

Lab 9: Likelihoods

When we have independent samples x_1, x_2, \dots, x_n from a common probability density $p(x)$, the joint probability density of the whole sample is

$$\prod_{i=1}^n p(x_i).$$

When we are not sure what the right density p is, but we think it belongs to some family (like the Gaussian, the exponential, the gamma, etc.), we write the parameters of the family as θ , and say that the **likelihood function** is

$$L(\theta) = \prod_{i=1}^n p(x_i; \theta).$$

Notice that the likelihood is a **function** of the unknown parameters θ , not the known data $x_{1:n}$. One way to estimate the parameters is to maximize the likelihood,

$$\hat{\theta}_{MLE} = \operatorname{argmax}_{\theta} L(\theta).$$

For several reasons, including numerical stability, we usually work with the log-likelihood instead,

$$l(\theta) = \log L(\theta) = \sum_{i=1}^n \log(p(x_i; \theta))$$

whose maximum is located at the same point as the maximum of L . Maximum likelihood estimation is generally the most statistically efficient way to find the parameters of a probability density, when true density really is in the family we've guessed. In this lab, we begin working with likelihood functions, continuing to use the data on the heart weight of cats from previous labs. (Load it now, please, as in Lab 4.)

1. Fit the gamma distribution to the cats' hearts, using the "method of moments" from Lab 4.
2. Calculate the log-likelihood of the shape and scale parameters you just estimated. The answer, rounded to the nearest integer, should be -326. Hint: `?dgamma`.
3. Write a function, `gamma.loglike`, which takes in a shape and a scale parameter (in that order), and returns the log-likelihood of the cats' hearts' masses. Check that when you run `gamma.loglike` with the estimate from question 1, you get the log-likelihood from question 2.
4. Using `curve`, make two plots of the log-likelihood. In the first, vary the shape parameter over the range 1 to 40, while holding the scale parameter fixed at the value you estimated in problem 1. In the second, vary the scale parameter over the range 0.01 to 1, holding shape fixed at your estimated value. At what values of the parameters do the two curves peak? Are the two maximum equal? Should they be? Hint: remember that `curve` will sweep out over any argument, so long as you call it x .
5. Make a contour plot of the log-likelihood, with the shape parameter on the horizontal axis (range 1 to 40) and the scale parameter on the vertical (range 0.01 to 1). Add a point indicating the location of your moment-based estimate from question 1. Hints: Consider using one of the surface functions from the previous lecture (code online). Also, you will probably want to increase the number of levels on the contour plot above the default of 10.
6. Is your contour plot from problem 5 compatible with your curves from problem 4? How can you tell?
7. Use the plot from the previous question to locate the region where the likelihood seems to be largest. Make a new plot which zooms in on this region by changing the ranges over which the shape and scale vary.
8. Based on your plots, how much higher can the likelihood get than -326? Where, roughly, does it reach this maximum?