CSCE 658: Randomized Algorithms

Lecture 20

Samson Zhou

Relevant Supplementary Material

- Chapter 1-3 of “The Algorithmic Foundations of Differential Privacy”, by Cynthia Dwork and Aaron Roth (https://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf)

Last Time: Differential Privacy

- [DMNS06] Given $\varepsilon > 0$ and $\delta \in (0,1)$, a randomized algorithm $A: U^* \to Y$ is (ε, δ)-differentially private if, for every neighboring pair D and D' of datasets, and for all $E \subseteq Y$,
$$\Pr[A(D) \in E] \leq e^{\varepsilon} \cdot \Pr[A(D') \in E] + \delta$$

$A($  $) \approx A($  $)$

Last Time: Differential Privacy

- [DMNS06] Given $\varepsilon > 0$ and $\delta \in (0,1)$, a randomized algorithm $A: U^* \to Y$ is (ε, δ)-differentially private if, for every neighboring pair D and D' of datasets, and for all $E \subseteq Y$,

$$\Pr[A(D) \in E] \leq e^{\varepsilon} \cdot \Pr[A(D') \in E] + \delta$$

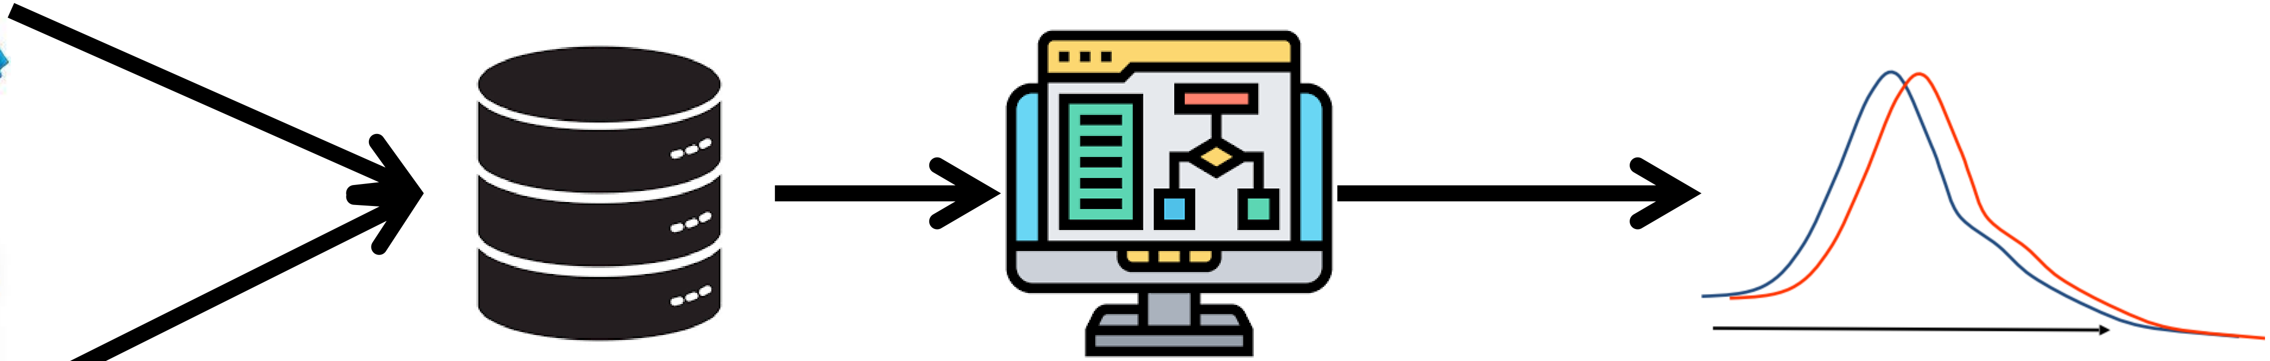

Sensitive dataset Algorithm Output distribution

Last Time: Differential Privacy

- [DMNS06] Given $\varepsilon > 0$ and $\delta \in (0,1)$, a randomized algorithm $A: U^* \to Y$ is (ε, δ)-differentially private if, for every neighboring pair D and D' of datasets, and for all $E \subseteq Y$,
$$\Pr[A(D) \in E] \leq e^{\varepsilon} \cdot \Pr[A(D') \in E] + \delta$$

- Implication: Deterministic algorithms cannot be differentially private unless they are a constant function

Last Time: Local Differential Privacy (LDP)

- [KLNRS08] Given $\varepsilon > 0$ and $\delta \in (0,1)$, a randomized algorithm $A: U^* \rightarrow Y$ is (ε, δ)-differentially private if, for every pairs of users’ possible data x and x' and for all $E \subseteq Y$,
$$\Pr[A(x) \in E] \leq e^{\varepsilon} \cdot \Pr[A(x') \in E] + \delta$$

- Algorithm takes a single user's data
- Compared to previous definition of DP, where algorithm takes all users' data

Randomized Response, Revisited

- How many people in this class have a pet?

