Homework 2 of CS291A Introduction to Differential Privacy (Fall 2021)

University of California, Santa Barbara

Assigned on Oct 25, 2021(Monday)

Due at 11:59 pm on Nov 15, 2021 (Monday)

Notes:

- Be sure to read "Policy on Academic Integrity" on the course syllabus.
- There are ***[100 points]*** in this homework, and a bonus ***[5 points]***.
- 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 Warm-up [35 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. (e),(f) and (g) are about concentrated DP, Renyi DP and their composition.

(a) (5 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) (5 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) (5 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.

(e) (5 pts) Rewrite the definition of Renyi-divergence of $\mathcal{M}(x)$ and $\mathcal{M}(x')$ for two fixed neighboring datasets x and x' in terms of the moments generating function of the privacy loss random variable.

(f) (5 pts) Show that if $\mathcal{M}_1$ satisfies ρ_1-zCDP and $\mathcal{M}_2$ satisfies ρ_2-zCDP conditioning on any possible output of $\mathcal{M}_1$, then the composition $(\mathcal{M}_1, \mathcal{M}_2)$ satisfy $(\rho_1 + \rho_2)$-zCDP.

(g) (5 pts) Show that if $\mathcal{M}$ satisfies $(\alpha, \epsilon(\alpha))$-Renyi Differential privacy for all $\alpha > 1$, then $\mathcal{M}$ also satisfies (ϵ, δ)-DP for any $0 < \delta < 1$ and

$$\epsilon(\delta) = \min_{\alpha > 1} \epsilon(\alpha) + \frac{\log(1/\delta)}{\alpha - 1}.$$

(h) (5 pts) Derive the optimal α and the expression of $\epsilon(\delta)$ above when $\mathcal{M}$ is a composition of k-Gaussian mechanisms, each satisfying ρ-zCDP.

2 Comparing statistical error and the additional error due to DP [25 pts]

(a) (5 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 $1 - \beta$.

(b) (5 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) (5 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 Q1(a) and Q1(b), just plug in the b and σ parameter you get in this question)?

(d) (10 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)

3 Equivalent definitions of DP [20 pts]

(a) (5 pts) Show that $\max_{D \simeq D'} H_{e^\epsilon}(\mathcal{M}(D) \| \mathcal{M}(D')) \leq \delta$ is implied by (ϵ, δ)-DP.

(Note: Lecture 7 has a proof for how a HS-divergence bound implies (ϵ, δ)DP. Your proof of the the above will conclude the equivalence of the two.)

(b) (5 pts) Show that (ϵ, δ)-DP implies a tradeoff function lower bound of the form

$$\max\{0, 1 - \delta - e^\epsilon \alpha, e^{-\epsilon}(1 - \delta - \alpha)\}.$$

(Note: This was covered several times and the instructor sketched a proof in Slide 22 of Lecture 7. This question asks you to write it down formally.)

(c) (10 pts) Consider the Leaky Randomized Response Mechanism from Slide 16 of Lecture 8, show that this is the worst-possible (ϵ, δ)-DP mechanism by showing that the *likelihood ratio test* (defined below) with particular choice of its two parameters attains every point of the tradeoff function of an (ϵ, δ)-DP mechanism.

The family of likelihood ratio test for testing two distributions P, Q based on a sample y is defined as follows:

$$\mathrm{LRT}_{\eta,c}(y) = \begin{cases} P & \text{if } \log \frac{p(y)}{q(y)} > \eta \\ Q & \text{if } \log \frac{p(y)}{q(y)} < \eta \\ \text{output } P \text{ with probability } c & \text{if } \log \frac{p(y)}{q(y)} = \eta \end{cases}$$

The two parameters are η, c.

(Hint: it helps to draw the tradeoff function and inspect where each parameter pair of LRT ends up at.)

(Remark: This would certify that leaky RR is a tight dominating pair of the (ϵ, δ)-DP mechanism and to prove the advanced composition of an arbitrary sequence of (ϵ, δ)-DP mechanism, it suffices to prove a tail bound of the sum of PLRVs from leaky RR.)

4 Differentially Private (Binary) Logistic Regression[20 Pts]

In this question, you will read two tutorials of DP linear regression (in Jupyter Notebooks) and modify the second Notebook to implement the noisyGD algorithm for DP logistic regression.

Let the feature space to be $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_2 \leq 1\}$ and the label-space to be $\mathcal{Y} = \{0, 1\}$. Logistic regression minimizes the cross-entropy loss for all training data points $(x_1, y_1), ..., (x_n, y_n)$, i.e.,

$$\hat{\theta} = \arg\min_\theta \sum_{i=1}^n \ell(\theta, (x_i, y_i))$$

where

$$\ell(\theta, (x, y)) = -\Big(y \log(\frac{e^{x^T\theta}}{1 + e^{x^T\theta}}) + (1 - y) \log(\frac{1}{1 + e^{x^T\theta}})\Big).$$

(a) (5 pts) Denote $\mathcal{L}(\theta, \mathrm{Data}) = \sum_{i=1}^n \ell(\theta, (x_i, y_i))$. What is the L_2 sensitivity of $\nabla_\theta \mathcal{L}(\theta, \mathrm{Data})$ at a given θ?

(b) (15 pts)Implement the noisy Gradient Descent mechanism for linear logistic regression by iteratively applying Gaussian mechanism by fill in the details in the provided notebook.

You will need to 1. implement noisy gradient descent linear logistic regression. 2. represent noisy GD in autodp and use calibrator to figure out the 3. Run experiments to demonstrate the privacy-utility tradeoff as we adjust the ϵ parameters.

To ensure that the method converges, you may use the theoretical learning rate

$$\min\{1/L, \sqrt{\frac{2(\mathcal{L}(\theta_0) - \mathcal{L}^*)}{d\sigma^2 LT}}\}$$

where L is the gradient Lipschitz constant of $\mathcal{L}$, σ is the std of the Gaussian noise added by the Gaussian mechanism to each coordinate (Not the noise multiplier!) and T is the number of iterations to run this algorithm.

(Hint: 1. The gradient Lipschitz constant L is an upper bound of the largest eigenvalue of the Hessian $\nabla^2\mathcal{L}$. What is a crude upper bound of it? 2. Suppose we initialize at 0 what is an upper bound of $\mathcal{L}(\theta_0) - \mathcal{L}^*$?)