CS395T: Continuous Algorithms
Homework VI

Kevin Tian

Due date: April 25, 2024, start of class (3:30 PM).

Please list all collaborators on the first page of your solutions. Unless we have discussed and I have specified otherwise, homework is not accepted if it is not turned in by hand at the start of class, or turned in electronically on Canvas by then. Send me an email to discuss any exceptions.

1 Problem 1

Let $D_\alpha(P\|Q) := \frac{1}{\alpha-1}\log(\int_\Omega P(\omega)^\alpha Q(\omega)^{1-\alpha}\mathrm{d}\omega)$ denote the Rényi divergence of order $\alpha \geq 1$ between two probability distributions P, Q supported on the same sample space Ω.

(i) Prove that if $\alpha \leq \beta$, then $D_\alpha(P\|Q) \leq D_\beta(P\|Q)$.

(ii) Prove that $\lim_{\alpha\to 1^+} D_\alpha(P\|Q) = D_{\mathrm{KL}}(P\|Q)$,[1] assuming $D_\alpha(P\|Q) < \infty$ for some $\alpha > 1$.

2 Problem 2

Let $\alpha, \epsilon \in (0, \frac{1}{10})$. For a database $\mathbf{D} \in \{-1,1\}^{n\times d}$, let $\mu(\mathbf{D}) \in [-1,1]^d$ denote its one-way marginals, i.e. $[\mu(\mathbf{D})]_j = \frac{1}{n}\sum_{i\in[n]} \mathbf{D}_{ij}$ is the average of the entries in the j^{th} column, for all $j \in [d]$. Let $\mathcal{M} : \{-1,1\}^{n\times d} \to [-1,1]^d$ be an ϵ-DP mechanism such that

$$\Pr\left[\|\mathcal{M}(\mathbf{D}) - \mu(\mathbf{D})\|_\infty \leq \alpha\right] \geq \frac{1}{2} \text{ for all } \mathbf{D} \in \{-1,1\}^{n\times d}.$$

In particular, we view databases as elements in $(\{-1,1\}^d)^n$, so $\mathbf{D}$ and $\mathbf{D}'$ are neighboring if they differ in one row. Prove that there is a 1-DP mechanism $\mathcal{M}' : \{-1,1\}^{m\times d} \to [-1,1]^d$ such that

$$\Pr\left[\|\mathcal{M}'(\mathbf{D}) - \mu(\mathbf{D})\|_\infty \leq \frac{1}{2}\right] \geq \frac{1}{2}, \text{ for all } \mathbf{D} \in \{-1,1\}^{m\times d},$$

for a value of $m = \Theta(\alpha\epsilon n)$, assuming $\alpha\epsilon n$ is sufficiently large.[2]

3 Problem 3

(i) Let $\mathcal{M} : \{-1,1\}^{n\times d} \to [-1,1]^d$ be a 1-DP mechanism. Let $\mathbf{D} \in \{-1,1\}^{n\times d}$ be a database, and suppose there are N databases $\{\mathbf{D}_k\}_{k\in[N]}$ such that each $\mathbf{D}_k$ is at Hamming distance at most Δ from $\mathbf{D}$.[3] Suppose that there are disjoint subsets $\{\mathcal{X}_k\}_{k\in[N]} \subseteq [-1,1]^d$ such that for all $k \in [N]$, $\Pr[\mathcal{M}(\mathbf{D}_k) \in \mathcal{X}_k] \geq \frac{1}{2}$. Prove that $N \leq 2\exp(\Delta)$.

(ii) Following the notation of Problem 2, let $d \geq 10$, and suppose $\mathcal{M} : \{-1,1\}^{n\times d} \to [-1,1]^d$ is a 1-DP mechanism, satisfying

$$\Pr\left[\|\mathcal{M}(\mathbf{D}) - \mu(\mathbf{D})\|_\infty < 1\right] \geq \frac{1}{2}, \text{ for all } \mathbf{D} \in \{-1,1\}^{n\times d}.$$

Prove that $n \geq \frac{d}{3}$.

[1] We use $\lim_{\cdot\to\cdot^+}$ to denote a one-sided limit from the right.

[2] The strategy for establishing the conclusion of this problem is particularly helpful for proving sample complexity lower bounds in differential privacy, as it reduces to the case of lower bounding the sample complexity of constant-accurate, constant-differentially private mechanisms. For instance, it boosts the guarantee in Problem 3.

[3] We say two databases have Hamming distance Δ if they differ in exactly Δ rows.

4 Problem 4

Please fill out this form: `https://forms.gle/SFXGe83HoZTMTe5w9`.

5 Problem 5

Please fill out this form: `https://forms.gle/mmrpUpYyzsKjgdtK6`. (This problem is not graded.)