Lecture 1: Introduction to Differential Privacy and the Laplace Mechanism

1 Introduction

Differential privacy is a rigorous, mathematical framework that enables the protection of individual privacy in datasets while still allowing for useful statistical analysis. It ensures that any output from a differentially private algorithm is nearly the same, whether or not an individual's data is included in the dataset. This provides a formal measure of privacy protection and makes it difficult for an adversary to infer information about an individual.

2 Definition of Differential Privacy

2.1 Adjacent Datasets

Two datasets D_1 and D_2 are said to be adjacent if they differ in the data of exactly one individual. Formally, they are adjacent if:

$$|D_1 \Delta D_2| = 1$$

2.2 ϵ-Differential Privacy

A randomized algorithm $\mathcal{A}$ satisfies ϵ-differential privacy if for any two adjacent datasets D_1 and D_2, and for any possible output $O \subseteq Range(\mathcal{A})$, the following inequality holds:

$$\frac{Pr[\mathcal{A}(D_1) \in O]}{Pr[\mathcal{A}(D_2) \in O]} \leq e^{\epsilon}$$

Here, ϵ is a privacy parameter. Smaller values of ϵ imply stronger privacy guarantees.

3 The Laplace Mechanism

3.1 Definition

The Laplace mechanism is a fundamental technique for achieving differential privacy. Given a function $f : \mathcal{D} \to \mathbb{R}^d$, where $\mathcal{D}$ is the domain of the dataset and d is the dimension of the output, the Laplace mechanism adds Laplace noise to the output of f.

Let b be the scale parameter of the Laplace distribution, which is given by:

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

Given a dataset D, the Laplace mechanism $\mathcal{A}$ is defined as:

$$\mathcal{A}(D) = f(D) + Lap(0|b)^d$$

3.2 Sensitivity

To ensure ϵ-differential privacy, we need to determine the appropriate scale parameter b. This is where the sensitivity of the function f comes into play. The sensitivity Δf of a function f is the maximum difference in the output of f when applied to any two adjacent datasets:

$$\Delta f = \max_{D_1, D_2 : |D_1 \Delta D_2| = 1} \|f(D_1) - f(D_2)\|_1$$

3.3 Achieving ϵ-Differential Privacy

To achieve ϵ-differential privacy, we choose the scale parameter b as:

$$b = \frac{\Delta f}{\epsilon}$$

With this choice of b, the Laplace mechanism $\mathcal{A}$ satisfies ϵ-differential privacy.

4 Proof: Laplace Mechanism is Differentially Private

The Laplace mechanism $\mathcal{A}(D) = f(D) + Lap(0|b)^d$ satisfies ϵ-differential privacy, where $b = \frac{\Delta f}{\epsilon}$ and Δf is the sensitivity of the function f.

Proof. Let D_1 and D_2 be any two adjacent datasets, and let $O \subseteq Range(\mathcal{A})$. We need to show that:

$$\frac{Pr[\mathcal{A}(D_1) \in O]}{Pr[\mathcal{A}(D_2) \in O]} \leq e^{\epsilon}$$

Let $y_1 = f(D_1)$ and $y_2 = f(D_2)$. Then, we have:

$$\frac{Pr[\mathcal{A}(D_1) \in O]}{Pr[\mathcal{A}(D_2) \in O]} = \frac{Pr[y_1 + Lap(0|b)^d \in O]}{Pr[y_2 + Lap(0|b)^d \in O]}$$

By defining $O' = \{x - y_1 : x \in O\}$, we can rewrite the probability ratio as:

$$\frac{Pr[Lap(0|b)^d \in O']}{Pr[Lap(0|b)^d \in O' + (y_1 - y_2)]}$$

Let $Lap_b(x) = Lap(x|b)$. Then, for any $x \in O'$, we have:

$$\begin{aligned}
\frac{Lap_b(x)}{Lap_b(x - (y_1 - y_2))} &= \frac{\frac{1}{2b} e^{-\frac{|x|}{b}}}{\frac{1}{2b} e^{-\frac{|x - (y_1 - y_2)|}{b}}} \\
&= e^{\frac{|x - (y_1 - y_2)| - |x|}{b}} \\
&\leq e^{\frac{|y_1 - y_2|}{b}} \\
&\leq e^{\frac{\Delta f}{b}} \\
&= e^{\epsilon}
\end{aligned}$$

The second-to-last inequality follows from the triangle inequality and the definition of sensitivity. Thus, we have:

$$\frac{Lap_b(x)}{Lap_b(x - (y_1 - y_2))} \leq e^{\epsilon}$$

Integrating both sides of this inequality over $x \in O'$, we obtain:

$$\frac{Pr[Lap(0|b)^d \in O']}{Pr[Lap(0|b)^d \in O' + (y_1 - y_2)]} \leq e^{\epsilon}$$

This completes the proof that the Laplace mechanism satisfies ϵ-differential privacy. □

5 Example: Querying the Mean of a Dataset

Suppose we have a dataset D of n real numbers, where each number is in the range $[0, M]$. We want to compute the mean of the dataset while preserving differential privacy. To do this, we can use the Laplace mechanism.

5.1 Function and Sensitivity

First, we define the function f that computes the mean of a dataset:

$$f(D) = \frac{1}{n} \sum_{i=1}^{n} D_i$$

Next, we compute the sensitivity of f. Since the range of each data point is $[0, M]$, the maximum difference in the mean when we add or remove a data point is $\frac{M}{n}$. Therefore, the sensitivity Δf is:

$$\Delta f = \frac{M}{n}$$

5.2 Applying the Laplace Mechanism

To achieve ϵ-differential privacy, we set the scale parameter b as:

$$b = \frac{\Delta f}{\epsilon} = \frac{M}{n\epsilon}$$

Now, we can apply the Laplace mechanism to compute the differentially private mean:

$$\mathcal{A}(D) = f(D) + Lap(0|b) = \frac{1}{n} \sum_{i=1}^{n} D_i + Lap\left(0|\frac{M}{n\epsilon}\right)$$

By adding Laplace noise with the appropriate scale parameter, we can compute the mean of the dataset while preserving ϵ-differential privacy.

6 Conclusion

In this lecture, we introduced the concept of differential privacy, defined ϵ-differential privacy, and analyzed the Laplace mechanism. We also demonstrated how to apply the Laplace mechanism to query the mean of a dataset while preserving privacy. In the next lectures, we will explore other mechanisms and techniques for achieving differential privacy in various settings.