

Lecture 4: Differential Privacy

CS 181E

Fall 2026


Review: Anonymization is hard

- Personal data showing up in unexpected places
- Errors suppressing data
- **Thinking only about personal data / PII is not sufficient**
- It's hard to reason about what is/is not identifiable
- Availability of **auxiliary information**
 - Extremely hard to control


Math to the rescue!

Review: Probability

- Probability is the math of describing how likely something is
- A **probability** is a number in the range $[0,1]$ that describes how likely something is.




$$\Pr[\textit{coin} = \textit{Heads}] = \frac{1}{2}$$
$$\Pr[\textit{coin} = \textit{Tails}] = \frac{1}{2}$$




$$\Pr[\textit{weather} = \textit{Raining}] = \frac{1}{100}$$
$$\Pr[\textit{weather} = \textit{Not Raining}] = \frac{99}{100}$$

Conditional Probability

- The **conditional probability** is the likelihood of an event happening given that another event has already occurred.





$$\Pr[\text{weather} = \text{Raining}] = \frac{1}{100}$$
$$\Pr[\text{weather} = \text{Raining} \mid \text{Cloudy}] = \frac{1}{10}$$



Independence

- Two random events are **independent** if the probability of event B is the same no matter what the outcome of event A




Combining Probabilities

- In general, the probability of both of two events A and B occurring is $\Pr(A \cap B) = \Pr(A) \cdot \Pr(B|A)$
- If two events A and B are independent, the probability of both events occurring is $\Pr(A \cap B) = \Pr(A) \cdot \Pr(B)$
- The probability of one of two events A and B occurring is $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$
- The probability of one of two mutually exclusive events A and B occurring is $\Pr(A \cup B) = \Pr(A) + \Pr(B)$

Exercise: Probabilities


Weather



$$\Pr[\text{weather} = \text{Rain}] = \frac{1}{100}$$

$$\Pr[\text{weather} = \text{Rain} \mid \text{Cloudy}] = \frac{1}{10}$$

Weather



- What is the probability that it is cloudy?


Probability Distributions

- a **sample space** is the set of events comprising all possible outcomes
 - possible outcomes must be mutually exclusive
 - the set of possible outcomes must be collectively exhaustive

Example: $S = \{Raining, Not Raining\}$

- a **probability distribution** is a function that assigns probabilities to events in a sample space such that:
 - the probability of each event is non-negative
 - the probability that one of the events occurs is 1

Example:




$$\begin{aligned}\Pr[\text{weather} = \text{Raining}] &= .01 \\ \Pr[\text{weather} = \text{Not Raining}] &= .99\end{aligned}$$

- **drawing** or **sampling** from a distribution returns a value in the sample space with the corresponding probabilities

Defining Privacy: Try #1

- You don't know anything more after querying the database than you already knew before




The Problem...

- Assume you have external information that I am half an inch taller than the median height of a woman in the United States (plus anything you can infer from seeing/interacting with me).
- How tall do you think I am?
- Assume a database has data on the heights of 8,545 people in the United States (but that I'm not in the dataset)
- Intuitively, this should be ok, since I'm not in the dataset
- but...

Simplifying the problem...


Goal

- Privacy no matter how the user interacts with the dataset




The Simpler Case

- Privacy when the user only gets to run one query over the dataset




Differential Privacy



A query Q is differentially private if the answer you get in the world with my data is included is pretty much the same as the answer you would get otherwise



Differential Privacy




Range is all possible outputs

A query Q is ϵ -differentially private if $\forall x \in \text{Range}(Q), \forall D, r \in D,$
 $\Pr[Q(D) = x] \leq e^\epsilon \cdot \Pr[Q(D - r) = x]$


Exercise: Calculating epsilon

- How differentially-private is the constant query $f(x) = 47$?

World 1



World 2



$$\forall x \in \text{Range}(f), D, r, \Pr[f(D) = x] \leq e^{???} \Pr[f(D - r) = x]$$

$$\forall x \in \text{Range}(f), D, r, \frac{\Pr[f(D) = x]}{\Pr[f(D - r) = x]} \leq e^{???}$$

$$\forall D, r, \frac{\Pr[f(D) = 47]}{\Pr[f(D - r) = 47]} = \frac{1}{1} = 1 = e^0 \Rightarrow \text{0-differentially private}$$

How do we achieve differential privacy?

