STATS 730: Midterm test

Monday 16 Sept 2014

Total marks is 50

Question 1: 11 mks:

1. Give an example of a *non-identifiable* model. Be sure to explain why it is not identifiable.
2. Give an example of an MLE that is approximately normal (in the sense that its distribution becomes arbitrarily close to that of a normal distribution when n increases), yet does not possess a mean.
3. Give an example where the bootstrap percentile method would not be a good approach for calculating a confidence interval. [Hint: In the justification of the percentile method there is a symmetry-of-distribution assumption. We have already seen a simple model in which this assumption does not hold.]

Question 2: 11 mks:

Consider a simple linear regression model with known variance,

$$\mu_i = a + bx_i \,, \qquad i = 1, ..., 10 \tag{1}$$

$$y_i = N(\mu_i, 1) \,. \tag{2}$$

For convenience, the variance of y_i has been assumed known, leaving only a and b to be estimated. The MLE of (a, b) was found to be $(\widehat{a}, \widehat{b}) = (-0.4, 1.0364)$ and the Hessian matrix of the log-likelihood was

$$\mathbf{H}(\widehat{\boldsymbol{\theta}}) = \begin{pmatrix} -10 & -55 \\ -55 & -385 \end{pmatrix}$$

which has inverse

$$\mathbf{H}(\widehat{\boldsymbol{\theta}})^{-1} = \begin{pmatrix} -0.4667 & 0.0667 \\ 0.0667 & -0.0121 \end{pmatrix}.$$

Other summary statistics (which may or may not be relevant) are

$$\overline{x} = 5.5, \quad \overline{y} = 5.7$$

$$\sum(x_i - \overline{x})^2 = 82.5, \quad \sum(y_i - \overline{y})^2 = 102.1$$
$$\sum(y_i - \widehat{a} - \widehat{b}x_i)^2 = 13.491, \quad \sum(y_i - x_i)^2 = 14$$

1. Calculate the Wald test statistic for the null hypothesis that the true regression line is $\mu_i = x_i$.
2. Calculate the LR test statistic for the above hypothesis. [NB. The two test statistics should be the same, up to the level of numerical accuracy, since the log-likelihood is quadratic.]
3. What is the degrees of freedom for these test statistics?

Question 3: 11 mks Let y be observed from a Binomial(n, p) distribution. We know that $\widehat{p} = y/n$ with approximate variance $\widehat{p}(1-\widehat{p})/n$. Suppose that we are interested in making inference about $\zeta = P(Y = 0) = (1-p)^n$.

1. Use the delta method to calculate the approximate variance of $\widehat{\zeta} = (1-\widehat{p})^n$.
2. How would you recommend constructing a Wald-based confidence interval (i.e., one based on the approximate normality of MLEs) for ζ?

Question 4: 11 mks Let $Y_i, i = 1, ..., n$ be iid from a Cauchy distribution with median $\theta \in \mathbb{R}$. That is, each Y_i has density function

$$f(y;\theta) = \frac{1}{\pi(1+(y-\theta)^2)}, \quad y \in \mathbb{R}$$

Assume that, within an R session, the data are in a vector `y`.

a). Write an R program to find $\widehat{\theta}$ using the `optim` function.

b). Write additional code to predict 1000 future values from $f(y;\theta)$, using pseudo-Bayesian prediction.

[Note: You may assume that n is sufficiently large that the likelihood is well behaved.]

Question 5: 6 mks: Subject to regularity conditions, it has been shown that the likelihood is higher at the true parameter value, θ_0, than at any other fixed parameter value with probability that becomes arbitrarily close to 1 as the sample size increases. Using this fact, show that the maximum likelihood estimator $\widehat{\theta}$ is consistent. (You may assume $\theta \in \mathbb{R}$ and that the likelihood is unimodal.)