



1. Above is some data similar to today's slides, except that shape has used to describe the classes rather than the blue/orange. Start by covering up the validation data (right) and then answer the questions below.

- a. Draw freehand an estimate of a good separating line for the points in the training data.
- b. Now, uncover the validation data and create the same line from (a) on the right.
- c. Next, compute a confusion matrix for the validation data in the space below:

d. Calculate the error rate for the training data and the validation data.

2. Six examples of predictive modelling tasks are described. For each, identify (i) the classes, (ii) what an observation is, (iii) two possible numeric features, and (iv) pick one class to be the reference class (the class you are trying to predict) and estimate the signs you would expect to be associated with each of the features.

a. Create a model to predict whether a basketball player makes a free-throw shot.

b. Predict whether it will rain tomorrow.

c. Estimate whether a user's cell phone will still be working one year from now.

d. Predict whether a tweet is associated with positive or negative feelings.

e. Predict whether a music song is classified as "rock" or "classical" music.

f. Create a model to determine if an image was taken indoors or outside.

3. Assume that we have built a logistic regression model to predict whether a text message is spam using two features: the count of the number of numeric digits 0-9 (we will call this X) and the count of the number of words with two or more letters that are written all in capitals (we will call this Y). Our training data yielded the following model:

$$\text{Prob}(\text{spam}) = F(-1.2 + 0.3 \times X + 0.2 \times Y)$$

Answer the following questions based on this model.

- a. If a message contains no words with two or more letters that are written all in capitals, after how many numerical digits is there a predicted probability of at least 0.5 that it is spam?
- b. If a message contains no numerical digits, after how many words with two or more letters that are written all in capitals is there a predicted probability of at least 0.5 that it is spam?
- c. Based on (a) and (b), draw a sketch of the boundary line implied by the logistic regression model.
- d. For each of three following messages, compute the probability that the text message is spam.
 - i. "Congratulations, you have won €12.000. Text us back ASAP to claim!"
 - ii. "Hi Joe. I owe you £30.50 for dinner last night. Do you have Venmo?"
 - iii. "I am taking DSST389 this semester. It is by far the WORST class I have ever taken."

