

# CS 4320: Machine Learning

## Assignment: Neural Network Classification of Images

In this assignment you will explore the use of convolution neural networks and other neural network architectures to classify images. The data set, [Typeface MNIST](#), is a collection of grayscale images of the digits 0-9 rendered with 2,990 different fonts. The classifier will be judged on its accuracy in predicting the digit from the image.

This is an exploratory assignment. You are required to train *many* different network architectures on this data set, and compare the resulting accuracy of the trained network with the structure of the network.

### Required Steps

- Download the data.
- Download the starter code.
- Unpack the data and split it 80%/20%, for training/testing.
- Further split the training data 70%/30%, for training-fit/validation.
- Repeatedly
  - Configure a network architecture, record its structure and number of parameters.
  - Fit the network using the training-fit data.
  - Measure the network accuracy on the training-fit and validation data; record them.
  - Make notes on your observations of connections between network architecture and accuracy.
- Identify the network architectures with the smallest number of parameters that can obtain validation accuracy of at least: 0.90, 0.95, 0.97, and 0.99. (Numbers as displayed by the `score` function of the starter code).
- *AFTER* all exploration has completed and the 4 different networks have been selected, measure their accuracy on the test data.
- Write a report that includes:
  - A description of the data.
  - An outline of any code changes made (besides network architecture.)
  - A description of each of the 4 different architectures with fewest parameters that meet the accuracy cutoffs.
  - A table with the training-fit, validation, and test accuracy scores for each of the 4 different architectures.
  - A discussion of your observations on the functionality of convolution layers, pooling layers, and dense layers for this particular problem. This will probably include some discussion of hyper-parameter settings for these layer types.
  - Conclusions with your advice to future image classification network designers.
- Commit and push your code in the git repository.
- Submit the report (as PDF) to Canvas.