Homework 7

36-705

Due: Thursday October 15 by 3pm.

1. For each of the following examples compute the Bayes estimator and its Bayes risk under squared loss:

   (a) $X \sim \text{Poisson}(\lambda), \lambda \sim \text{Gamma}(\alpha, \beta)$.

   (b) $X \sim N(\theta, \sigma^2)$ where σ^2 is known and $\theta \sim N(a, b^2)$.

2. Suppose that $X_1, \ldots, X_n \sim N(\mu, 1)$. Suppose we want to estimate μ^2. It turns out the best unbiased estimator in this problem is $\overline{X}_n^2 - 1/n$.

   (a) Calculate the Cramér-Rao lower bound on the variance of any unbiased estimator of μ^2. To do this, you need the following general form of the Cramér-Rao lower bound: If $W = W(X_1, \ldots, X_n)$ is a function of the data and $\mathbb{E}_\theta[W]$ is a differentiable function of θ then

   $$\text{Var}_\theta(W) \geq \frac{\left(\frac{d}{d\theta}\mathbb{E}_\theta[W]\right)^2}{I_n(\theta)}.$$

   (b) Show that the estimator above has a larger variance than the lower bound. This shows that in some problems the Cramér-Rao lower bound is not achievable.

3. Let $X \sim \text{Bin}(n, p)$ where $0 < p < 1$ and suppose we want to estimate $\theta = 1/p$.

   (a) Show that there are no unbiased estimators for θ.

   (b) Find the MLE and show that it is consistent.

4. Let $X_1, \ldots, X_n \sim N(\mu, 1)$. Let $\mu \sim N(0, b^2)$ be the prior for μ.

   (a) Find the Bayes estimator.

   (b) Find the risk $R(\mu)$ of the Bayes estimator.

   (c) Find $\sup_\mu R(\mu)/R_n$ where R_n is the risk of the minimax estimator.

   (d) What happens as $n \to \infty$?