Homework 2 of DSC 291 Introduction to Differential Privacy (Fall 2024)

University of California, San Diego

Assigned on Oct 18, 2024(Friday)

Due at 11:59 pm on Nov 1, 2024 (Friday)

Notes:

- Be sure to read "Policy on Academic Integrity" on the course syllabus.
- There are ***[100 points]*** and a bonus ***[20 points]*** in this homework.
- You need to submit your homework via Gradescope.
- Contract the instructor if you spot typos. Any updates or correction will be posted on the course Announcements page and piazza, so check there occasionally.

---

0. Acknowledgment *[0 points]*

For each question in this HW, please list all your collaborators and reference materials (beyond those specified on the website) that were used for this homework.

1. **List of Collaborators** List the names of all people you have collaborated with and for which question(s).
2. **List of Acknowledgements.** If you find an assignment's answer or use a another source for help, acknowledge for which question and provide an appropriate citation (there is no penalty, provided you include the acknowledgement). If not, then write "none".

1 Privacy loss random variables [30 Pts + 15 Bonus point]

A central concept in DP analysis is the so-called privacy loss random variable $\boldsymbol{\varepsilon}$. Let randomized algorithm $\mathcal{M} : \mathcal{X}^* \to \mathcal{Y}$, where $\mathcal{X}^*$ is the space of the datasets and $\mathcal{Y}$ is the space of the outputs. The privacy loss random variable is defined for a fixed pair of neighboring datasets $x, x' \in \mathcal{X}^*$ as follows.

$$\boldsymbol{\varepsilon}_{\mathcal{M}}^{x,x'} = \log(\frac{p(\mathbf{y})}{p'(\mathbf{y})}) \text{ where random variable } \mathbf{y} \sim \mathcal{M}(x).$$

In the above, p, p' are the probability density function (or probability mass function if $\mathcal{Y}$ is discrete) of $\mathcal{M}(x)$ and $\mathcal{M}(x')$ respectively. For simplicity, you may ignore measure theoretic-considerations and just consider a finite discrete output space.

(a) (10 pts) Show that for all neighboring pair x, x' $\mathbb{P}(\boldsymbol{\varepsilon}_{\mathcal{M}}^{x,x'} > \epsilon) = 0$ *if and only if* $\mathcal{M}$ satisfies ϵ-DP.

(b) (10 pts) Prove the following useful lemma.

Lemma 1 (Tail bound to (ϵ, δ)-DP conversion)**.** *Let* $\boldsymbol{\varepsilon}_{\mathcal{M}}^{x,x'}$ *be the privacy loss RV defined above. If*

$$\mathbb{P}(\boldsymbol{\varepsilon}_{\mathcal{M}}^{x,x'} > \epsilon) \leq \delta$$

for all *pair of neighboring* x, x' *then* $\mathcal{M}$ *satisfies* (ϵ, δ)*-DP.*

(Hint: For any set of the output space $\mathcal{S}$, consider both $\mathcal{S}$ and its complement $\mathcal{S}^c$. Follow the definition of DP.)

(c) (10 pts) (non-adaptive composition) Use the above to show that for two *fixed* mechanisms $\mathcal{M}_1, \mathcal{M}_2$. Let's say $\mathcal{M}_1, \mathcal{M}_2$ satisfies that for any pair of neighboring input x, x',

$$\mathbb{P}(\boldsymbol{\varepsilon}_{\mathcal{M}_1}^{x,x'} > \epsilon) \leq \delta$$

and

$$\mathbb{P}(\boldsymbol{\varepsilon}_{\mathcal{M}_2}^{x,x'} > \epsilon) \leq \delta$$

show that the composition $\mathcal{M} = (\mathcal{M}_1, \mathcal{M}_2)$ satisfies that $(2\epsilon, 2\delta)$-DP.

(Hint: Use Lemma 1. Apply the **Law of Total Probability** on the event $\varepsilon_{\mathcal{M}_1}^{x,x'} > \epsilon$)

(d) (5 pts) (Bonus, Adaptive composition) Let's say that $\mathcal{M}_2$ depends on the realized output of $\mathcal{M}_1$, but for each potential output $y_1 \in \mathcal{Y}_1$, $\mathcal{M}_2$ satisfies $\mathbb{P}(\boldsymbol{\varepsilon}_{\mathcal{M}_2}^{x,x'} > \epsilon) \leq \delta$ for any pair of neighboring input x, x', . Show that the composition $\mathcal{M} = (\mathcal{M}_1, \mathcal{M}_2)$ satisfies $(2\epsilon, 2\delta)$-DP.

