EE378C: Homework #2

Due on Wednesday, May 5, 2021

Please hand in your homework via Gradescope before 11:59 PM. Typing your solution using LaTeX is highly recommended.

1. This problem concerns the local asymptotic minimax theorem applied to the entropy estimation example covered in Lecture 7. Recall the following setup: the learner draws n iid samples $X_1, \cdots, X_n$ from a discrete distribution $P = (p_1, \cdots, p_k)$, and aims to estimate the entropy $H(P) = \sum_{i=1}^{k} -p_i \log p_i$. With a slight abuse of notation, we also use P to denote the free parameter $(p_1, \cdots, p_{k-1})$, which belongs to the parameter set $\mathcal{P}_k = \{(p_1, \cdots, p_{k-1}) \in \mathbb{R}_+^{k-1} : \sum_{i=1}^{k-1} p_i \le 1\}$ with a non-empty interior in $\mathbb{R}^{k-1}$.

   (a) For a fixed P in the interior of $\mathcal{P}_k$, find the expression of the Fisher information $I(P)$ and the inverse Fisher information $I(P)^{-1}$ in the above model with $n = 1$.
   Hint: the following Woodbury matrix identity might be useful: for invertible A, C,

   $$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1}.$$

   (b) Use the local asymptotic minimax theorem to show that for any P_0 in the interior of $\mathcal{P}_k$ and any sequence of estimators $\widehat{H}_n$ based on n samples, it holds that

   $$\lim_{C\to\infty} \liminf_{n\to\infty} n \cdot \sup_{P\in\mathcal{P}_k : \|P-P_0\|_2 \le C/\sqrt{n}} \mathbb{E}_P[(\widehat{H}_n - H(P))^2] \ge \mathsf{Var}_{X\sim P_0}(\log P_0(X)),$$

   where for $P = (p_1, \cdots, p_k)$, the variance is defined as

   $$\mathsf{Var}_{X\sim P}(\log P(X)) \triangleq \sum_{i=1}^{k} p_i \log^2 p_i - \left(\sum_{i=1}^{k} p_i \log p_i\right)^2.$$

   (c) Find a suitable P_0 in (b) to conclude that

   $$\liminf_{n\to\infty} n \cdot \inf_{\widehat{H}_n} \sup_{P\in\mathcal{P}_k} \mathbb{E}_P[(\widehat{H}_n - H(P))^2] \ge c \cdot \log^2 k,$$

   where $c > 0$ is an absolute constant independent of (n, k).

2. In this problem we analyze the reduction scheme from Lecture 6 that lets us approximately map

   - Bern(1/2) $\to \mathcal{N}(0, 1)$
   - Bern(1) $\to \mathcal{N}(\mu, 1)$

   Formally, we observe a bit $B \in \{0, 1\}$ and use it to create a distribution that approximates a Gaussian random variable. Consider the following algorithm $RK(B)$:

   - If $B = 1$, sample $Z \sim \mathcal{N}(\mu, 1)$. Output Z with probability 1.

- If $B = 0$, set $T = 0$ and $Z = 0$.
  While $T \le N$, sample $Y_T \sim \mathcal{N}(0,1)$.
  With probability $\max\left\{0, 1 - \frac{\mathcal{N}(\mu,1)(Y_T)}{2\mathcal{N}(0,1)(Y_T)}\right\}$, set $Z = Y_T$ and break.
  Otherwise, increment T by 1.
  Output Z.

Here $\mathcal{N}(a,1)(x)$ represents the pdf of the distribution $\mathcal{N}(a,1)$ at x. Denote the distribution of the output of this algorithm when $B \sim \mathrm{Bern}(x)$ as $RK(\mathrm{Bern}(x))$.

(a) Compute $\|RK(1) - \mathcal{N}(\mu,1)\|_{\mathrm{TV}}$.

(b) Define $S = \{x \in \mathbb{R} : 2\mathcal{N}(0,1)(x) \ge \mathcal{N}(\mu,1)(x)\}$ and a distribution φ supported on S specified by the pdf

$$\varphi(x) = \frac{\mathcal{N}(0,1)(x) - \frac{1}{2}\mathcal{N}(\mu,1)(x)}{p} \cdot \mathbb{1}(x \in S),$$

where $p := \mathbb{P}_{X\sim\mathcal{N}(0,1)}[X \in S] - \frac{1}{2}\mathbb{P}_{X\sim\mathcal{N}(\mu,1)}[X \in S]$ is the normalizing constant. Show that $\|RK(0) - \varphi\|_{\mathrm{TV}} = (1-p)^N$.

Hint: Show that if $P_{X|A}$ denotes the conditional distribution of X given $X \in A$, then $\|P_X - P_{X|A}\|_{\mathrm{TV}} = P_X(A^c)$.

(c) Show that

$$\|RK(\mathrm{Bern}(1/2)) - \mathcal{N}(0,1)\|_{\mathrm{TV}} \le \frac{1}{2}(1-p)^N + p - \frac{1}{2}.$$

Note: as $\mu \to 0$ and $N \to \infty$, one can show that $p \to 1/2$, which means that we have achieved the desired approximate reduction.

