Homework 6: I Made You a Likelihood Function, But I Ate It

36-350, Statistical Computing

Due at 11:59 pm on Tuesday, 8 October 2013 (note day)

Agenda: Using functions as arguments *and* as return values; maximum likelihood estimation; standard errors of estimates.

This homework follows on from this week's lab. You can use the solution code from the lab, or from any earlier homework, lab, or lecture, with acknowledgment.

1. (25) Using any optimization code you like, maximize the likelihood of the gamma distribution on the cats' hearts. Start the optimization at the estimate you get from the method of moments.
   (a) (10) What command do you use to maximize the log-likelihood? Explain its arguments.
   (b) (5) What is the estimate?
   (c) (5) What is the log-likelihood there? The gradient?
   (d) (5) Do the location and the value of the maximum match what you would expect from the plot from the lab? Explain.
2. (35) We need standard errors for the estimated parameters. As we saw in homework 4, if we believe the model is accurate, we can get standard errors by simulating from the fitted model, and re-estimating on the simulation output.
   (a) (10) Write a function, `make.gamma.loglike`, which takes in a data vector `x` and returns a log-likelihood function. Check that if `x` is `cats$Hwt`, then `make.gamma.loglike` returns a function which matches your old `gamma.loglike` at multiple parameter values.
   (b) (10) Write a function, `gamma.mle`, which takes in a data vector `x`, and returns a shape and a scale parameter, estimated by maximizing the log-likelihood of the gamma distribution. It should use your `make.gamma.loglike` function from the previous part. Check that if `x` is `cats$Hwt`, then `gamma.mle` matches the answer you got in problem 1b.

(c) (10) Modify the code from homework 4 to use your `gamma.mle` function, rather than the method-of-moments estimator. (You can use the code from the solutions to homework 4, with acknowledgments.) In addition to giving the modified code, explain in words what you had to change, and why.

(d) (5) What standard errors do you get from running 10^4 simulations?

3. (20) An alternative to using simulation is to use the jack-knife.

   (a) (10) Calculate jack-knife standard errors for the MLE of the gamma distribution. Your code should be able to work with an arbitrary data vector, not just `cats$Hwt`, and you will want to use functions from problems 2a and 1b. You may write your own function, or borrow code from solutions to previous exercises and/or lecture; if you do borrow, you need to acknowledge the source and explain how the borrowed code works.

   (b) (5) What are the jackknife standard errors for the MLE? (If you do not have two, one for the shape and one for the scale parameters, something is wrong.)

   (c) (5) Do your jackknife standard errors for the MLE match those you got in problem 2d? Should they?

4. (5) Are the maximum-likelihood and the method-of-moments estimates compatible with each other? Explain.

5. (5) What is the likelihood (not the log-likelihood) at the maximum? Should you be worried about using the model because the maximum likelihood is so small? Explain.