(Hint: The analysis should be more or less the same as Part (c))

(e) (10 pts) (Bonus, Advanced Composition) Let $\mathcal{M}_1, ..., \mathcal{M}_k$ be a sequence of randomized algorithms, $\mathcal{M}_i$ may depend on the realized output of $\mathcal{M}_1, ..., \mathcal{M}_{i-1}$. Assume that for all i, $\mathcal{M}_i$ satisfies that $\mathbb{P}(\boldsymbol{\varepsilon}_{\mathcal{M}_1}^{x,x'} > \epsilon) \leq 0$ show (using Azuma-Hoeffding's inequality) that the composition $(\mathcal{M}_1, ..., \mathcal{M}_k)$ satisfies (ϵ', δ)-DP with

$$\epsilon' = k\epsilon(e^\epsilon - 1) + \sqrt{2k\log(1/\delta)}\epsilon.$$

(Hint: The challenge is in bounding the expected value of the privacy loss RV with $\epsilon(e^\epsilon - 1)$. Notice that the expected value of the privacy loss random variable is the KL-divergence, which is always larger than 0.)

2 Practice on Laplace and Gaussian tail bounds [30 pts + 5 pts bonus]

(I have covered many of these in the lectures. The point of this question is to make sure you understand the lecture. If you don't follow the lecture, feel free to read the textbook, understand the proof and write them down.)

(a) and (b) ask you to practice manipulating the tail bounds of probability distributions. (c) and the bonus (d) are about privacy loss random variables.

(a) (10 pts) Using the tail bound of Laplace distribution and union bound, work out a high probability bound (with probability $1-\beta$) of the L_∞ error (maximum error in all coordinates) of Laplace mechanism satisfying ϵ-DP when applying it to a query f : DataSet Space $\to \mathbb{R}^d$ satisfying that its L1-sensitivity $\Delta_1(f) = 1$.

(b) (10 pts) Repeat the above for Gaussian mechanism under the same assumptions for (ϵ, δ)-DP. For simplicity, you many assume $\epsilon < 1$ and use the expression for the classical Gaussian mechanism (i.e., the expression with $\sigma = \frac{\Delta_2}{\epsilon}\sqrt{2\log(1.25/\delta)}$). Also, you may use the subgaussian tail bound that we covered in the lecture.

(c) (10 pts) Show that the Privacy loss random variable of the Gaussian mechanism is Gaussian.

(d) (5 pts bonus) Work out the distribution of the privacy loss random variable for the Laplace mechanism.

3 Comparing statistical error and the additional error due to DP [40 pts]

(a) (8 pts) Assume a dataset $\phi_1, ..., \phi_n \in \mathbb{R}^d$ is drawn i.i.d. from an unknown distribution supported on $[0, 1]^d$. Apply Hoeffding's inequality and union bound to derive the L_∞-norm error of the mean estimator $\|\frac{1}{n}\sum_i \phi_i - \mathbb{E}[\phi_1]\|_\infty$ with probability $1 - \beta$.

(Hint: apply Hoeffding's inequality to every coordinate, then union bound)

(b) (8 pts) Derive the L1 and L2 sensitivity of the mean estimator $\frac{1}{n}\sum_i \phi_i$ under the "Replace One" neighboring relationship (i.e., n is public).

(c) (8 pts) Work out the amount of noise to add for Laplace mechanism (the b parameter) and Gaussian mechanism (the σ parameter) to achieve ρ-zCDP with $\rho = 1$. What is the high-probability bound you get (you should've already calculated the expression from Q2(a) and Q2(b), just plug in the b and σ parameter you get in this question)?

(d) (16 pts) You may wonder whether the comparison above is fair. After all, we are comparing upper bounds with upper bounds. It is possible that the bounds are not tight. Set up an experiment, for the cases when each coordinate of ϕ_1 is drawn iid from a Bernoulli distribution with parameter $0 < \mu < 1$. Choose $n = 10,000$, $d = 5$. Plot the statistical error $\|\frac{1}{n}\sum_i \phi_i - \mathbb{E}[\phi_1]\|_\infty$ and the additional error from Laplace mechanism and Gaussian mechanism against $\mu = [0.001, 0.01, 0.02, ..., 0.98, 0.99, 0.999]$. Repeat the same plot for $d = 100$.

(For the plot to look pretty, you may need to use logarithmic scale for the y-axis)