

COGSCI 131 – Assignment 9
DUE: April 22th at 5 pm

In this assignment, you will implement the Metropolis algorithm to sample from a posterior distribution on a psychological parameter, based on data on bcourses. Note that this is a Bayesian data analysis in that we have an assumed/fixed cognitive model (the approximate number model from class) and you are inferring one of its parameters, W from behavioral data. Thus, you are measuring a cognitive parameter from data.

Assume that a Weber ratio W has prior, $P(W) = e^{-W}$. A likelihood function has been provided that computes the probability of answering correctly given W and the behavioral stimuli (n_1 , n_2 , and a). To use this function, you will need to load the data and extract these columns to provide them to the likelihood.

Recall that in the Metropolis-Hastings algorithm, you start with an initial W , you propose a change to a new value W' (e.g. by adding a sample from a $\text{Normal}(0,1)$ distribution to W) and you accept this change always if it is better, and otherwise with probability $P(W'|D)/P(W|D)$.

NOTE: Please pay attention to which terms do and do not have logs in this assignment. Your implementation will need to always represent logs of probabilities.

1. [10pts, HELP] In 2-3 sentences, explain what “ W ” measures and how. Explain it in terms that a 15-year-old could understand.
2. [5pts, HELP] In Q4, we will implement functions that compute the logarithm of the prior and the logarithm posterior. Explain in 1-2 sentences why we use logs.
3. [5pts, HELP] Explain in a 1-2 sentences why the form of the acceptance ratio $P(W'|D)/P(W|D)$ is convenient for us – that is, why it saves us from doing an integral.
4. [10pts, SOLO] Write functions to compute something proportional to the log prior and log posterior in this model. Ensure that your prior correctly handles cases when $W < 0$, which should be zero probability. In sentence explain what your function should return when $W < 0$, and why.
5. [25pts, HELP] Implement the Metropolis algorithm, starting from a random W , and plot (a) the posterior score of W over the first 300 samples; (b) the value of W over the first 300 samples, and (c) a histogram of the samples of W over the first 10,000 samples after 1000 samples of “burn in.”
6. [15pts, HELP] Use your sampler to determine the probability that W is in the interval $[0.60, 0.65]$. What is that probability?
7. [15pts, SOLO] Run your sampler on the prior (assuming no data) and plot histograms of the prior samples and the posterior samples (from Q4c) in the same graph. What relationship between these distributions does your plot show, and what does it mean in terms of beliefs we should have about the value of W ?