Assignment 10

Paul Gribble
Winter, 2019

due: Apr 11, 2019

Maximum Likelihood Estimation

In the notes on MLE (https://www.gribblelab.org/stats2019/notes/MLE.pdf) the final section is an example of applying MLE to the estimation of psychometric functions. Read this carefully, you will be working with it for this assignment. There is also some example code in https://www.gribblelab.org/stats2019/code/MLEcode.R (section 3).

In the example in the notes, MLE is used to estimate two parameters, β_0 and β_1. From those two parameters, one can calculate three summary parameters of the psychometric function: threshold, slope and acuity (see equations 12, 13 and 14).

1. Implement the code from the notes to estimate *acuity*. (2 points)
2. Estimate the empirical sampling distribution of *acuity* using a parametric bootstrap. Plot a histogram of the distribution of *acuity* and report the 95 % confidence interval. (5 points)

Here are some hints:

In a parametric bootstrap, you generate sample data by sampling from the estimated model (the probability distribution plus estimated parameters) a large number of times (e.g. 10,000). You then estimate the statistic of interest (in this case, *acuity*) for each sample. So for example in your parametric bootstrap, to generate one new sample, you would:

- for each X position in the real data, compute the probability p of responding "right" using the logistic function and the estimate of β_0 and β_1 you obtained from the real data.
- simulate a participant's response using that probability p. In R you can generate a simulated response for a single trial using `rbinom(1,1,p)`
- so this for all positions X from the original dataset
- now you have the original positions X and simulated responses R
- once you've done that, use the code from the notes to re-estimate the β parameters, and re-estimate *acuity*
- do this 10,000 times and compute & report the 95 % confidence interval for acuity