

ARTIFICIAL NEURAL NETWORKS

Due date: 4/16/2020

Other than black voodoo magic, Artificial Neural Networks (ANN) are a very practical way of finding an *unknown* non-linear mapping between some input (i.e. hand-written letters) and the desired output (i.e. digitized words). ANNs are good for two types of problems: classification and non-linear regression. In this exercise we will construct an unbiased backpropagating neural network with 3 layers to help with character recognition.

Assignment:

- a) Create a training set of digits 0-9 by downloading them from the MNIST dataset. Have at least 15-20 examples for each number.
- b) Assemble a 3-layer network. 20 or 30 hidden units might be a good start. Train the network with the created training set. Plot the learning curve and show recognition convergence for the training set.
- c) Test the performance of a neural network on a *distinct* set of digits (preferably hand-written by you and your quarantine inmates). Which numbers are most frequently confused? What is the network's recognition rate?
- d) Fiddle with the network topology. How is recognition affected by the varying number of hidden units and the learning rate parameter? Plot training curves for several examples and comment on optimization.
- e) **Extra credit:** think of another ANN application, this time for non-linear regression, and implement it. Present your results and test their robustness.