Exercise Sheet #7

Problem 1 (*MNIST - Again and again and again...*)

Implement a multilayer perceptron (MLP) with ReLU activation function and cross entropy loss function to classify the handwritten digits in the famous MNIST dataset. You can access the dataset through PyTorch (see here).

- Experiment with different optimizers and hyperparameter values. Try out different hidden layer sizes.
- Generate a plot of the loss over training epochs.
- Evaluate the trained model using the test dataset. What is the prediction accuracy?

Bonus: Repeat the exercise using other network architectures, e.g. consider

- a convolutional neural network (CNN).
- a recurrent neural network (RNN).