- Generate a random integer:
  - If it is even, answer truthfully
  - Otherwise, proceed below
- Generate another random integer:
  - If it is even, answer YES
  - Otherwise if it is odd, answer NO

Randomized Response, Revisited

- $\Pr[Y_i = X_i] = \frac{3}{4}$ and $\Pr[Y_i = 1 - X_i] = \frac{1}{4}$
- $\mathrm{E}[Y_i] = \frac{3}{4} \cdot X_i + \frac{1}{4} \cdot (1 - X_i) = \frac{X_i}{2} + \frac{1}{4}$
- Let $Y = \frac{Y_1 + \ldots + Y_n}{n}$ and $X = \frac{X_1 + \ldots + X_n}{n}$

- $\mathrm{E}[Y] = \frac{X}{2} + \frac{1}{4}$
- Report $2\left(Y - \frac{1}{4}\right)$ for true fraction

Randomized Response, Revisited

- Answer is correct in expectation, but what is its variance?
- Each answer is incorrect with probability $\frac{1}{4}$
- The variance is $O(n)$

Randomized Response, Revisited

- Answer is correct in expectation, but what is its variance?
- Each answer is incorrect with probability $\frac{1}{4}$
- The variance is $O(n)$

- By Chebyshev, we have additive error $O(\sqrt{n})$ with probability 0.99
- By anti-concentration, error is $\Omega(\sqrt{n})$

