

## Section 7

---


1. **Binary Tree:** Consider the following function for constructing binary trees:

```
def random_binary_tree(p):  
    """  
    Returns a dictionary representing a random binary tree structure.  
    The dictionary can have two keys, "left" and "right".  
    """  
    if random_bernoulli(p): # returns true with probability p  
        new_node = {}  
        new_node["left"] = random_binary_tree(p)  
        new_node["right"] = random_binary_tree(p)  
        return new_node  
    else:  
        return None
```

The if branch is taken with probability  $p$  (and the else branch with probability  $1 - p$ ). A tree with no nodes is represented by `nullptr`; so a tree node with no left child has `nullptr` for the left field (and the same for the right child).

Let  $X$  be the number of nodes in a tree returned by `randomTree`. You can assume  $0 < p < 0.5$ . What is  $E[X]$ , in terms of  $p$ ?

## 2. Flo. Tracking Menstrual Cycles



Let  $X$  represent the length of a menstrual cycle: the number of days, as a continuous value, between the first moment of one period to the first moment of the next, for a given person.  $X$  is parameterized by  $\alpha$  and  $\beta$  with probability density function:

$$f(X = x) = \beta \cdot (x - \alpha)^{\beta-1} \cdot e^{-(x-\alpha)^\beta}$$

- For a particular person,  $\alpha = 27$  and  $\beta = 2$ . Write a simplified version of the PDF of  $X$ .
- For a particular person,  $\alpha = 27$  and  $\beta = 2$ . Write an expression for the probability that they have their period on day 29. In other words, what is the  $P(29.0 < X < 30.0)$ ?
- For a particular person,  $\alpha = 27$  and  $\beta = 2$ . How many times more likely is their cycle to last **exactly** 28.0 days than exactly 29.0 days? You do not need to give a numeric answer. Simplify your expression.
- A person has recorded their cycle length for 12 cycles stored in a list:  $m = [29.0, 28.5, \dots, 30.1]$  where  $m_i$  is the recorded cycle length for cycle  $i$ . Use MLE to estimate the parameter values  $\alpha$  and  $\beta$ . Assume that cycle lengths are IID.  
You don't need a closed form solution. Derive any necessary partial derivatives and write up to three sentences describing how a program can use the derivatives in order to choose the most likely parameter values.

Note: Flo is a real “AI based” app that helps people track their period lengths. The real world distribution of periods is thought to be a mixture distribution between a normal and a weibull distribution [1]. This problem only has you estimate parameters for a simplified Weibull [2].

### 3. Why Boba Cares About MAP:

You don't understand why there's no boba place within walking distance around campus, so you decide to start one. In order to estimate the amount of ingredients needed and the time you will spend in the business (you still need to study), you want to estimate how many orders you will receive per hour. After taking CS109, you are pretty confident that incoming orders can be considered as independent events and the process can be modeled with a Poisson.

Now the question is - what is the  $\lambda$  parameter of the Poisson? In the first hour of your soft opening, you are visited by 4 curious students, each of whom made an order. You have a prior belief that  $f(\Lambda = \lambda) = K \cdot \lambda \cdot e^{-\frac{\lambda}{2}}$ . What is the MLE estimate? What is inference of  $\lambda$  given the observation? What is the Maximum-A-Posteriori (MAP) estimate of  $\lambda$ ? Through your process try to identify what is a point-estimate, and what is a distribution.

#### 4. Vision Test MLE

You decide that the vision tests given by eye doctors would be more precise if we used an approach inspired by logistic regression. In a vision test a user looks at a letter with a particular font size and either correctly guesses the letter or incorrectly guesses the letter.

You assume that the probability that a particular patient is able to guess a letter correctly is:

$$p = \sigma(\theta + f)$$

Where  $\theta$  is the user's vision score and  $f$  is the font size of the letter. This formula uses the sigmoid function:

$$\sigma(z) = \frac{1}{1 + e^{-z}}$$
$$\frac{\partial \sigma(z)}{\partial z} = \sigma(z)[1 - \sigma(z)]$$

Explain how you could estimate a user's vision score ( $\theta$ ) based on their 20 responses  $(f^{(1)}, y^{(1)}) \dots (f^{(20)}, y^{(20)})$ , where  $y^{(i)}$  is an indicator variable for whether the user correctly identified the  $i$ th letter and  $f^{(i)}$  is the font size of the  $i$ th letter. Solve for any and all partial derivatives required by the approach you describe in your answer.

## 5. Multiclass Bayes

Note: we don't expect folks to get to this problem in section. It is here just for students who would like some extra review!

In this problem we are going to explore how to write Naive Bayes for multiple output classes. We want to predict a single output variable  $Y$  which represents how a user feels about a book. Unlike in your homework, the output variable  $Y$  can take on one of the *four* values in the set  $\{\text{Like, Love, Haha, Sad}\}$ . We will base our predictions off of three binary feature variables  $X_1, X_2$ , and  $X_3$  which are indicators of the user's taste. All values  $X_i \in \{0, 1\}$ .

We have access to a dataset with 10,000 users. Each user in the dataset has a value for  $X_1, X_2, X_3$  and  $Y$ . You can use a special query method **count** that returns the number of users in the dataset with the given *equality* constraints (and only equality constraints). Here are some example usages of **count**:

<b>count</b> ( $X_1 = 1, Y = \text{Haha}$ )	returns the number of users where $X_1 = 1$ and $Y = \text{Haha}$ .
<b>count</b> ( $Y = \text{Love}$ )	returns the number of users where $Y = \text{Love}$ .
<b>count</b> ( $X_1 = 0, X_3 = 0$ )	returns the number of users where $X_1 = 0$ , and $X_3 = 0$ .

You are given a new user with  $X_1 = 1, X_2 = 1, X_3 = 0$ . What is the best prediction for how the user will feel about the book ( $Y$ )? You may leave your answer in terms of an argmax function. You should explain how you would calculate all probabilities used in your expression. Use **Laplace estimation** when calculating probabilities.