

## Assignment 4

# Neural Networks

## Handwritten Digit Recognition

**Instructions.** This assignment focuses on training a feed-forward neural network to classify handwritten digits from the classic MNIST dataset (28×28 grayscale images of digits 0–9). You will experiment with different network architectures and training settings and consider the trade-offs that arise when deploying neural networks.

Build models using the [Week 4 web app](#) and answer the questions below. Stick to 3 sentences maximum per question. Upload your completed assignment to BruinLearn as a PDF by the end of the day (11:59pm PT).

**Question 1.** Train a neural network with a single hidden layer of 8 nodes, keeping all other parameters at their defaults (ReLU activation, learning rate 0.001, batch size 32, 50 epochs, early stopping enabled). Report the final test accuracy. Looking at the row for the digit “2” in the confusion matrix, which other digit is the digit “2” most commonly mistaken for by the model? Give a brief explanation as to why.

**Question 2.** Set up a network with the following parameters: Layer 1: 64 nodes, Layer 2: 32 nodes, ReLU activation, batch size 32, 50 epochs, **no** early stopping. Next, run two versions of the model with two different learning rates: one with  $lr=0.0001$  (slider all the way left) and one with  $lr=0.1$  (slider all the way right). Report the validation accuracy for both models. Explain why a higher learning rate might perform worse than a low learning rate (Hint: Consider how it impacts the model weight updates).

**Question 3.** Try different combinations of hyperparameters (architecture, learning rate, batch size, epochs, early stopping on/off) to get the highest validation accuracy you can. Report the settings you used and the validation accuracy you achieved. What was the single biggest lever in improving your accuracy, and why do you think that was? (Many acceptable answers).

**Question 4.** USPS sorts handwritten mail by scanning each envelope, running a neural network to read the address, cross-checking against the list of valid addresses and ZIP codes, and routing low-confidence reads to human reviewers. What tradeoff does USPS face between model accuracy, human labor costs, and the risk of misrouted mail?