Privacy and Noise

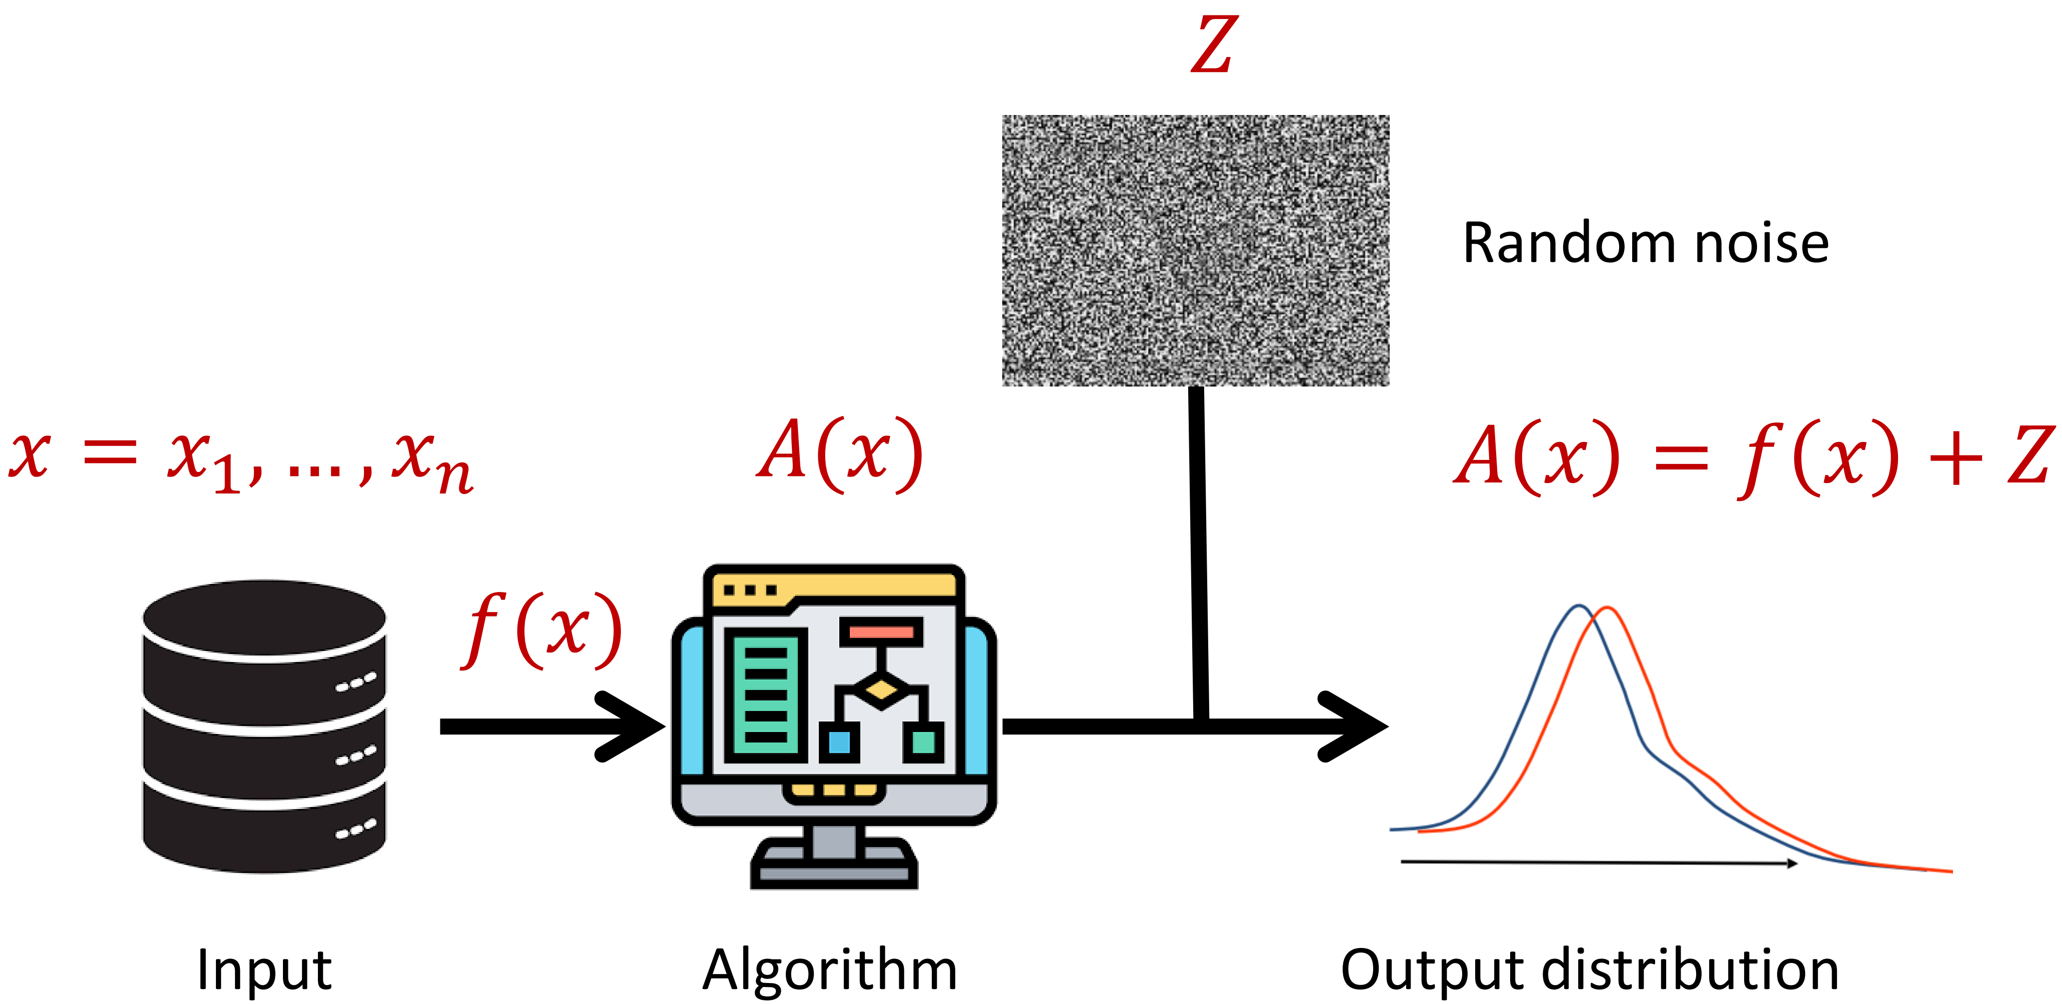

Proving Differential Privacy

- For pure DP, we want to show that for all for all E, $\Pr[A(D) \in E] \leq e^{\varepsilon} \cdot \Pr[A(D') \in E]$
- We have $\Pr[A(D) \in E] = \sum_{Z \in E} \Pr[A(D) = Z]$

- It suffices to show that for all Z in the support of E, $\Pr[A(D) = Z] \leq e^{\varepsilon} \cdot \Pr[A(D') = Z]$

Differential Privacy Properties

- What properties would we like from a rigorous definition of privacy?

