Intelligent Systems: Reasoning and Recognition

James L. Crowley and Nachwa Aboubakr

MoSIG M1 Winter Semester 2019-2020

Exercise 5 12 March 2020

Recognizing Digits using Neural Networks

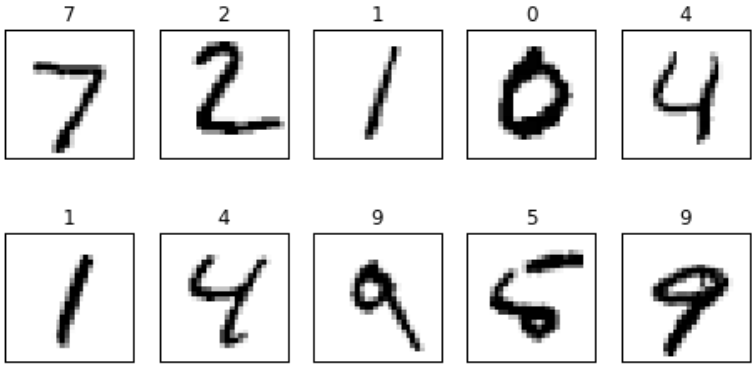

The MNIST (Modified National Institute of Standards and Technology) database is a large collection of handwritten digits. The MNIST database contains 60,000 training images and 10,000 testing images. The database was created by "re-mixing" samples of digits from NIST's original datasets taken from American Census Bureau employees and American high school students. The black and white images from NIST were normalized to fit into a 28x28 pixel bounding box and anti-aliased, which introduced gray-scale levels.

Your task is to design and evaluate neural network architectures that can recognize hand-drawn digits using the gray-scale MNIST images. Your networks can consist of convolution layers, fully connected layers, an auto-encoder, or any combination of these or other neural network techniques. You may use weights that are pre-trained on a different dataset or trained from scratch. Your must provide objective performance evaluation, using metrics such as error rates, ROC curves, accuracy, precision and recall. It is recommended that you use Keras with the tensorflow backend libraries to help you in completing this exercise.

This project should be performed in teams of 3 students, and should be described by a written report with descriptions and performance evaluation results for the techniques that are tested. Reports may be written in French or English. Programming teams are given freedom in their choice of techniques to evaluate. The following is an indicative barometer for grading. Actual grades will depend on a subjective appreciation for the amount of effort deployed and the depth of understanding displayed in the results, and the clarity of the report. Creativity is encouraged and will be rewarded! Reports are due on the 3 April.

Grade	Evaluation Criteria
10-12	Construct and describe a multi-layer network to recognize MNIST digits. Provide results of performance evaluation using ROC and/or precision-recall plots. Provide a clear description of the network architecture, learning and testing procedures. Discuss and explain results.
12-14	Construct, describe and evaluate a multi-layer network architectures to recognize MNIST digits using a variety of training techniques. Test the effects of variations in learning rate, batch size, no. of epochs, kernel size, etc. Provide results of performance evaluation using error rates, ROC curves and precision-recall plots. Explain the learning and testing procedures.
14-18	Construct, describe and comparatively evaluate the effects of different hyper-parameters in the design of a network for recognizing MNIST digits. Build a baseline network to recognize the MNIST digits and then use this network to explore the effects of variations of hyper-parameters such as number of layers, non-linear activation functions, loss functions, types of pooling, spatial-extent (size) and number of filters. Provide results of performance evaluation using error rates, ROC curves and precision-recall plots. Explain the learning and testing procedures. Discuss and explain lessons learned from the results.
18-20	All of the above plus additional unexpected insights or results.