

1. A good way to understand penalised logistic regression is by thinking about a case where we have two very highly correlated feature variables. Consider the task of collecting data from 100 people where we are trying to predict whether someone is over 167cm (about the average height in the U.S.) based on two variables: the measured length of their left leg in centimetres (X) and the measured length of their right leg in centimetres (Y). The two quantities X and Y will be very similar, but due to measurement error and some natural variation, they will not be identical for all people.

Fitting a logistic regression model to data of this form, here is what the results might look like (I completely made this up but believe the numbers are realistic):

$$\text{Prob (over 167cm)} = \sigma(0.05 + 3.20 \times X - 1.44 \times Y)$$

Answer the following questions based on this model.

- a. Look at the signs of the coefficients. What does the model seem to imply about the relationship between the length of someone's right leg (Y) and their height?
- b. In most cases we would expect the values of X and Y for a specific observation to be very similar. Simplify the logistic regression model by assuming that X and Y are identical and writing the Prob (over 167cm) as a function of X alone (i.e., not Y).
- c. Another way to simplify the model would be to force the coefficients for X and Y to be the same (i.e., we would have $b = c$). Assuming the X and Y are so similar that their differences are not important, how could we re-write the logistic regression with coefficients for X and Y that are equal?
- d. Compute the elastic net complexity scores for the original logistic regression model, the model from (b) and the model from (c). Likely all the models perform nearly equally as well. Which model(s) would seem to be preferred by the elastic net?
- e. There is another complexity score used to fit what is called **ridge regression** equal to the sum of the squared values of the coefficients. Compute the ridge penalty for the three models. Which model(s) would seem to be preferred by penalised regression using the ridge penalty?

2. In this question, you will derive a formal form of the logistic elastic net model. To begin, let's consider a prediction task with a single feature variable X and only three training data points. We will assume that the classes are the numbers 0 and 1 and the reference class we are trying to predict is 1. The training data is given and denoted by the following ($L1$, $L2$, and $L3$ are the labels):

$$\begin{array}{ll} X1 = 2 & L1 = 1 \\ X2 = 1 & L2 = 0 \\ X3 = 5 & L1 = 1 \end{array}$$

We will be describing a logistic regression of the form:

$$\text{Prob}(L = 1) = F(a + b \times X)$$

In the questions below, we will use the following notation to denote the probability of the logistic regression predictions:

$$P1 = F(a + b \times X1)$$

$$P2 = F(a + b \times X2)$$

$$P3 = F(a + b \times X3)$$

Based on this data, answer the following questions.

- Logistic regression tries to find parameters (here, a and b) that maximizes the probability of observing the training data. Write down the probability of observing the label $L1$ in terms of the predicted probability $P1$. [Hint: The solution is very simple]
- Repeat the previous question for the probability of observing the label $L2$ in terms of the predicted probability $P2$. [Hint: If you not used to think about probabilities, try think about a concrete example. For instance, if $P2$ (the probability that $L2$ is 1) is equal to 0.8, what is the probability that $X2$ is equal to 0?]
- The probability of independent events occurring at the same time is equal to the product of the individual probabilities. A fair coin toss has a 0.5 probability; the probability that it comes up head three times in a row is $0.5 \times 0.5 \times 0.5 = 0.125$. Write down (in terms of $P1$, $P2$, and $P3$) the probability of observing the actual labels for the three data points given the model.
- A number raised to the power of 1 is equal to itself; any non-zero number raised to the power of 0 is equal to one. Using these rules, write an equation for the probability of observing $L1$ given the probability $P1$ that is true regardless of the value of $L1$. No conditional statements allowed!

- e. Combine the previous questions to write a single quantity for the probability of observing the training data given the model in terms of $P_1, P_2, P_3, L_1, L_2,$ and L_3 that is true regardless of the values of the labels.
- f. The goal of logistic regression is to find parameters that maximize the probability you wrote in the previous question. A trick to solving it is to notice that the probability must be non-negative and therefore we can take the logarithm of the equation. The maximizing value of the logarithm is the same value that maximizes the original probability. Take logarithm of the value you have in (e) below. Simplify by applying the following rules: $\log(a \times b) = \log(a) + \log(b)$ and $\log(a^b) = b \times \log(a)$
- g. Rewrite the solution to (f) using summation notation. This should make it easier to read and much more compact.
- h. Now, write (g) in terms of the values X_i along with the parameters \mathbf{a} and \mathbf{b} .
- i. Since we have gotten this far, finish by extrapolating the previous question and write down the optimisation task that describes logistic regression in terms of features X_{ij} and parameters b_j , where i is an index of the observations and j is an index of the features.