Post-processing of Differential Privacy

- Ability to handle post-processing
  - If mechanism A has privacy loss ε and we release $g(A(D))$, then we have privacy loss ε

Post-processing of Differential Privacy

- For a fixed W, want to show

$$\Pr\left[g\left(A(D)\right) = W\right] \le e^{\varepsilon} \cdot \Pr\left[g\left(A(D')\right) = W\right]$$

- Suppose mechanism A has privacy loss ε
- Consider all possible Z such that $g(Z) \to W$ (under some randomness)
- Then for neighboring datasets D and D', we have $\Pr[A(D) = Z] \le e^{\varepsilon} \cdot \Pr[A(D') = Z]$ for all Z

Composition of Differential Privacy

- Ability to compose multiple private mechanisms
- Privacy loss measure ε accumulates across multiple computations and datasets
  - If mechanism M_1 has privacy loss ε_1 and mechanism M_2 has privacy loss ε_2, then releasing the results of both M_1 and M_2 has privacy loss $\varepsilon_1 + \varepsilon_2$

Composition of Differential Privacy

- Let $g(\cdot,\cdot)$ be a composition function
- Suppose $g(X, Y) = Z$, so that $g\big(M_1(D), M_2(D)\big) = Z$ if $M_1(D) = X$ and $M_2(D) = Y$
- What is $\Pr[g(M_1(D) = X]$? What is $\Pr[M_2(D) = Y]$?
- What is $\Pr[g(M_1(D') = X]$? What is $\Pr[M_2(D') = Y]$?

Composition of Differential Privacy

- We have:

$$\frac{\Pr[g(M_1(D) = X] \cdot \Pr[M_2(D) = Y]}{\Pr[g(M_1(D') = X] \cdot \Pr[M_2(D') = Y]} \leq e^{\varepsilon_1} \cdot e^{\varepsilon_2}$$

- So:

$$\Pr\left[g\left(M_1(D), M_2(D)\right) = Z\right] \leq e^{\varepsilon_1 + \varepsilon_2} \cdot \Pr\left[g\left(M_1(D), M_2(D)\right) = Z\right]$$

Basic Composition of Differential Privacy

- If mechanisms $M_1, \ldots, M_k$ are $(\varepsilon_1, \delta_1), \ldots, (\varepsilon_k, \delta_k)$-DP, then for any function g, then the composition mechanism $g(M_1, \ldots, M_k)$ is $(\varepsilon_1 + \cdots + \varepsilon_k, \delta_1 + \cdots + \delta_k)$-DP

Differential Privacy Properties

- Ability to handle post-processing
  - If mechanism A has privacy loss ε and we release $g(A(D))$, then we have privacy loss ε
- Privacy loss measure ε accumulates across multiple computations and datasets
  - If mechanisms $M_1, \dots, M_k$ are $(\varepsilon_1, \delta_1), \dots, (\varepsilon_k, \delta_k)$-DP, then for any function g, then the composition mechanism $g(M_1, \dots, M_k)$ is $(\varepsilon_1 + \cdots + \varepsilon_k, \delta_1 + \cdots + \delta_k)$-DP

Towards Differential Privacy

- Suppose everyone has an integer from $[1, 10]$

- How to privately compute the max of the integers?
- How to privately compute the sum of the integers?
- How to privately release a histogram of the integers?

Privacy and Noise

- Goal: release private approximation to $f(x)$
- Intuition: $f(x)$ can be released accurately if the function f is not sensitive to changes by any of the individuals $x = x_1, \dots, x_n$
- Sensitivity: $\sigma_f = \max_{\text{neighbors } x, x'} \|f(x) - f(x')\|_1$

Sensitivity

- Sensitivity: $\sigma_f = \max_{\text{neighbors } x, x'} \|f(x) - f(x')\|_1$

- Suppose a study is conducted that measures the height of individuals, ranging from 1 to 300 centimeters

- What is the sensitivity of the maximum height query?
- What is the sensitivity of the average height query?