3. In class we see how the two-point method is used to establish the lower bound for the *expected* loss $\mathbb{E}_\theta[L(\theta,T)]$. In some scenarios we are also interested in the *high probability* upper bound of the following form: $L(\theta,T) < \varepsilon$ with probability at least $1-\delta$. In this problem we show how to adapt the two-point method to proving lower bounds for the high probability result, i.e. show that $L(\theta,T) \ge \varepsilon$ with probability at least δ under $X \sim P_\theta$ for some $\theta \in \Theta$. In particular, we are interested in the risk dependence on the error probability δ.

(a) Find a loss function $L_0(\theta,a)$ (which may depend on L and ε), such that

$$\sup_{\theta\in\Theta} \mathbb{E}_\theta[L_0(\theta, T(X))] \le \delta$$

if and only if the estimator T satisfies $L(\theta,T) < \varepsilon$ with probability at least $1-\delta$ for every $\theta \in \Theta$.

(b) Consider a Bernoulli model $X_1, \cdots, X_n \sim \mathrm{Bern}(p)$ with unknown $p \in [0,1]$ and loss $L(p,a) = |p-a|$. By applying the two-point method to the loss function L_0, argue that there exists an estimator T with $L(\theta, T) < \varepsilon$ with probability at least $1-\delta$ for every $\theta \in \Theta$, where $\varepsilon, \delta \in (0, 1/4)$, *only if*

$$n \ge c \cdot \frac{\log(1/\delta)}{\varepsilon^2},$$

for some absolute constant $c > 0$ independent of (ε, δ). In other words, given n samples, any estimator suffers a loss at least $\Omega(\sqrt{\log(1/\delta)/n})$ with probability at least δ for the worst-case $p \in [0,1]$.

Hint: recall the following relationship between TV *and* KL*:*

$$\|P - Q\|_{\mathrm{TV}} \le 1 - \frac{1}{2}\exp(-D_{\mathrm{KL}}(P\|Q)).$$

(c) Now consider the uniformity testing problem covered in Lecture 8. Show that if the test error is required to be at most $\delta \in (0, 1/4)$ under both $H_0 : P = \mathsf{Unif}([k])$ and $H_1 : \|P - \mathsf{Unif}([k])\|_{\mathrm{TV}} \ge \varepsilon$, the number of samples required is at least

$$n = \Omega\left(\frac{1}{\varepsilon^2}\sqrt{k \log\left(\frac{1}{\delta}\right)}\right).$$

Note: the dependence on δ in (b) and (c), albeit different, is both tight.

4. Consider a multi-armed bandit problem with K arms and two possible scenarios: the reward of arm $i \in [K]$ follows the distribution μ_i in the first scenario, and the distribution ν_i in the second scenario; the rewards across different times are independent. Now consider a generic policy $\pi = (\pi_1, \cdots, \pi_T)$, where for each time t, the action $\pi_t \in [K]$ depends causally on the historic observations $(\pi_1, r_{1,\pi_1}, \pi_2, r_{2,\pi_2}, \cdots, \pi_{t-1}, r_{t,\pi_{t-1}})$, where $r_{t,i} \sim \mu_i$ or ν_i denotes the random reward of arm i at time t. Let $P^T_{\mu,\pi}$ be the probability distribution of all observations under policy π and the first scenario, and $P^T_{\nu,\pi}$ is defined similarly under the second scenario. Moreover, for any $i \in [K]$, let $N_i = \sum_{t=1}^T \mathbb{1}(\pi_t = i)$ be the number of times that arm i is pulled. Show that

$$D_{\mathrm{KL}}(P^T_{\mu,\pi} \| P^T_{\nu,\pi}) = \sum_{i=1}^K \mathbb{E}_{P^T_{\mu,\pi}}[N_i] \cdot D_{\mathrm{KL}}(\mu_i \| \nu_i).$$

5. Consider the following Gaussian sequence model $X_i = \theta_i + Z_i$ for $i \in [p]$, where the parameter vector $\theta = (\theta_1, \cdots, \theta_p)$ could take any value in $\mathbb{R}^p$, and $Z_1, \cdots, Z_p \sim \mathcal{N}(0,1)$ are iid standard normal noises. Consider the target of estimating $\theta_{\max} = \max_{i \in [p]} \theta_i$, with the loss function $L(\theta, T) = (T - \theta_{\max})^2$.

(a) Show that there exists an absolute constant $c > 0$ independent of p such that

$$\inf_T \sup_{\theta \in \mathbb{R}^p} \mathbb{E}_\theta[(T - \theta_{\max})^2] \ge c \cdot \log p.$$

(b) Propose an estimator T such that

$$\sup_{\theta \in \mathbb{R}^p} \mathbb{E}_\theta[(T - \theta_{\max})^2] \leq C \cdot \log p,$$

for some absolute constant $C < \infty$ independent of p.

Note: a careful analysis could give the tight constant $1/2$*:*

$$\inf_T \sup_{\theta \in \mathbb{R}^p} \mathbb{E}_\theta[(T - \theta_{\max})^2] = \left(\frac{1}{2} + o_p(1)\right) \cdot \log p.$$