Privacy in Statistics and Machine Learning Spring 2023
In-class Exercises for Lecture 10 (Advanced Composition)
February 23, 2023

Adam Smith (based on materials developed with Jonathan Ullman)

Problems with marked with an asterisk (*) are more challenging or open-ended.

1. (Exercise 2.1 From Lecture Notes) What is the privacy loss $I_{\mathbf{x},\mathbf{x}'}(y)$ when A is the Laplace mechanism in one dimension? To make things concrete: assume $f(\mathbf{x}) = 0$ and $f(\mathbf{x}') = 1$ and we add noise $\mathsf{Lap}(1/\varepsilon)$.

   Write out $I_{\mathbf{x},\mathbf{x}'}(y)$ as a function of y. Show that its expectation is $\Theta(\varepsilon^2)$ (assuming $\varepsilon \leq 1$) when the input is drawn according to $A(\mathbf{x})$.

2. Suppose we add *uniform* noise to a count query, that is, we release $A_\lambda(\mathbf{x}) = f(\mathbf{x}) + U_{[-\lambda,\lambda]}$ where f counts how many records staistfy some condition, and $U_{[-\lambda,\lambda]}$ is uniformly distributed in the interval $[-\lambda, \lambda]$.

   (a) Assuming that $f(\mathbf{x}) = 0$ and $f(\mathbf{x}') = 1$ and $\lambda > 1/2$, what is the distribution of $I_{\mathbf{x},\mathbf{x}'}(Y)$ when $Y \sim A_\lambda(x)$?

   (b) How large must λ be to satisfy (ε, δ)-DP? Do both ε and δ matter in setting λ? When $\delta < 1/n$, will this mechanism produce useful information?

3. (*) Prove the Simulation Lemma (Lemma 3.1 from the lecture notes) for the special case of the Laplace mechanism (with $\delta = 0$).

4. **(Differentially private top-k selection)** Suppose we have d candidate items and a score function $q : [d] \times \mathcal{U}^n \to \mathbb{R}$. In the selection problem of Lecture 6, we aimed to find a single high-score item. Suppose we now want to find $k < \frac{d}{2}$ high-score items.

   Given an algorithm that outputs a set of k items $S = A(\mathbf{x})$, we measure error as follows: let $q_{(k)}(\mathbf{x})$ be the score of the k-th best item (so $q_{(1)}$ is the maximum score). The error of the algorithm is

   $$q_{(k)}(\mathbf{x}) - \min_{j \in S} q(j; \mathbf{x}) \,.$$

   What error guarantee can you prove for the algorithm that proceeds by repeating the exponential mechanism k times to sample items without replacement? (E.g., “With probability at least 2/3, the error is at most blah. More generally, with probability at least $1 - \beta$, ...”)

5. **Histograms.** Consider the following algorithm for releasing histograms.

```
Algorithm 1: Stable Histogam($\mathbf{x}; \varepsilon, \delta$)
  Input: $\mathbf{x}$ is a multi-set of values in $\mathcal{U}$.
  for every $z \in \mathcal{U}$ that appears in $\mathbf{x}$ do $\tilde{c}_z = \#\{i : x_i = z\} + \mathsf{Lap}(1/\varepsilon)$s
  Release the set of pairs $\{(z, \tilde{c}_z) : \tilde{c}_z > \tau\}$ where $\tau = 1 + \frac{\ln(1/\delta)}{\varepsilon}$.
```

(a) Show that for any domain $\mathcal{U}$, Algorithm 1 is (ε, δ)-differentially private when neighboring data sets are allowed to differ by the insertion or deletion of one value.

Hint: The delicate part of this result is that we add noise only to counts of non-empty bins. (For example, if we were counting how many people live on each square mile of land in Alaska, most of the bins would be empty, but others would have lots of people.) There are two kinds of adjacent data sets: those where the set of nonempty bins changes, and those where it does not. You may need the following simple concentration bound for Laplace random variables: If $Y \sim \mathsf{Lap}(\lambda)$, then for every $t > 0$, we have $\Pr(Y > \lambda t) \leq \frac{1}{2}\exp(-t)$.

(b) Prove that the Stable Histograms algorithm is not $(\varepsilon', 0)$ differentially private for any finite positive value ε'. [*Hint*: Give two neighboring data sets and a histogram y such that y is a possible output for only one of the two data sets.]

6. **(Composing the Gaussian mechanism)**

(a) Consider a version of the Simulation Lemma that is specific to the Gaussian mechanism: show that for every function $f : \mathcal{U}^n \to \mathbb{R}$ with global sensitivity Δ, for every pair of neighboring datasets $\mathbf{x}, \mathbf{x}'$, there is a randomized algorithm F such that

- if $U \sim N(0, \sigma^2)$ then $F(U) \sim A_{f,\sigma}(\mathbf{x})$, and
- if $V \sim N(\Delta, \sigma^2)$ then $F(V) \sim A_{f,\sigma}(\mathbf{x}')$,

where $A_{f,\sigma}(\mathbf{x}) = f(x) + Z$ where $Z \sim N(0, \sigma^2)$.

[*Hint:* Consider F of the form $F(z) = az + b + N(0, \rho^2)$. Use a and b to get the means right, and use ρ to adjust the variance.]

(b) Use part (a) to show that the adaptive composition of k executions of the Gaussian mechanism with Δ-sensitive queries satisfies (ε, δ)-DP for $\sigma = \frac{\sqrt{2\ln(1/\delta)}}{\varepsilon}\Delta\sqrt{k}$. That is, it satisfies the same guarantee as does a single execution of the multi-dimensional Gaussian mechanism on a k-dimensional function with ℓ_2-sensitivity $\Delta\sqrt{k}$.

7. (*) Complete the sketched proof of the Simulation Lemma in the lecture notes.