Laplace Mechanism

- Goal: Algorithm computes $f(x)$ and releases $f(x) + Z$, where $Z \sim \mathrm{Lap}\left(\frac{\sigma_f}{\varepsilon}\right)$

- Laplacian distribution: Probability density function for $\mathrm{Lap}(b)$ is

$$p(x) = \frac{1}{2b}\exp\left(-\frac{|x|}{b}\right) = \frac{1}{2b}e^{\left(-\frac{|x|}{b}\right)}$$

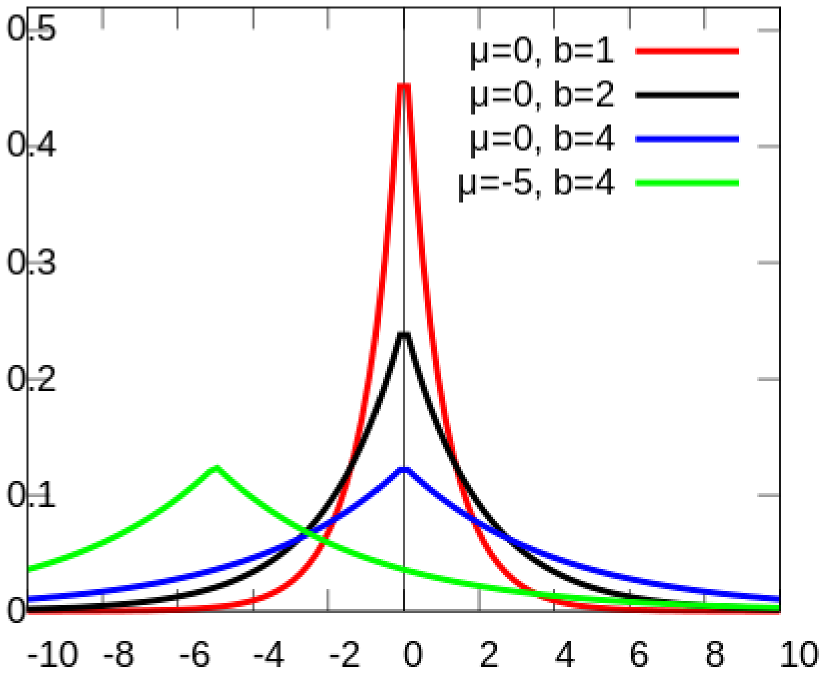

Laplace Mechanism

- What does the Laplace mechanism do in the following cases?

- Suppose a study is conducted that measures the height of individuals, ranging from 1 to 300 centimeters

- What is the sensitivity of the maximum height query?
- What is the sensitivity of the average height query?

Laplace Mechanism

- Theorem: Laplace mechanism is ε-differentially private (pure DP)

Laplace Mechanism

- Suppose the "answer" for dataset D is z and the "answer" for dataset D is z'
- Let $y = z + \mathrm{Lap}\left(\frac{\sigma_f}{\varepsilon}\right)$ and let $y' = z' + \mathrm{Lap}\left(\frac{\sigma_f}{\varepsilon}\right)$
- Want to show that for any fixed x, $\Pr[y = x] \leq e^{\varepsilon} \cdot \Pr[y' = x]$

Laplace Mechanism

- Want to show that for any fixed x, $\frac{\Pr[y=x]}{\Pr[y'=x]} \leq e^{\varepsilon}$
- Let Z be a draw from a Laplace distribution
- $\Pr[y = x] = \Pr[Z = x - z]$
- $\Pr[y' = x] = \Pr[Z = x - z']$

Laplace Mechanism

- $\Pr[y = x] = \Pr[Z = x - z]$
- $\Pr[y' = x] = \Pr[Z = x - z']$
- $p(x) = \frac{\varepsilon}{2\sigma_f} \exp\left(-\frac{\varepsilon|x|}{\sigma_f}\right) = \frac{\varepsilon}{2\sigma_f} e^{\left(-\frac{\varepsilon|x|}{\sigma_f}\right)}$ for a Laplace distribution with scale parameter $\frac{\sigma_f}{\varepsilon}$
- $\frac{\Pr[Z=x-z]}{\Pr[Z=x-z\prime]} = e^{\left(-\frac{\varepsilon|x-z|}{\sigma_f}\right)} / e^{\left(-\frac{\varepsilon|x-z\prime|}{\sigma_f}\right)} \leq e^{\left(-\frac{\varepsilon|z-z\prime|}{\sigma_f}\right)}$

Differential Privacy from Laplace Mechanism

- Suppose everyone has an integer from $[1, 10]$

- How to privately compute the max of the integers?
- How to privately compute the sum of the integers?
- How to privately release a histogram of the integers?

Laplace Mechanism for Vectors

- Given a function $f: U \to \mathbb{R}^n$, the Laplace mechanism is defined by $f(D) + v$, where each entry of v is drawn from $\mathrm{Lap}\left(\frac{\sigma_f}{\varepsilon}\right)$

- Theorem: Laplace mechanism is ε-differentially private (pure DP)
- Proof follows by generalizing the previous proof, decomposing along the coordinates