


## Section 7: Machine Learning

---

### 1. 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 an expression for the probability that they have their period on day 29. In other words, what is  $P(29.0 < X < 30.0)$ ? You do not need to simplify.
- 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? 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  $i$ . For this problem, assume that  $\beta = 2$ . Use MLE to estimate the parameter value for  $\alpha$ . 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. Here is the PDF again:

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

As a way to test your knowledge, after section you can drop the assumption that  $\beta = 2$  and use MLE to solve for both  $\alpha$  and  $\beta$ .

2. **The Most Important Features** Let's explore saliency, a measure of how important a feature is for classification. We define the saliency of the  $i$ th input feature for a given example  $(\mathbf{x}, y)$  to be the absolute value of the partial derivative of the log likelihood, with respect to that input feature  $|\frac{\partial LL}{\partial x_i}|$ . Below, we show both input images and the corresponding saliency of their features (in this case, pixels) for an image classification model:



Consider a trained logistic regression classifier with weights  $\theta$  that predicts binary class labels,  $y$ . In this question, we allow the values of  $\mathbf{x}$  to be real numbers, which doesn't change the algorithm at all (neither training nor testing).

- a. (4 points) What is the Log Likelihood of a single training example  $(\mathbf{x}, y)$  for this logistic regression classifier?
- b. (8 points) Calculate the saliency of a single feature  $(x_i)$  for one training example  $(\mathbf{x}, y)$ .
- c. (5 points) Show that the ratio of saliency for features  $i$  and  $j$  is the ratio of the absolute value of their weights  $\frac{|\theta_i|}{|\theta_j|}$ .



#### 4. Vision Test MLE (Optional)

**Why is this question helpful?** In this problem you will be practicing all the mechanics and math behind logistic regression. The problem has all the same structural features (including a similar assumption), but it is not itself, logistic regression. Wahoo!

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:

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

Explain how you could estimate a user's vision score ( $\theta$ ) based on their 20 responses,  $(f^{(1)}, y^{(1)})$  to  $(f^{(20)}, y^{(20)})$ , 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 all partial derivatives necessary.