Privacy Enhancing Technologies FS2026
Exercise Sheet 8 (version 1.1)

Florian Tramèr

Problem 1: Conceptual Questions. For each of the following statements, say whether it is TRUE or FALSE. Write at most one sentence to justify your answer.

(a) Which of the following mechanisms are 1-DP (i.e., $\varepsilon = 1$) for an input database $D = \{x_1, \ldots, x_n\} \in \{0,1\}^n$ of size $n \geq 1$?

i. Release $\texttt{Enc}(D)$, where $\texttt{Enc}$ is a one-time pad encryption scheme (i.e., $\texttt{Enc}(D)$ samples a mask r uniformly random from $\{0,1\}^n$ and returns $D \oplus r$).

ii. Release $\texttt{Enc}(D)$, where $\texttt{Enc}$ is an encryption scheme (with a fixed key) that is semantically secure for polynomial-time adversaries (e.g., AES).

iii. Let A be a 1-DP algorithm for computing the mean of a database. Let $z \leftarrow A(D)$. Release a denoised version $z' = (z + z_{\text{denoise}})/2$, where $z_{\text{denoise}} = \frac{1}{n}\sum_{i=1}^{n} x_i$.

iv. Release $f(D)$, where f is a function that is $\frac{1}{2}$-DP.

v. To release the sum of elements in D, first sample independent $\texttt{Laplace}(\mu = 0, b = 1)$ noise Y_1 and Y_2, compute $v_0 = \sum_{i=1}^{\lfloor n/2 \rfloor} x_i + Y_1$ and $v_1 = \sum_{i=\lfloor n/2 \rfloor + 1}^{n} x_i + Y_2$, and output $v_0 + v_1$.

vi. Release $\frac{1}{n}\sum_{i=1}^{n} x_i + \texttt{U}[-1,1]$ where $\texttt{U}[-1,1]$ is the uniform distribution over $[-1,1]$.

Problem 2: Additive Differential Privacy. Let $M(D)$ be an algorithm that is $(0,\delta)$-DP for a database $D \in \mathcal{X}^n$, where $\mathcal{X}$ is some finite set. Show that:

(a) If $\delta \geq \frac{1}{2n}$, then M is not private. Specifically, give an example of a mechanism that is $(0,\delta)$-DP but leaks one individual's information $x_i \in \mathcal{X}$ in the database (for some arbitrary index i) with constant (non-zero) probability.

(b) If $\delta < \frac{1}{2n}$, then M is not useful. Specifically, show that there exists a database $D \in \{0,1\}^n$, such that $M(D)$ incurs error $E = \Omega(1)$ (with constant, non-zero probability) when computing the mean of elements of D.

Problem 3: On the tightness of Dinur-Nissim (part 1) Consider the Randomized Response mechanism from the lecture with $\gamma = 1/4$ applied to a database $D = (x_1, \ldots, x_n) \in \{0,1\}^n$. That is, we release a fixed database $D' = (x'_1, \ldots, x'_n) \in \{0,1\}^n$, where $x'_i = x_i$ with probability $3/4$ and $1 - x_i$ with probability $1/4$. Since the adversary gets access to the whole database of noisy responses D', they can ask as many (noisy) counting queries as they want.

(a) Show that, with high probability $(1 - o(1))$ over the randomness in the Randomized Response mechanism, there is a predicate $\phi : [n] \to \{0,1\}$ (which selects a subset of the

database) for which the adversary incurs error

$$e(\phi) = \left| \frac{1}{n} \sum_{i=1}^{n} \phi(i) x_i - \frac{1}{n} \sum_{i=1}^{n} \phi(i) x'_i \right| \in \Omega(1)$$

Hint: Use Hoeffding's inequality.

(b) Show that the adversary can recover (in expectation) 3/4 of the entries of D.

(c) Is this consistent with Theorem 1 of Dinur and Nissim?

Problem 4: Private k-Means. Consider the following algorithm for clustering a database of n vectors into k clusters, known as Lloyd's algorithm (or more simply, k-means):

Algorithm 1: Lloyd's algorithm with random initialization

Input: Data set $D \in \mathcal{X}^n$, where $\mathcal{X} = \{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$, parameter k
$c_1^{(0)}, \ldots, c_k^{(0)} \leftarrow$ random points in $\mathcal{X}$;
for $t = 1$ **to** T **do**
  for $j = 1$ **to** k **do**
    $S_j \leftarrow \{i : c_j^{(t-1)}$ is the closest center to $x_i\}$;
    $c_j^{(t)} \leftarrow \frac{1}{|S_j|} \sum_{i \in S_j} x_i$;
return $c_1^{(T)}, \ldots, c_k^{(T)}$;

The algorithm starts with random centers and iteratively updates them to minimize the sum of squared distances to the points in the cluster.

We'll now aim to make this algorithm differentially private.

(a) Suppose you are given the cluster centers $c_1^{(t-1)}, \ldots, c_k^{(t-1)}$ of the previous iteration $t-1$, so that you can compute S_j at step t.

First, we'll look at computing all the values $|S_j|$ for $j = 1, \ldots, k$ while satisfying ε_1-DP. Show how to do this with the Laplace mechanism, so that the maximal noise across all values is on the order of $\tilde{O}(\frac{k}{\varepsilon_1})$ with high probability (i.e., going to 0 as n goes to infinity).

(b) It turns out that this is highly suboptimal! Show how to compute all values $|S_j|$ for $j = 1, \ldots, k$ with the Laplace mechanism, while satisfying ε_1-DP and with maximal noise across all values on the order of $\tilde{O}(\frac{1}{\varepsilon_1})$ with high probability.

Hint: View all the $|S_j|$ as a single vector.

(c) In the same setting as above, show how to release all the values $\sum_{i \in S_j} x_i$ for $j = 1, \ldots, k$ with ε_2-DP using the Laplace mechanism (with as little noise as possible).

(d) What is the privacy level ε of the entire algorithm, as a function of ε_1 and ε_2?