Université d'Orléans **2025/2026**
Graduate School Orléans Numérique (GSON)

Introduction to Deep Learning
Problem Sheet 1

1) Dataset

Take a look at the MNIST dataset:

1. Unzip the archive (`mnist.zip`). You may use `shutil.unpack_archive`. Inspect a few images.
2. How is the dataset organized? How many images are there in total and per digit? What is their size in pixels?
3. Load one image into a NumPy array using `PIL.Image.open`. How is the image encoded? What are the minimum and maximum pixel values? How are white, black, and gray represented? Visualize the image using `matplotlib.pyplot.imshow`.
4. Write a function that loads the entire dataset into two NumPy arrays: `x` (images) and `y` (labels). Ensure that the grayscale values are floating-point numbers in $[0, 1]$. Shuffle the examples using the same permutation for both `x` and `y`.
5. Split both arrays into training (70%), validation (15%), and test (15%) sets. **Note:** any preprocessing (e.g., normalization/standardization) must be *fit on the training set only* and then applied to validation/test (avoid data leakage).

2) Small neural network

Create a small neural network in TensorFlow using `tensorflow.keras.models.Sequential`:

$$\text{Input} \to \text{Flatten} \to \text{Dense(16, ReLU)} \to \text{Dense(10, Softmax)}.$$

1. Initialize your model (default initializer), *compile* it with `Adam` and `SparseCategoricalCrossentropy`, then *before training* evaluate it on the **validation** set:
   - Sparse categorical cross-entropy.
   - Accuracy.
2. If softmax outputs are uniform for 10 classes, what values do you expect for cross-entropy and accuracy?
3. How many weights and biases does the model have?
4. Train the model for 20 epochs
5. Plot training and validation loss, and report the best validation accuracy and the test accuracy.
6. Is this model sufficient? Train again, replacing the single hidden layer of 16 units with *two* hidden layers of 64 and 32 units. How do the curves and accuracies change?

3) Analyze the behaviour of your network

Model:

$$\text{Flatten} \to \text{Dense(64, ReLU)} \to \text{Dense(32, ReLU)} \to \text{Dense(10, Softmax)}.$$

1. Create a confusion matrix using the test set. Comment on the most frequent confusions.
2. Visualize some of the misclassified images.
3. Create two PNG images manually: one with a digit and one with a cross. (Ensure 28×28 grayscale, scaled to $[0, 1]$.) Classify them with your network and comment.

4) Overfitting

Change the following:

- Data size: train on only 5,000 samples.
- Use a larger model (e.g., Dense(512)–Dense(512)–Softmax).

Train and plot training and validation loss. Identify the epoch where validation loss increases while training loss continues to decrease.

5) Framing: Classification vs. Regression

Frame the problem as a regression problem.

1. Train the network with two hidden layers of 64 and 32 units, but use a *single scalar* output with a *linear* activation and use MSE as the loss function. Report test RMSE and also compute a "rounded accuracy" by rounding predictions to the nearest integer; compare with the classification baseline.
2. Explain why, for ambiguous inputs where $p(Y \mid X = x)$ is multi-modal, the regression predictor may yield values between classes that are not valid labels. When is regression an appropriate modeling choice, and when is classification necessary?