Statistical Learning - Milan, Fall2019

Homework problem 10

Basic Neural Nets

1. Assume we are given a modeling problem with $x \in \mathbb{R}^p$ and $y \in \{0, 1\}$, which are can treat as a regression or classification problem (but prediction is always by comparing predictions to 0.5 and predicting either 0 or 1). For the following popular models, describe a neural network that implements them:

   - Standard linear regression
   - Logistic regression

   Explain in what sense the network implements them. Specifically, do we expect to get the same fitted model from the network as from the regular model when applied to data? Why yes or why not?

2. The code prob10.r reads the South African heart dataset, divides it into training and test sets, and uses Keras to apply a NN with one hidden neuron and logit (=sigmoid) activation. It also applies and tests logistic regression. Use this skeleton to:

   (a) Implement all four models described in the previous part

   (b) Prepare 2*2 confusion tables of predicted vs. actual labels

   (c) Briefly discuss the results compared to your expectations from the previous section

3. Implement a more complex architectures (e.g., a hidden layer with three nodes, and then an output layer, see commented code in the file) and apply it to the data. You may play with some of the parameters if necessary. Discuss its test-set performance.

4. Implement a network with a hidden layer with three nodes and an output layer, where all activations are linear. What form does the final model have? What functions of the original x variables are being fitted?

Resources for this problem:
Keras help
Keras in R