Name: Patrik Morvai, Bsc student, second year

Project type: thesis project

Topic: Neural Network implementation on FPGAs

Supervisors: Dr. Zoltán Kincses, Dr. Balázs Bánhelyi

The problem of teaching a computer to recognize handwritten digits

- Extremely difficult to solve using the conventional approach to programming
- The idea was to take a large number of handwritten digits, known as training examples and the develop a system which can learn from those training examples
- The neural network uses the examples to automatically infer rules for recognizing handwritten digits
- By increasing the number of training examples, the network can learn more about handwriting, and so improve its accuracy

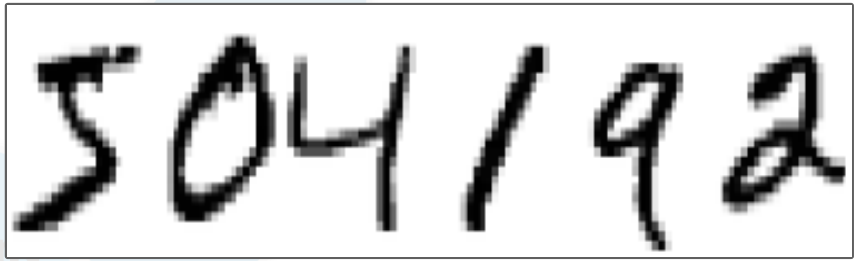

Experiments

A simple network to classify handowrriten digits

- We split the problem of recognizing handwritten digits into two sub-problems:

1. Breaking an image containing many digits into a sequence of separate images, each containing a single digit

- Our approach was to test various image segmentations and use the digit classifier to score them.
- A high score indicates the classifier is confident in all segments, while a low score suggests difficulty in one or more segments, likely due to incorrect segmentation.

2. Classifying each individual digit

- Our training data for the the network consists of many 28 by 28 pixel images of scanned handwritten digits, and so the input layer contains 784 neurons.
- The input pixels are greyscale, with a value of 0.0 representing white, a value of 1.0 representing black, and in between values representing gradually darkening shades of grey.

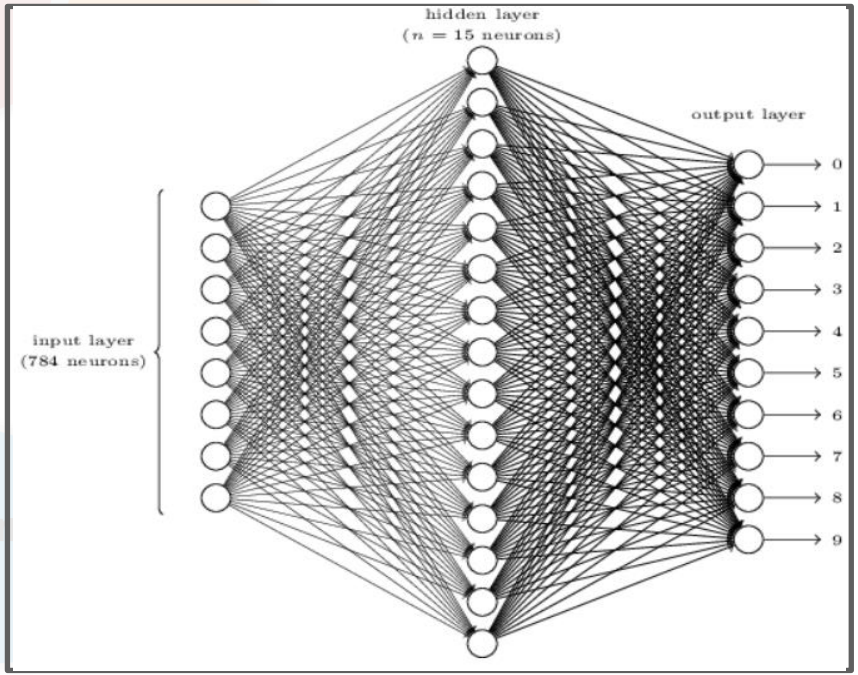

Results & future work

Implementing our network to classify digits

- Using stochastic gradient descent and MNIST training data
- The biases and weights in the Network object are all initialized randomly, using a function to generate Gaussian distributions with mean 0 and standard deviation 1
- Note that the Network initialization code assumes that the first layer of neurons is an input layer, and omits to set any biases for those neurons, since biases are only ever used in computing the outputs from layers

Future work:

- Our goal is to push the accuracy of classifying the digits as high as possible
- Unfortunately, 100% accuracy is not possible because there are training examples that even we humans would struggle to recognize