- **Goal:** Bound how likely the two worlds are to be different
- **Intermediate step:** Bound how much the two worlds can be different

Sensitivity

- The sensitivity Δ of a query Q is the maximum the answer to Q can possibly change between two databases that differ only by one person
- Q = number of people taller than 6 ft $\Delta = 1$
- Q = maximum height of a person $\Delta = 48$

Exercise: Sensitivity

- Assume you have a database containing the heights of 100 users specified in inches. You should assume that all heights are between 48 in and 96in.
- What is the sensitivity of the following queries?
 1. The number of people who are 5' 4"
 2. The median height in the dataset
 3. The mean height in the dataset


How do we achieve differential privacy?

- Don't answer the query exactly, give an approximate answer
- Approximate answer is generated by adding random noise
- Distribution of random noise is chosen very carefully
- How much noise depends on the sensitivity of the query

Laplacian Distribution

- $Lap(b)$ is the probability distribution with the property that

$$\Pr[Lap(b) = x] = \frac{1}{2b} \cdot e^{-\frac{|x|}{b}}$$



Laplacian Mechanism

- Given a query Q on a database D that has sensitivity Δ , respond with $f(D) = Q(D) + Y$ where Y is drawn from the distribution $Lap(\frac{\Delta}{\epsilon})$

- Theorem:** this mechanism satisfies ϵ -differential privacy

$$\begin{aligned}
 \frac{\Pr[f(D) = x]}{\Pr[f(D-r) = x]} &= \frac{\Pr[Q(D) + Y = x]}{\Pr[Q(D-r) + Y = x]} = \frac{\Pr[Y = x - Q(D)]}{\Pr[Y = x - Q(D-r)]} \\
 &= \frac{\frac{1}{2b} \cdot e^{-\frac{|x-Q(D)|}{b}}}{\frac{1}{2b} \cdot e^{-\frac{|x-Q(D-r)|}{b}}} = \frac{\frac{1}{2(\Delta/\epsilon)} \cdot e^{-\frac{|x-Q(D)|}{\Delta/\epsilon}}}{\frac{1}{2(\Delta/\epsilon)} \cdot e^{-\frac{|x-Q(D-r)|}{\Delta/\epsilon}}} = \frac{e^{-\frac{|x-Q(D)|}{\Delta/\epsilon}}}{e^{-\frac{|x-Q(D-r)|}{\Delta/\epsilon}}} \\
 &= e^{\frac{|x-Q(D-r)|}{\Delta/\epsilon} - \frac{|x-Q(D)|}{\Delta/\epsilon}} = e^{\left(\frac{\epsilon}{\Delta}\right) \cdot (|x-Q(D-r)| - |x-Q(D)|)} \\
 &\leq e^{\left(\frac{\epsilon}{\Delta}\right) \cdot (|x-Q(D-r) - x + Q(D)|)} = e^{\left(\frac{\epsilon}{\Delta}\right) \cdot (|Q(D) - Q(D-r)|)} \\
 &\leq e^{\left(\frac{\epsilon}{\Delta}\right) \cdot \Delta} = e^{\epsilon} \Rightarrow \epsilon\text{-differentially private}
 \end{aligned}$$

What if people ask more than one query?

Composition

- **Theorem:** If f_1 satisfies ϵ_1 -differential privacy and f_2 satisfies ϵ_2 -differential privacy, then $f_1; f_2$ satisfies $(\epsilon_1 + \epsilon_2)$ -differential privacy.

$$\begin{aligned} \frac{\Pr[(f_1(D) = x_1) \text{ and } (f_2(D) = x_2)]}{\Pr[(f_1(D - r) = x_1) \text{ and } (f_2(D - r) = x_2)]} &= \frac{\Pr[f_1(D) = x_1] \cdot \Pr[f_2(D) = x_2]}{\Pr[f_1(D - r) = x_1] \cdot \Pr[f_2(D - r) = x_2]} \\ &= \left(\frac{\Pr[f_1(D) = x_1]}{\Pr[f_1(D - r) = x_1]} \right) \cdot \left(\frac{\Pr[f_2(D) = x_2]}{\Pr[f_2(D - r) = x_2]} \right) \\ &\leq e^{\epsilon_1} \cdot e^{\epsilon_2} \\ &\leq e^{\epsilon_1 + \epsilon_2} \end{aligned}$$

- This is worst case. Usually do better in practice.
- Can also do better with coordinated noise

Other Limitations

- Have to trust whoever is storing the database
- Whoever stores database has to be online for analysis
- Alternative: **Local differential privacy**