Privacy in Statistics and Machine Learning Spring 2023
In-class Exercises for Lecture 6 (Exponential Mechanism and RNM)
February 7, 2023

Adam Smith (based on materials developed with Jonathan Ullman)

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

1. Suppose we run the exponential mechanism to choose between just two output (say, election candidates Alice and Bob). Suppose Alice gets a votes and Bob gets b votes (where $a + b = n$ is the total number of votes).

   What are:

   (a) The output set $\mathcal{Y}$?

   (b) The score function q?

   (c) The sensitivity bound Δ for q?

   (d) The "odds ratio" $\frac{\Pr(Y=\text{'Alice'})}{\Pr(Y=\text{'Bob'})}$, assuming Y is the outcome of the exponential mechanism with for this problem with input $\varepsilon = 0.1$? (Express your answer as a function of $a - b$.

   (e) How big of a margin $a - b$ must Alice for her name to be output with probability at least 95%?

2. Suppose we run the exponential mechanism (or report-noisy-max/RNM) with outcome set $\mathcal{Y}$ and score function $q : \mathcal{Y} \times \mathcal{U}^n \to \mathbb{R}$ with sensitivity Δ. The theorems in the notes show that we expect the error $q_{\max} - q(A(\mathbf{x}))$ to be $O(\Delta \ln(d)/\varepsilon)$, but it might be much better.

   Specifically, fix a data set $\mathbf{x}$. Suppose the "true winner" for $\mathbf{x}$, the outcome y^* with score $q_{\max}$, is substantially better than all other outcomes, namely, for every $y \neq y^*$,

   $$q(y) < q_{\max} - \frac{2\Delta(\ln(d) + t)}{\varepsilon}$$

   Show that the algorithm will output y^* with probability at least $1 - e^{-t}$.

3. Suppose that, after running the exponential mechanism with privacy parameter ε, we use the Laplace mechanism to estimate the error $q_{\max} - q(A(\mathbf{x}))$ with noise $\text{Lap}(2\Delta/\varepsilon)$. What is the total privacy cost of the combined algorithm if we analyze it using Basic Composition from Lecture 5?

4. The exponential mechanism is often used in private machine learning. Suppose our data set consists of pairs (x_i, z_i) where x_i is an image (e.g., represented as a grid of pixels), and z_i is a label (perhaps indicating whether the picture is of Beyoncé (labeled "1") or Harry Styles (labeled "-1", even though we're not judgy that way).

   We are given a collection of k possible classifiers $f_1, ..., f_k$ (perhaps representing different settings of weights in a neural network) among which we want to choose and output one with low *training*

error. We define the training error of a classifier f as the fraction of misclassified points in the input data set:

$$\text{error}(f) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}\{z_i = f(x_i)\}\,.$$

How can you cast this as a selection problem? If you use the exponential mechanism or report noisy max, what bounds do the lecture notes imply for the (expected, asymptotic) difference between the error of the classifier that is output and the error of the best classifier? Express your answer as a function of k, n and ε.

5. Suppose you have a graph with a fixed vertex set V, and where each individual data point x_i is an undirected edge $\{u, v\} \in V \times V$. For example, the nodes might represent locations, and an edge $\{(u, v\}$ might represent the locations between wich an individual travels most often.

   Consider the problem of finding a near-minimum cut in the graph. This is a partition of V into two disjoint sets A, B of nodes. The *weight* of the cut is the number of edges that cross from A to B (so $u \in A$ and $v \in B$ or vice versa). The weight of a cut can be as large as the size of the data set n, and n can be as large as $\Omega(|V|^2)$.

   (a) Use the exponential algorithm (or report noisy max) to design an algorithm that returns a cut with expected weight $\mathsf{min\text{-}weight} + O(|V|/\varepsilon)$ It's OK if your algorithm runs in time polynomial in $2^{|V|}$.

   (b) **(**)** There can be multiple distinct minimum cuts in a graph. However, one neat (and highly non-trivial to prove) fact is that if w^* is the number of edges in the minimum cut, the number of distinct cuts with weight $\leq cw^*$ is at most $O(|V|^{2c})$. Using this fact, prove that the error of the exponential mechanism (or RNM) is actually much better, and it outputs a cut with expected weight $\mathsf{min\text{-}weight} + O(\log(|V|)/\varepsilon)$

6. **(*)** Prove that Report Noisy Max with exponential noise (Alg. 2 in the notes) is differentially private.

7. Show that the accuracy guarantees we showed for the exponential mechanism (and RNM) are basically tight in general. Specifically,

   (a) give an input to the approval voting problem with d candidates, on which $q_{\max} = n = \frac{\ln(d)}{2\varepsilon}$ but the algorithm A_{EM} will return a candidate who received 0 votes with constant probability (independent of d).

   (b) **(**)** Consider the family of data sets $\left\{\mathbf{x}^{(1)}, ..., \mathbf{x}^{(d)}\right\}$ defined as follows: in $\mathbf{x}^{(j)}$, one candidate j receives $q_{\max} = n = \frac{\ln(d)}{2\varepsilon}$ votes and all others receive 0 votes. Show that for **every** ε-differentially private A algorithm, if we choose J uniformly at random in $[d]$, then with constant probabililty $A(\mathbf{x}^{(J)})$ will return a candidate other than J. [That is, A will fail to find the winner for many datasets of this form.]