CSC2412 F19: Assignment 1

Due: October 25, by 11:59pm

Guidelines:

- Your assignment solution must be submitted as a *typed* PDF document. Scanned handwritten solutions, solutions in any other format, or unreadable solutions will **not** be accepted or marked. You are encouraged to learn the LaTeX typesetting system and use it to type your solution.
- To submit this assignment, use the MarkUs system, at URL `https://markus.teach.cs.toronto.edu/csc2412-2019-09`
- This is an *individual assignment*. You may consult any of the reading material posted on the course website. However, your solutions should show your individual work.
- You may use any result discussed in class or covered in the assigned reading by just referring to it. You do not need to reproduce proofs that we have covered in the lectures.
- Unless stated otherwise, you should justify all your answers using rigorous arguments. Your solution will be marked based both on its completeness and correctness, and also on the clarity and precision of your explanation.

For all questions you can assume that the notion of neighbouring datasets used is the one based on replacement, and the size of the dataset is public. I.e. datasets $X, X' \in \mathcal{X}^n$ are neighbouring if there exists an $i \in \{1, \dots, n\}$ such that $x'_j = x_j$ whenever $j \neq i$.

Question 1. (5 marks)

Suppose you are given a private dataset $X \in \mathcal{X}^n$, where $\mathcal{X} = \{1, \dots, N\}$. I.e. the dataset X consists of n integers $x_1, \dots, x_n$ between 1 and N. Describe an ε-differentially private algorithm, based on the exponential mechanism, which outputs a number $y \in \mathcal{X}$ such that if $n \geq \frac{C_1}{\varepsilon} \ln(|N|/\beta) + C_2$, then

$$\mathbb{P}\left(\min_{i=1}^{n} x_i \leq y \leq \max_{i=1}^{n} x_i\right) \geq 1 - \beta.$$

Above C_1 and C_2 are constants independent of n, N, β, and ε. Justify why your algorithm is ε-differentially private, and why it satisfies the property above. Specify the constants C_1, C_2 in your answer.

[Solution]

Use the exponential mechanism with output range $\mathcal{X}$ and utility function $u : \mathcal{X}^n \times \mathcal{X}$ defined by

$$u(X, y) = \min\{|\{i : x_i \leq y\}|, |\{i : x_i \geq y\}|\}.$$

Note that y is between $\min_{i=1}^{n} x_i$ and $\max_{i=1}^{n} x_i$ exactly when $u(X, y) \geq 1$. Also, $\mathrm{OPT}(X) = \max_{y \in \mathcal{X}} u(X, y)$ satisfies

$$\mathrm{OPT}(X) \geq \left\lfloor \frac{n}{2} \right\rfloor,$$

because we can pick y to be the $\lfloor \frac{n}{2} \rfloor$-th element in the sorted order of X.

The sensitivity Δu of u is at most 1. To see this, note that for any two neighbouring datasets X and X',

$$\big| |\{i : x_i \leq y\}| - |\{i : x'_i \leq y\}| \big| \leq 1,$$

because in the worst case an element of X that is less than or equal to y is replaced in X' by one that is greater than y, or vice versa. Similarly,

$$\big| |\{i : x_i \geq y\}| - |\{i : x'_i \geq y\}| \big| \leq 1.$$

Therefore,

$$|u(X, y) - u(X', y)| \leq \max\{\big| |\{i : x_i \leq y\}| - |\{i : x'_i \leq y\}| \big|, \big| |\{i : x_i \geq y\}| - |\{i : x'_i \geq y\}| \big|\} \leq 1.$$

Because we took X and X' to be arbitrary neighbouring datasets, this shows that $\Delta u \leq 1$.

Now, by the utility guarantee for the exponential mechanism shown in class, we have that for the random output y generated by the mechanism

$$\mathbb{P}\left(u(X, y) \geq \mathrm{OPT}(X) - \frac{2}{\varepsilon} \ln(|N|/\beta)\right) \geq 1 - \beta.$$

Since $\mathrm{OPT}(X) \geq \lfloor \frac{n}{2} \rfloor \geq \frac{n-1}{2}$, if

$$n \geq \frac{4}{\varepsilon} \ln(|N|/\beta) + 3,$$

then $\mathrm{OPT}(X) - \frac{2}{\varepsilon} \ln(|N|/\beta) \geq 1$, and we have

$$\mathbb{P}(u(X, y) \geq 1) \geq 1 - \beta.$$

As remarked above, $u(X, y) \geq 1$ implies that y is between $\min_{i=1}^{n} x_i$ and $\max_{i=1}^{n} x_i$. This proves the desired property of the algorithm.

The fact that the algorithm is ε-DP follows from the privacy analysis of the exponential mechanism from class.

Question 2. (10 marks)

For this question you can use the following identity: for a Laplace random variable $w \sim \mathrm{Lap}(b)$, we have for any $t \geq 0$

$$\mathbb{P}(|w| \geq t) = e^{-t/b}.$$

The goal in this question is to design an algorithm which estimates the mean of a dataset of numbers. The estimate should be accurate whenever the numbers are bounded, but the algorithm should be private even if the numbers are arbitrary.

Part a. (4 marks)

Suppose that the dataset $X = (x_1, \dots, x_n)$ consists of integers, which can be positive or negative, and are not *a priori* bounded. Describe an ε-differentially private algorithm $\mathcal{A}$ such that, if $x_i \in [-B, +B]$ for every i, and $n \geq \frac{2B \ln(1/\beta)}{\varepsilon\alpha}$, then

$$\mathbb{P}\left(\left|\mathcal{A}(X) - \frac{1}{n}\sum_{i=1}^{n} x_i\right| \geq \alpha\right) \leq \beta. \tag{1}$$

You can assume that the parameter B is known to the algorithm. Note that while (1) needs to hold only if $x_i \in [-B, +B]$ for all i, the algorithm $\mathcal{A}$ needs to be ε-differentially private for **every** dataset, i.e. even if $x_i \notin [-B, B]$ for some values of i. Justify your answer.

