Math 466/566 - Homework 5 Solutions

1. Book, chapter 7, problem 4.

Solution: The expected value of the sample mean is always the population mean, so the sample mean is always an unbiased estimator. The variance of a Poisson RV is equal to its mean, θ. So the variance of the sample mean is θ/n. To find the Cramer-Rao bound we must compute $I(\theta)$.

$$f(x|\theta) = \frac{\theta^x e^{-\theta}}{x!}$$

where $x = 0, 1, 2, 3, \cdots$. So

$$\ln(f(x|\theta)) = x \ln(\theta) - \theta - \ln(x!)$$

$$\frac{\partial \ln(f(x|\theta))}{\partial \theta} = \frac{x}{\theta} - 1$$

$$\frac{\partial^2 \ln(f(x|\theta))}{\partial \theta^2} = \frac{-x}{\theta^2}$$

Since the Poisson RV is discrete, $I(\theta)$ is given by a sum, not an integral

$$I(\theta) = \sum_{x=0}^{\infty} \frac{x}{\theta^2} f(x|\theta) = \frac{1}{\theta^2} \sum_{x=0}^{\infty} x\, f(x|\theta)$$

Note that the last sum is just the expected value of X and so is θ. So $I(\theta) = 1/\theta$. So Cramer-Rao says the variance of an unbiased estimator is at least θ/n. So the sample mean has minimal variance.

2. Book, chapter 7, problem 5. **Solution:** Some computation gives

$$\frac{\partial^2 \ln(f(x|\theta))}{\partial \theta^2} = 2\frac{(x-\theta)^2 - s^2}{(s^2 + (x-\theta)^2)^2}$$

So

$$I(\theta) = \int_{-\infty}^{\infty} 2\frac{(x-\theta)^2 - s^2}{(s^2 + (x-\theta)^2)^2}\, f(x|\theta)\, dx = \frac{2s}{\pi} \int_{-\infty}^{\infty} \frac{(x-\theta)^2 - s^2}{(s^2 + (x-\theta)^2)^3}\, dx$$

A nasty integral, but we can simplify it a bit. A simple change of variables $x \to x + \theta$ shows

$$I(\theta) = \frac{2s}{\pi} \int_{-\infty}^{\infty} \frac{x^2 - s^2}{(s^2 + x^2)^3} \, dx$$

Then another change of variables ($x \to sx$) shows

$$I(\theta) = \frac{2}{\pi s^2} \int_{-\infty}^{\infty} \frac{x^2 - 1}{(1 + x^2)^3} \, dx$$

You can do the integral using tables, a software package or even contour integration if you've taken complex variables. I think you get $I(\theta) = 1/(2s^2)$. So Cramer Rao says the variance of any unbiased estimator is at least $2s^2/n$.

3. Consider the exponential distribution $f(x|\theta) = \theta \, e^{-\theta x}$ where $\theta > 0$. As always, we have a random independent sample $X_1, X_2, X_3, \cdots, X_n$. The mean of this distribution is $\mu = 1/\theta$.

(a) Find the maximum likelihood estimators of the mean μ and of θ.

Solution:

$$f(x_1, x_2, \cdots, x_n) = \theta^n \, \exp(-\theta \sum_{i=1}^{n} x_i)$$

So

$$\ln(f(x_1, x_2, \cdots, x_n)) = n \ln(\theta) - \theta \sum_{i=1}^{n} x_i$$

Take derivative with respect to θ and set it to zero to find the maximum:

$$\frac{n}{\hat{\theta}} - \sum_{i=1}^{n} x_i = 0$$

So the MLE for θ is

$$\hat{\theta} = \left(\frac{1}{n} \sum_{i=1}^{n} x_i \right)^{-1} = \frac{1}{\bar{X}_n}$$

Since MLE's satisfy the principle of functional invariance, the MLE of $\mu = 1/\theta$ is

$$\hat{\mu} = \bar{X}_n$$

(b) By appealing to a theorem, show that for large n, the MLE for θ is approximately normal, with mean θ and variance θ^2/n.

Solution: We use theorem 8.5 in the book. It says that $\hat{\theta}$ is approximately normal with mean θ and variance $[nI(\theta)]^{-1}$. To compute $I(\theta)$,

$$I(\theta) = -\int \frac{\partial^2 \ln(f(x|\theta)}{\partial \theta^2} f(x|\theta)\, dx = \int \frac{1}{\theta^2} f(x|\theta)\, dx = \frac{1}{\theta^2}$$

So the variance is approximately θ^2/n.

4. Consider the geometric density $f(x|p) = p(1-p)^x$ where $x = 0, 1, 2, \cdots$. We have a random independent sample $X_1, X_2, X_3, \cdots, X_n$. Find the maximum likelihood estimator of the mean and of p.

Solution:

$$f(x_1, x_2, \cdots, x_n) = p^n\, (1-p)^{\sum_{i=1}^n x_i}$$

$$\ln(f(x_1, x_2, \cdots, x_n)) = n \ln(p) + \sum_{i=1}^n x_i\, \ln(1-p)$$

Take derivative with respect to p and set it to zero to find the maximum:

$$n\frac{1}{\hat{p}} - \sum_{i=1}^n x_i \frac{1}{1-\hat{p}} = 0$$

Solving for $\hat{p}$, we find the MLE for p is

$$\hat{p} = \frac{1}{1 + \bar{X}_n}$$

The mean of the geometric distribution is given by $\mu = (1-p)/p$. So by functional invariance, the MLE for the mean is

$$\hat{\mu} = \frac{1-\hat{p}}{\hat{p}} = \bar{X}_n$$

5. Consider the uniform distribution on $[0, \theta]$. We have a random sample $X_1, X_2, \cdots, X_n$.

(a) Find the maximum likelihood estimator of θ. Hint: don't use derivatives. Just try to maximize the likehood given $X_1, \cdots, X_n$.

Solution: The likelihood function is θ^{-n} when $x_1, x_2, \cdots x_n$ all belong to $[0, \theta]$. Otherwise it is zero. So we can write it as

$$f(x_1, \cdots, x_n|\theta) = \theta^{-n} \, 1(x_i \le \theta, \, i = 1, \cdots, n) = \theta^{-n} \, 1(\max x_i \le \theta)$$

We want to maximize this as a function of θ. This is equivalent to maximizing θ^{-n} subject to the constraint $\max x_i \le \theta$. The max occurs at $\hat{\theta} = \max x_i$. So $\hat{\theta}$ is $X_{(n)}$, the largest order statistic.

(b) Find the MLE of the mean $\mu = \theta/2$.

Solution: By the principal of functional invariance, the MLE of the mean is $\hat{\mu} = X_{(n)}/2$.

(c) (**566 only**) Now suppose that we have the uniform distribution on $[\theta_1, \theta_2]$ with both θ_1 and θ_2 unknown. Find the MLE's of θ_1 and θ_2 and of the mean $\mu = (\theta_1 + \theta_2)/2$.

Solution: Now we must maximize $(\theta_2 - \theta_1)^{-n}$ as a function of θ_1 and θ_2 subject to the constraints $\theta_1 \le \min x_i$ and $\theta_2 \ge \max x_i$. The max occurs at

$$\hat{\theta}_1 = X_{(1)}, \quad \hat{\theta}_2 = X_{(n)}$$

By functional invariance, the MLE of μ is $(X_{(1)} + X_{(n)})/2$.

6. (**566 only**) Book, chapter 7, problem 6.