STAT 830

Problems: Assignment 4

1. Suppose we have data $X_1, \ldots, X_n$ and our model is that the data are iid with a Normal$(\mu, 1)$ distribution. Let $Y_i = 1(X_i > 0)$. Let $\psi = P_\mu(Y_i = 1)$.

   (a) Find the MLE, $\hat\psi$, of ψ and an approximate 95% confidence interval for ψ.

   (b) Let $\tilde\psi$ be the sample mean of the Y_i. Show that $\tilde\psi$ is a consistent estimator of ψ.

   (c) Use the delta method to compute an approximate standard error of $\hat\psi$.

   (d) Compute the ratio of the standard error of the MLE of ψ and the approximate SE of $\tilde\psi$ and take the limit of this ratio as $n \to \infty$. The SE of $\tilde\psi$ comes from fact that it is an average – an estimate of a binomial probability.

   (e) Suppose that with some large sample size n you got a certain SE for the MLE of ψ. How many times larger would n have to be to get the same SE, approximately, for $\tilde\psi$? (One over this number is called the Asymptotic Relative Efficiency of $\tilde\psi$.)

   (f) Suppose the data really have the distribution

   $$F(x) = \exp(-\exp(-x)).$$

   What are the limits of $\tilde\psi$ and the limit of the MLE you found in part a)? (The limits I mean are limits as the sample size n goes to ∞.) How do these compare with $P(X > 0)$ for this distribution?

2. Let $X_1, \ldots, X_n$ be a sample from the Normal$(\mu, 1)$ distribution and let $\theta = e^\mu$. Let $\hat\theta$ be the MLE of θ. In this problem I want you to compare several approximations to the distribution of $\hat\theta$.

   (a) First compute $P_\mu(\hat\theta \le t)$ exactly using the true distribution of the sample mean.

(b) Next generate a single sample of size $n = 100$ from the Normal(5,1) distribution; use `setseed(1943)` to make sure everyone uses the same data set. Generate 1000 non-parametric bootstrap replicates of $\hat{\theta}_n$; in the original version of this question I used the undefined symbol T_n for this (perhaps to emphasize that these are computed from bootstrap samples).

(c) Also compute the MLE of μ and generate 1000 parametric bootstrap replicates. That is, draw 1000 samples of size 100 from the Normal$(\hat{\mu}, 1)$ distribution and compute $\hat{\theta}_n$ for each of these 1000 samples.

(d) Use the delta method to get a standard error for $\hat{\theta}_n$ and then use that formula to get an estimated SE for your initial sample of 100.

(e) Plot 4 cdfs on the same plot. Let c_1 be the 0.001 percentile of the true distribution of $\hat{\theta}_{100}$ in the first part for $\mu = 5$ and let c_2 be the 0.999 percentile. Plot the true cdf for $c_1 \le t \le c_2$. Then add the empirical cdfs for the 1000 sample values from parts b) and c). Then add a normal approximations: the cdf of the normal distribution given by the delta method.

(f) The plot shows 3 approximations to the truth based on a single data set. Which is best? Explain your answer.

3. Suppose $X_1, \ldots, X_n$ are a sample of size n from the density

$$f_{\alpha,\beta}(x) = \frac{1}{\beta\Gamma(\alpha)} \left(\frac{x}{\beta}\right)^{\alpha-1} \exp(-x/\beta)\, 1(x > 0).$$

In the following question the digamma function ψ is defined by $\psi(\alpha) = \frac{d}{d\alpha}\log(\Gamma(\alpha))$ and the trigamma function ψ' is the derivative of the digamma function. From the identity $\Gamma(\alpha+1) = \alpha\Gamma(\alpha)$ you can deduce recurrence relations for the digamma and trigamma functions. You can consult Lecture 6 material where a lot of this is redone. You don't have to redo anything I already told you in the videos or the notes.

(a) If α is known find the mle for β.

(b) When both α and β are unknown what equation must be solved to find $\hat{\alpha}$, the mle of α?

(c) Evaluate the Fisher information matrix.

(d) Here is a sample of size 40

```
3.547 1.228 2.052 1.556 2.487 0.469 2.707 0.395
0.770 0.666 4.242 1.474 1.277 2.519 0.578 2.989
1.900 1.422 3.701 1.278 2.820 0.224 0.482 1.426
2.146 2.975 2.792 0.846 3.190 1.680 0.686 1.634
0.969 4.010 1.792 1.287 0.730 0.849 2.447 2.147
```

Use this data in the following questions. First take $\alpha = 1$ and find the mle of β subject to this restriction.

(e) Now use $\mathrm{E}(X) = \alpha\beta$ and $\mathrm{Var}(X) = \alpha\beta^2$ to get method of moments estimates $\tilde{\alpha}$ and $\tilde{\beta}$ for the parameters. (This was done in class so I just mean get numbers.)

(f) Do two steps of Newton Raphson to get MLEs.

(g) Compute standard errors for the MLEs and compare the difference between the estimates in the previous 2 questions to the SEs.

(h) Do a likelihood ratio test of $H_o : \alpha = 1$.

4. Suppose $X_1, \ldots, X_n$ are Uniform$[0, \theta]$ and let $T_n = \max\{X_1, \ldots, X_n\}$ be the MLE of θ. Let

$$\phi(X_1, \ldots, X_n) = 1(T_n > c)$$

be a test function to test $H_o : \theta = 1$.

(a) Find the power function of ϕ.

(b) Find a choice of c so that the level of this test is $\alpha = 0.05$.

(c) The P-value corresponding to this test is a certain function of T_n. Plot this function of T_n against T_n over the range $0.9 \le T_n \le 1.1$ for $n = 20$.

5. For $n = 1, 2, \ldots$ let $\lambda_n = 1/n$ and suppose that $X_n \sim$Poisson(λ_n).

(a) Prove that X_n converges to 0 in probability.

(b) Prove that even $Y_n = nX_n$ converges to 0 in probability.

(c) Show that Y_n does not converge to 0 in pth mean for either $p = 1$ or $p = 2$.

6. Suppose that for each positive integer n we have

$$P\left(X_n = \frac{1}{n}\right) = 1 - \frac{1}{n^2} = 1 - P(X_n = n^2).$$

(a) Does X_n converge in probability and if so to what limit?

(b) Does X_n converge in quadratic mean and if so to what limit?

(c) BONUS question only: does X_n converge almost surely and if so to what limit?

Due: November 24, 2020.