

# Economics 706/420

October 1, 2024

R. Davidson

## Assignment 1

This is exercise 8 from Chapter 7 of Géron's book, reproduced here for ease of reference.

Load the MNIST dataset (introduced in Chapter 3), and split it into a training set, a validation set, and a test set (e.g., use 50,000 instances for training, 10,000 for validation, and 10,000 for testing). Then train various classifiers, such as a random forest classifier, an extra-trees classifier, and an SVM classifier. Next, try to combine them into an ensemble that outperforms each individual classifier on the validation set, using soft or hard voting. Once you have found one, try it on the test set. How much better does it perform compared to the individual classifiers?

I know we didn't cover SVM classifiers, and so feel free to leave it out and substitute something else if you prefer.