 1/n^2$


A: sensitivity $\sqrt{2}$ in bounded DP

$$\sigma^2 = 4 \ln(1.25 n^2) / \epsilon^2$$

$$\Delta_2 \doteq \max_{D, D'} \|f(D) - f(D')\|_2$$



$f(D) + Y$ is (ϵ, δ) -DP if
 $Y \sim N(0, \sigma^2)$

$$\sigma^2 = 2 \ln\left(\frac{1.25}{\delta}\right) \Delta_2^2 / \epsilon^2$$



General Discrete Mechanisms

- A general mechanism that takes inputs and outputs from discrete sets can be written in matrix form by listing its inputs as rows, and its outputs as columns
 - this is similar to how we wrote mechanism when we talked about statistical inference attacks



	y_1	y_2	...	y_m
x_1
x_2	...	$\Pr(y_2 x_2)$
...
x_n

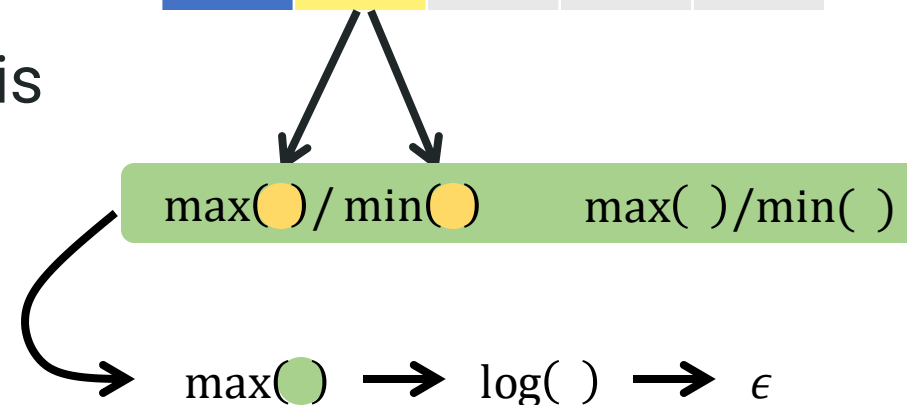
you get the idea...

General Discrete Mechanisms

- Computing ϵ for a mechanism in matrix form is very easy!
1. For every column (output), take the largest value and divide it by the smallest
 - This is computing $\max_{x,x'} \Pr(y|x) / \Pr(y|x')$ for a given y .
 2. Take the largest one of those ratios
 - This value is \leq than any $\Pr(y|x) / \Pr(y|x')$
 3. Compute the natural logarithm of this, and this will give you ϵ .
 - Since ϵ is the value such that

$$\frac{\Pr(y|x)}{\Pr(y|x')} \leq e^\epsilon$$

	y_1	y_2	...	y_m
x_1
x_2
...
x_n



General Discrete Mechanism: example

Q: Alice uses the generalized randomized response to report a differentially private version of her location to a location-based service provider. Her possible locations are points of interest $\{x_1, x_2, \dots, x_n\}$. The mechanism reports her real location with probability p and any other location with probability q .

- What is the ϵ -DP level this provides? (note that it will be dependent on p and n).
- You can assume $p > 1/n$.
- You should check that, when setting $n = 2$, you get the same formula for ϵ as for the RR mechanism.

General Discrete Mechanism: example

Q: Alice uses the generalized randomized response to report a differentially private version of her location to a location-based service provider. Her possible locations are points of interest $\{x_1, x_2, \dots, x_n\}$. The mechanism reports her real location with probability p and any other location with probability q .

- What is the ϵ -DP level this provides? (note that it will be dependent on p and n).
- You can assume $p > 1/n$.
- You should check that, when setting $n = 2$, you get the same formula for ϵ as for the RR mechanism.

A: $q = \frac{1-p}{n-1}$. Since $p > \frac{1}{n}$, then $p > q$, and the maximum ratio for any output will be

$$\frac{p}{q} = \frac{p(n-1)}{1-p} \rightarrow \epsilon = \log\left(\frac{p(n-1)}{1-p}\right)$$

When $n = 2$, we are back to randomized response!