Advanced Computational Econometrics
Chapter 7: Neural networks

We consider again the `Boston` dataset. We want to predict variable `medv` using a neural network with one hidden layer (use package `nnet`). Experiment with different numbers of hidden units, different regularization coefficients, and different sets of predictors. Compare the result of best model (with minimal cross-validation error) to that of random forests and SVR. Repeat the experiments with 10 different training/test splits and test the significance of the results.