CS7880: Rigorous Approaches to Data Privacy, Spring 2017
POTW #6

Instructor: Jonathan Ullman

Due Fri, Mar 17, 11:59pm
(Email to jullman+PrivacyS17@gmail.com)

- **You may work on this homework in pairs if you like. If you do, you must write your own solution and state who you worked with.**
- Solutions must be typed in LaTeX.
- Aim for clarity and brevity over low-level details.

Problem 1 (Differential Privacy Prevents Reconstruction/Reidentification)**.**

We have seen several examples of attacks on privacy, and claimed informally that differential privacy prevents these attacks. In this problem we will make these informal claims rigorous.

(a) Suppose that our dataset $x = (x_1, \ldots, x_n) \in X^n$ is chosen uniformly at random. Show that if $A : X^n \to R$ is (ε, δ)-differentially private, then for every attacker $B : R \to X$

$$\mathbb{P}\left[B(A(x)) \in \{x_1, \ldots, x_n\}\right] \leq n \cdot \left(\frac{e^{\varepsilon}}{|X|} + \delta\right),$$

where the probability is over the random choice of x and the random coins of A and B. Thus, if $|X| \gg ne^{\varepsilon}$ and $\delta \ll n$, no attacker can "identify" any row of a random dataset from the output of a differentially private algorithm

(b) Suppose that our dataset $x = (x_1, \ldots, x_n) \in \{0,1\}^n$ is chosen uniformly at random. Show that if $A : \{0,1\}^n \to \mathbb{R}$ is (ε, δ)-differentially private, then for every attacker $B : R \to \{0,1\}^n$

$$\mathbb{E}\left[\frac{1}{n} Ham(x, B(A(x)))\right] \geq e^{-\varepsilon} \cdot \left(\frac{1}{2} - \delta\right)$$

where the probability is over the random choice of x and the random coins of A and B. Thus, if there is an algorithm $B(A(x))$ such that $\frac{1}{n} Ham(x, B(A(x))) \leq \rho$, and $e^{-\varepsilon}(1/2 - \delta) > \rho$, then A is not (ε, δ)-differentially private. (In class we saw reconstruction attacks where ρ was $\frac{1}{10}$, although we could have easily made ρ an arbitrarily small constant or even had $\rho \to 0$ under slightly stronger accuracy assumptions.)