CS109: Probability for Computer Scientists

Lecture 11 — Inference

January 30

Name to Age

Let B be the year that someone was born and let N be their name. From the social security dataset, I am able to give you the joint distribution $P(B = b, N = n) \approx \frac{\text{count}(b,n)}{\text{len(dataset)}}$.

(a) How would you compute $P(B = 1964 | N = \text{Michael})$ using the joint distribution?

(b) Describe how you would compute the full PMF $P(B = b | N = \text{Michael})$, over all birth years b, using code.

Hidden Chambers

a. Imagine the entire 100 meter path is limestone. In that case, the rate of muons arriving per month on the detection plate is

$$100 \cdot e^{-100/40} = 8.2.$$

Assume each muon arrives independently of any other muon and at a constant rate. What is the probability that in one month you would observe 12 muons?

b. Let X be your belief in the meters of limestone above the detection plate. Your prior belief is that any number of meters from 0 to 100 is equally likely:

$$X \sim \text{Uni}(0, 100).$$

After one month, your detection plate has been hit by 12 muons. What is your updated belief in X? You may leave your answer with integrals or sums.

Stanford Eye Test

Let $A \in \{0.00, 0.01, \ldots, 0.99, 1.00\}$ be the (discretized) ability that someone can see. Our prior belief in A is given to you as a dictionary for each possible ability. In this eye test, we show a user a letter at a certain font size and we observe if they get it right or wrong. Then we update our belief in their ability to see.

We observe that a user gets the first letter we show them incorrect. We define $Y = 0$ as the event that the user gets the first letter we show them incorrect.

a. Write an expression in math for the posterior $P(A = a|Y = 0)$.

b. Is this expression going to result in a number or a dictionary?

c. Describe in code how you would solve this expression. Assume you have access to a function **calc_likelihood**. (The likelihood function is super neat - a bit outside of the scope of today's class but happy to talk about it in OH).

1-D Tracking

Imagine you have a self driving car with one LiDAR sensor. A LiDAR sensor is a way to measure distance to other objects. You don't need to know how LiDAR works, just that it will give you a measure of distance. You are trying to detect how far away an object is from the self driving car.

Let T be the true distance from the car to the object. Our prior belief is that $T \sim N(\mu = 1, \sigma^2 = 3)$.

Our LiDAR sensor can measure distance but it isn't perfectly accurate. It has some noise due to measurement error within the instrument. Let X be the distance given by the LiDAR. We say that X is equal to the true distance plus some noise: $X = t + \text{Noise}$. Let M be the noise and $M \sim N(\mu = 0, \sigma^2 = 1.5)$.

a. What is the likelihood function $f(X = x|T = t)$? This is asking, if you knew a value for t, how could you express X. Specifically, use linearity of expectation and linearity of variance to find the parameters of this distribution.

b. We observe a LiDAR measurement of 4 meters. Write out the equation for the probability density $f(X = 4|T = t)$.

c. We want to update our belief in the true distance to the object $f(T = t|X = 4)$. Write your answer in terms of a constant K that you do not need to solve for.