Privacy in Statistics and Machine Learning — Spring 2025
In-class Exercises for Lecture 10 (Recap)
February 25, 2025

Adam Smith (based on materials developed with Jonathan Ullman)

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

1. **(Details of the simulation lemma)** Recall the proof sketch of the simulation lemma for $\delta = 0$. Let X and Y be random variables taking values in a discrete set $\mathcal{Z}$ whose distributions are ε-indistinguishable. For each element z, the system of equations on page 6 of the notes solves to

$$\begin{pmatrix} \Pr[F(0) = z] \\ \Pr[F(1) = z] \end{pmatrix} = \frac{1}{e^\varepsilon - 1} \begin{pmatrix} e^\varepsilon P_X(z) - P_Y(z) \\ e^\varepsilon P_Y(z) - P_X(z) \end{pmatrix}$$

   (a) Consider the Laplace mechanism example from a previous class: assume $\mathbf{x}$ and $\mathbf{x}'$ are enighboring datasets with $f(\mathbf{x}) = 0$, $f(\mathbf{x}') = 1$ and $A(\mathbf{x}) = f(\mathbf{x}) + \mathrm{Lap}(1/\varepsilon)$. Give a randomized algorithm F such that $A(\mathbf{x}) \sim F(RR_\varepsilon(0))$ and $A(\mathbf{x}') \sim F(RR_\varepsilon(1))$.
   Plot the densities of $F(0)$ and $F(1)$.

   (b) (*) Complete the proof sketch of the Simulation Lemma (Lemma 3.1 from the lecture notes) for the case $\delta = 0$ by showing that the probabilities (or densities) defined above are always are nonnegative and add (or integrate) to 1.

2. **(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-dimennsional Gaussian mechanism on a k-dimensional function with ℓ_2-sensitivity $\Delta\sqrt{k}$.

   (c) Show that randomized response RR_ε can be simulated from $U \sim N(0, 1)$ and $V \sim N(\Delta, 1)$ for some $\Delta = \Theta(\varepsilon)$, assuming $\varepsilon \leq 1$. (That is, give F such that $RR_\varepsilon(0) \sim F(U)$ and $RR_\varepsilon(1) \sim F(V)$.)
   (How does the required Δ behave when $\varepsilon \gg 1$?)

3. (**Stable Histograms**, from Lecture 9). 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 $\mathcal{U}$.

1 **for** *every* $z \in \mathcal{U}$ *that appears in* $\mathbf{x}$ **do**

2 $\quad \tilde{c}_z = \#\{i : x_i = z\} + \mathsf{Lap}(1/\varepsilon)$

3 Release the set of pairs $\{(z, \tilde{c}_z) : \tilde{c}_z > \tau\}$ where $\tau = 1 + \frac{\ln(1/\delta)}{\varepsilon}$.

This algorithm is remarkable because it adds noise only to the counts of *nonempty* bins, and the noise magnitude and threshold are independent of the total number of bins— in principle, the number of bins could be infinite. 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. This algorithm would reveal noisy counts only for sufficiently poulated areas.

(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. 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.]

4. (**More on node stability**) Recall from Homework 1 that two graphs are *node neighbors* if one can be obtained from the other by removing a node and all of its edges. Let $f(G)$ denote the number of triangles in an undirected graph G. You showed in the homework that, on the set of graphs with at most n vertices, the global sensitivity of f is $\binom{n-1}{2}$.

Now, given a parameter k and a graph G, consider the following linear program:

Variables:	x_t for every triangle t in G	
Constraints:	For every triangle t,	$0 \leq x_t \leq 1$, and
	for every node u,	$\sum_{t:t \text{ contains } u} x_t \leq k.$
Objective:	Maximize $\sum_t x_t$.	

Let $f_k(G)$ denote the value of this linear program (that is, the maximum possible value of the objective function). Show that

(a) Show that the value of this linear program equals the number of triangles in G *if and only if* every node u is contained in at most k triangles; and

(b) f_k has global sensitivity k (with no assumption on the neighboring inputs).