

Applied machine learning: Neural Networks

Problem sheet 4

In this exercise, you will build, train, and evaluate a neural network to classify the handwritten digits of MNIST dataset.

Tasks

1. Data Preparation:

- (a) Use the code from the first problem sheet to load the dataset into a 2D numpy array. As before, normalize the pixel values of the images to the range $[0, 1]$.
- (b) Shuffle the data and split it into training, validation, and test sets:

2. Neural network: Build and Train neural network using Tensroflow with the following structure:

- Input layer with $28 \times 28 = 784$ neurons.
- Two hidden layers:
 - First hidden layer: 100 neurons with ReLU activation.
 - Second hidden layer: 50 neurons with ReLU activation.
- Output layer with 10 neurons (one for each digit) and softmax activation.
- Use the loss function `sparse_categorical_crossentropy`.
- Train the model for 10 epochs on the training data.

3. Evaluation:

- (a) Predict the labels for the test set using the trained model.
- (b) Calculate the model's accuracy on the test set.
- (c) Compare the predicted labels with the true labels.
- (d) Randomly select 10 images from the test set.
- (e) Display each image along with its predicted label and the true label.