

Advanced Computational Econometrics

Chapter 7: Neural networks

1 Boston dataset

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.

2 Time series prediction

We consider the CP time series data set. The learning task is to predict the series at time $t + 1$ given the data at times $t - W + 1, \dots, t$, where W is the window size.

1. Read the data. Apply the logarithmic and differencing transformations. Create a matrix with W columns containing the predictors, and a vector containing the response. Split the data into a training set and a test set (use the last 176 observations for the test).
2. Experiment with different network and window sizes, and with different values of the weight decay hyperparameter. Does the neural network perform better than the naive persistence prediction $\hat{y}_{t+1} = y_t$?