Worksheet 09 (Solutions)

1. Find the MLE estimator for the estimation of the parameter λ from i.i.d. observations of an exponentialy distributed random variable.

Solution: We have the following log-likelihood:

$$\begin{aligned} l(\lambda; x_1, \ldots, x_n) &= \sum_{i=1}^{n} \log \left[\lambda \cdot e^{-\lambda x_i} \right] \\ &= \sum_{i=1}^{n} \left[\log(\lambda) - \lambda x_i \right] \end{aligned}$$

The derivative with respect to λ is:

$$\frac{\partial}{\partial \lambda} l(\lambda; x_1, \ldots, x_n) = \sum_{i=1}^{n} \left[\frac{1}{\lambda} - x_i \right]$$

Setting this equal to zero (and putting a hat on the parameter), gives:

$$\begin{aligned} \sum_{i=1}^{n} \frac{1}{\hat{\lambda}} &= \sum_{i=1}^{n} x_i \\ \frac{n}{\hat{\lambda}} &= \sum_{i=1}^{n} x_i \\ \frac{n}{\sum_{i=1}^{n} x_i} &= \hat{\lambda} \\ \frac{1}{\frac{1}{n} \cdot \sum_{i=1}^{n} x_i} &= \hat{\lambda} \end{aligned}$$

In other words, the MLE is just one divided by the sample mean. That makes a lot of sense (but, again, not maybe very interesting) given that λ is the inverse of the mean for the exponential distribution.

2. Find the MLE estimator for the estimation of the variance from i.i.d. observations of an exponentialy distributed random variable. Hint: This is easily derived from the previous result. Should not require any new derivatives.

Solution: We know that the variance of an exponentially distributed random variable is λ^{-2}. We already have the MLE for λ, so the MLE of the variance is just this value to the -2 power:

$$\begin{aligned} \text{MLE} &= \left[\frac{1}{\frac{1}{n} \cdot \sum_{i=1}^{n} x_i} \right]^{-2} \\ &= \bar{X}^2. \end{aligned}$$

Notice that this **is** quite different than the typical estimator that we use for estimating the variance of a sample (S_X^2), taking into account the special structure of the exponential distribution.

3. Find the MLE estimator for the estimation of the parameter p from i.i.d. observations of a Bernoulli distributed random variable. Hint: When you set the derivative equal to zero, multiple by $\frac{1}{n}$ to write the equation in terms of just $\bar{X}$ and $\hat{p}$.

Solution: We have the following log-likelihood:

$$\begin{aligned} l(p; x_1, \ldots, x_n) &= \sum_{i=1}^{n} \log \left[p^{x_i} \cdot (1-p)^{1-x_i} \right] \\ &= \sum_{i=1}^{n} \left[x_i \cdot \log(p) + (1-x_i) \cdot \log(1-p) \right] \end{aligned}$$

The derivative with respect to p is:

$$\frac{\partial}{\partial p} l(p; x_1, \ldots, x_n) = \sum_{i=1}^{n} \left[\frac{x_i}{p} + \frac{(-1) \cdot (1-x_i)}{1-p} \right]$$

Setting this equal to zero (and putting a hat on the parameter), gives the following

$$\frac{1}{\hat{p}} \cdot \sum_{i=1}^{n} x_i = \frac{1}{1-\hat{p}} \cdot \sum_{i} (1-x_i)$$

Dividing both side by n as in the hint gives:

$$\frac{1}{\hat{p}} \cdot \bar{x} = \frac{1}{1-\hat{p}} \cdot (1-\bar{x})$$

And then, solving gives:

$$\begin{aligned} \bar{x}(1-\hat{p}) &= \hat{p}(1-\bar{x}) \\ \bar{x} - \bar{x} \cdot \hat{p} &= \hat{p} - \bar{x} \cdot \hat{p} \\ \bar{x} &= \hat{p} \end{aligned}$$

And again, we see that the MLE is just the sample mean.

4. Find the MLE estimator for the estimation of the parameters μ and σ^2 from i.i.d. observations of a normally distributed random variable. Hint: We want to think of σ^2 as a single parameter (not the square of a parameter). I recommend using $v = \sigma^2$ to keep this clear. Also, find $\hat{\mu}$ first. You can find the MLE for the mean without knowing the MLE of the variance.

Solution: This is where things get a bit more interesting. We have the following log-likelihood:

$$\begin{aligned} l(\mu, v; x_1, \ldots, x_n) &= \sum_{i=1}^{n} \log\left[\frac{1}{\sqrt{2\pi v}} \cdot e^{-\frac{1}{2v}[x_i-\mu]^2}\right] \\ &= \sum_{i=1}^{n} (-1/2) \cdot \log(2\pi v) - \frac{1}{2v}[x_i - \mu]^2 \end{aligned}$$

The derivative with respect to μ is:

$$\frac{\partial}{\partial \mu} l(\mu, v; x_1, \ldots, x_n) = \sum_{i=1}^{n} \frac{1}{v}[x_i - \mu]$$

Setting this equal to zero gives:

$$\begin{aligned} 0 &= \sum_{i=1}^{n} [x_i - \hat{\mu}] \\ \hat{\mu} &= \bar{x}. \end{aligned}$$

Which is similar to the other results. The more interesting one is the variance. We see that the derivative is:

$$\begin{aligned} \frac{\partial}{\partial v} l(\mu, v; x_1, \ldots, x_n) &= \sum_{i=1}^{n} \frac{-1/2}{2\pi v} \cdot (2\pi) + \frac{1}{2v^2}[x_i - \mu]^2 \\ &= \sum_{i=1}^{n} \frac{-1}{2v} + \frac{1}{2v^2}[x_i - \mu]^2 \end{aligned}$$

Setting this equal to zero and plugging in the value that we know for $\hat{\mu}$, we get:

$$\begin{aligned} \sum_{i=1}^{n} \frac{1}{2\hat{v}} &= \frac{1}{2\hat{v}^2} \sum_{i=1}^{n} [x_i - \hat{\mu}]^2 \\ \frac{2n\hat{v}^2}{2\hat{v}} &= \sum_{i=1}^{n} [x_i - \hat{\mu}]^2 \\ \hat{v} &= \frac{1}{n} \sum_{i=1}^{n} [x_i - \hat{\mu}]^2 \\ &= \frac{1}{n} \sum_{i=1}^{n} [x_i - \bar{x}]^2 \end{aligned}$$

So, this is very similar, but not quite the same, as the estimator S_X^2 that we have been using so far.

5. What is the bias of the MLE estimator for the variance from a normal distribution with unknown mean and variance? Hint: Use what we know about S_X^2 to make this relatively easy.

Solution: We know that the $\hat{v}$ can be written in terms of S_X^2 as follows:

$$\hat{v}_{MLE} = \frac{n-1}{n} \cdot S_X^2$$

So, the expected value is:

$$\begin{aligned} \mathbb{E}\hat{v}_{MLE} &= \frac{n-1}{n} \cdot \mathbb{E}S_X^2 \\ &= \frac{n-1}{n} \cdot v \end{aligned}$$

And the bias is:

$$\begin{aligned} \mathbb{E}\hat{v}_{MLE} - v &= \frac{n-1}{n} \cdot v - v \\ &= v \cdot \left[\frac{n-1}{n} - 1\right] \\ &= v \cdot \left[\frac{n-1}{n} - \frac{n}{n}\right] \\ &= v \cdot \left[\frac{-1}{n}\right] \\ &= \frac{-v}{n} \end{aligned}$$

So, the bias is not zero, but (as we know will be true of all MLE estimators) will limit to zero in the limit of $n \to \infty$.