

Homework exercise 1

Due date: 5 April 2016 in class

1. Flexibility of modeling approaches and its implications

In this problem we will investigate both theoretically and empirically the meaning and results of flexibility in predictive modeling.

- (a) For each of these methods, say which version is more flexible and which less flexible, and give a short explanation:
 - i. Least squares regression with few explanatory variables vs least squares regression with many variables and/or interactions
 - ii. k-Nearest Neighbor (k-NN) with $k=1$ vs k-NN with large k (say, $k=100$)
 - iii. Non-parametric regression with big smoothness penalty (that gives a smooth function) vs small smoothness penalty (that gives a wiggly function)
- (b) Draw a schematic graph with the x-axis being the flexibility of the model, and the y-axis being squared error, and draw a graph for each ones of the relevant quantities, showing how it behaves when the flexibility grows:
 - i. Prediction squared bias
 - ii. Prediction variance
 - iii. Irreducible error
 - iv. Test/Prediction squared error
 - v. Training squared error

The important aspects of the schematic is to show for each one if it is increasing, decreasing, constant or non-monotone, and preserve the order between the graphs (for example, we know that the irreducible error is always smaller than the prediction error).

Note: if drawing a graph seems too difficult, a text explanation of the behavior of these curves is also acceptable. It has to explain all the properties of the graphs clearly.

- (c) For the k-NN case, try it out empirically: The file kNN-HW1.r pointed from the class home page does the following:
 - i. Draws a single test set of size $n_{tr}=100$, given by the matrix X_{te} , and vectors f_{te} and y_{te} .
 - ii. Draws 200 training sets of size $n_{te}=100$ (X_{tr} , f_{tr} and y_{tr}), and for each one applies k-NN regression to the test set with $k \in \{1, 5, 10, 20, 50, 100\}$.
 - iii. Collects all the predictions in the array `test.preds` of test predictions $\hat{f}(x)$, whose first coordinate is the observation number, second is k (the neighbors number) and third is the set index (1-200). For example, the vector of test set predictions for the 10th training set and 5 neighbors is `test.preds[, "5", 10]`
 - iv. Also given is the array `test.errs` (which gives $y - \hat{f}(x)$ for test observations), with similar coordinates.
 - v. At the end it gives some examples how to use `test.preds` and `test.errs` to do some calculations.

Your task is to calculate the four first quantities from the previous item for this example (excluding training error) and plot them on the y-axis against the value of k on the x-axis.

Hint: Some of the quantities may simply be read from the code (such as the irreducible error), others have to be calculated from the matrices `test.preds` and `err.preds` and/or the vectors `f.te` and/or `y.te`.

2. ISLR 3.3 (p120)
3. ISLR 3.4 (p120)