[Solution]

Let $g : \mathbb{Z} \to \mathbb{Z}$ be defined by $g(x) = x$ if $x \in [-B, +B]$, and $g(x) = 0$ otherwise. Then define $f : \mathbb{Z}^n \to \mathbb{R}$ by

$$f(X) = \frac{1}{n}\sum_{i=1}^{n} g(x_i).$$

Clearly if $\forall i : x_i \in [-B, +B]$, then $f(X) = \frac{1}{n}\sum_{i=1}^{n} x_i$. Therefore, it is sufficient to design an ε-DP algorithm $\mathcal{A}$ such that

$$\mathbb{P}(|\mathcal{A}(X) - f(X)| \geq \alpha) \leq \beta.$$

Notice that the sensitivity of f is $\frac{2B}{n}$: for any two neighbouring datasets X and X' differing in their i-th entry, we have

$$f(X) - f(X') = \frac{g(x_i) - g(x_i')}{n} \leq \frac{2B}{n},$$

since $g(x_i), g(x_i') \in [-B, +B]$. We can then just define $\mathcal{A}$ to be the Laplace noise mechanism: $\mathcal{A}(X) = f(X) + w$, where $w \sim \mathrm{Lap}\left(\frac{2B}{\varepsilon n}\right)$. We proved that this mechanism is ε-DP in class.

To show (1), we notice that $\mathcal{A}(X) - f(X) = w$, so

$$\mathbb{P}(|\mathcal{A}(X) - f(X)| \geq \alpha) = \mathbb{P}(|w| \geq \alpha) = e^{-\frac{\varepsilon n}{2B}},$$

with the right hand side at most β whenever $n \geq \frac{2B \ln(1/\beta)}{\varepsilon\alpha}$.

Part b. (6 marks)

Suppose that the dataset $X = (x_1, \dots, x_n)$ consists of integers in $[-N, +N]$, where N is some large integer. Describe an ε-differentially private algorithm $\mathcal{A}$ such that, if there exists some integer $z \in [-N, +N]$ for which $x_i \in [z - B, z + B]$ for every i, and

$$n \geq \max\left\{\frac{C_1 B \ln(2/\beta)}{\varepsilon\alpha}, \frac{C_2 \ln(2|2N+1|/\beta)}{\varepsilon} + C_3\right\}$$

then

$$\mathbb{P}\left(\left|\mathcal{A}(X) - \frac{1}{n}\sum_{i=1}^{n} x_i\right| \geq \alpha\right) \leq \beta. \tag{2}$$

Above C_1, C_2, and C_3 are constants independent of n, N, β, ε, B, and z. You can assume that the parameter B is known to the algorithm, but the parameter z is **not** known. Once again, the algorithm $\mathcal{A}$ needs to be

ε-differentially private for **every** $X \in \{-N, \dots, +N\}^n$. Justify your answer, and specify the constants C_1, C_2, and C_3.

Hint: You can use Question 1 to help you solve this subquestion.

[Solution]

First we use the algorithm from Question 1, with privacy parameter $\frac{\varepsilon}{2}$, to find an integer $y \in [-N, +N]$ such that, with probability at least $1 - \frac{\beta}{2}$, we have

$$\min_{i=1}^{n} x_i \leq y \leq \max_{i=1}^{n} x_i \tag{3}$$

By Question 1, the algorithm will have this property as long as

$$n \geq \frac{8 \ln(2|2N+1|/\beta)}{\varepsilon} + 3.$$

Notice that if (3) holds, and if $\forall i : x_i \in [z - B, z + B]$ for some z, then $|y - z| \leq B$, and, therefore, $\forall i : x_i \in [y - 2B, y + 2B]$.

Next we use a variant of the algorithm from the previous subproblem. We define a function $g : \mathbb{Z} \to \mathbb{Z}$ by

$$g(x) = \begin{cases} x & x \in [y - 2B, y + 2B] \\ 0 & \text{otherwise} \end{cases},$$

and $f(X) = \frac{1}{n} \sum_{i=1}^{n} g(x_i)$. As observed above, if (3) holds, and also $\forall i : x_i \in [z - B, z + B]$ for some z, then $f(X) = \frac{1}{n} \sum_{i=1}^{n} x_i$.

Analogously to the previous subquestion, the sensitivity of $f(X)$ is $\frac{4B}{n}$. We define the algorithm $\mathcal{A}$ to be the Laplace noise mechanism applied to $f(X)$ with privacy parameter $\frac{\varepsilon}{2}$: $\mathcal{A}(X) = f(X) + w$ where $w \sim \text{Lap}\left(\frac{8B}{\varepsilon n}\right)$. Analogously to above, if $n \geq \frac{8B \ln(2/\beta)}{\varepsilon \alpha}$, then

$$\mathbb{P}(|\mathcal{A}(X) - f(X)| \geq \alpha) \leq 1 - \frac{\beta}{2}.$$

Then, by the composition theorem, the composition of the algorithm from Question 1 and $\mathcal{A}$ will be ε-DP, and by the union bound both (3) and $|\mathcal{A}(X) - f(X)|$ will hold with probability at least $1 - \beta$, which implies (2).