STAT215: Homework 2

Due: Wednesday, Feb 14

1. (10 pt) Suppose we take one observation, X, from the discrete distribution,

$$\begin{array}{cccccc} x & -2 & -1 & 0 & 1 & 2 \\ \Pr(X=x|\theta) & (1-\theta)/4 & \theta/12 & 1/2 & (3-\theta)/12 & \theta/4, \quad 0 \le \theta \le 1 \end{array}$$

Find a real-valued statistic $T(X)$ with $\mathsf{E}[T(X)] \equiv \theta$ (T is an unbiased estimator of θ, but might take values outside $[0,1]$). Obtain the maximum likelihood estimator (MLE) $\hat{\theta}(x)$ of θ and show that it is not unique. Is any choice of MLE unbiased?

2. (10 pt) Bickel & Doksum pg 153, problem 2.3.8. Nine points (out of 10) for a correct solution in $p = 1$ dimensions. Hint: In $p = 1$ dimension, $|x-y|^\alpha = \big[(x-y)^2\big]^{\alpha/2}$; in $\mathbb{R}^p$, $|x-y|^\alpha = \big[\sum_{j=1}^p (x_j - y_j)^2\big]^{\alpha/2}$.

3. (10 pt) Let $(X_1, \cdots, X_n)$ be a random sample from the uniform distribution on the interval $(\theta - \frac{1}{2}, \theta + \frac{1}{2})$, where $\theta \in \mathbb{R}$ is unknown. Let $X_{(j)}$ be the j^{th} order statistic.

   (a) Show that $(X_{(1)} + X_{(n)})/2$ is strongly consistent for θ, *i.e.*, that $\lim_{n\to\infty}(X_{(1)} + X_{(n)})/2 = \theta$ a.s.

   (b) Show that $\overline{X}_n := (X_{(1)} + \cdots + X_{(n)})/n$ is L^2 consistent.

4. (10 pt) Let $\{X_i\} \overset{\text{iid}}{\sim} \mathsf{No}(\mu, \sigma^2)$ be a random sample from the normal distribution with unknown mean $\mu \in \mathbb{R}$ and known variance $\sigma^2 > 0$. For fixed $t \neq 0$, find the MLE $T_n(X)$ of $e^{t\mu}$ and find a constant c_n for which $S_n(X) := c_n T_n(X)$ is an unbiased estimator of $e^{t\mu}$. Show that the variance of S_n is larger than the Information Inequality lower bound, but that the ratio of it's variance to the lower bound converges to 1 as $n \to \infty$. Recall that MLE's in different parametrizations $\eta = g(\theta)$ are related simply by $\hat{\eta} = g(\hat{\theta})$ for any smooth function g, like $g(x) = e^{tx}$.

5. (10 pt) Let X have density function $f(x|\theta)$, $x \in \mathbb{R}^d$, and let θ have (proper) prior density $\pi(\theta)$. Let $\delta^\pi(x)$ denote the Bayes estimate of θ under $\pi(\theta)$ for squared-error loss (*i.e.*, $\delta^\pi(x) := \mathsf{E}[\theta \mid X = x]$) and suppose δ^π has finite Bayes risk $r(\pi, \delta^\pi) = \mathsf{E}[|\delta^\pi(X) - \theta|^2]$. Show that for any other estimator $\delta(x)$,

$$r(\pi, \delta) - r(\pi, \delta^\pi) = \int \big(\delta(x) - \delta^\pi(x)\big)^2 f(x)\,dx$$

where $f(x)$ is the marginal density of X. If $f(x|\theta)$ is normal $\mathsf{No}(\theta, 1)$, consider the collection of estimators of the form $\delta(X) := cX + d$. Show that whenever $0 \le c < 1$

these estimators are all proper Bayes estimators, and hence admissible [Hint: Find a prior π for which $\delta^\pi(X) = cX + d$]. Show that if $c > 1$ the resulting estimator is inadmissible.

6. (10 pt) Bickel & Doksum pg 197, problem 3.2.1

7. (10 pt) Bickel & Doksum pg 197, problem 3.2.4