

# MNIST Assignment

In this assignment, you will implement Python code to perform classification on the MNIST dataset using the deep learning framework Keras. There are two parts to the assignment. The *first* is to run the experiment with the instructions below, copy those results to a text file, and submit them to the instructor-specified location/grader. The *second* part of the assignment will be to run your code in the presence of a grader to confirm it runs as expected. In order to receive full credit, please follow the instructions below exactly.

## Part I

1. Navigate to [https://keras.io/examples/mnist\\_cnn/](https://keras.io/examples/mnist_cnn/). Copy and paste the code into a PyCharm project. If you choose to use a Jupyter Notebook or Google CoLab, your submitted .txt document must look the same as if you ran it in PyCharm.

2. Immediately following the line containing `epochs = 12`, add the line

```
np.random.seed(0)
```

Keras gets its source of randomness from numpy's random number generator. To ensure reproducible results, this starts random computation off at the same place every run.

3. Immediately **following** the line containing `model.add(Dense(num_classes, activation='softmax'))` and **before** the line containing `model.compile...`, add the line

```
print(model.summary())
```

This will print a summary of your model's architecture and parameters.

4. Change the line that has `model.fit(...` to

```
history = model.fit(...
```

5. Delete the last three lines of code, i.e. `score = ...`, `print('Test...`, and `print('Test...` and replace them with

```
print([(key, round(value[0], 4)) for key, value in history.history.items()])
```

6. In PyCharm, we will use the “Terminal” to run our code. The **only** reason for this is so that the output is easier for the grader to read. As shown in **Figure 1** below, go to the bottom of the PyCharm window and click “Terminal”.




Figure 1: Open Terminal in PyCharm.

7. In the Terminal window (see **Figure 2**), activate your anaconda environment by typing

`source activate your_env_name` (Mac/Linux)

OR

`activate your_env_name` (Windows)

8. Run your code in the Terminal by typing

`python your_code_name.py`




Figure 2: Activate Environment and Run Code in Terminal.

9. Finally, copy and paste your Terminal window output to a text file, FirstLastName\_MNIST.txt. Their may be some TensorFlow warnings at the start of your code. If your code continues to run correctly after those appear, simply disregard them.

To recap, your output should have, in this order

- (i) Your model summary,
- (ii) 12 progress bars, one for each epoch
- (iii) A list with four tuples, e.g. “[('val\_loss', 1.8283), ('val\_acc', 0.436), ('loss', 1.9965), ('acc', 0.331)]”

## Part II

For this part of the assignment, you will exhibit to the grader that your code is working properly.

1. In the same code you built in **Part I**, change the number of epochs to 1, i.e. **epochs = 1**.
2. Immediately following the line containing **(x\_train, y\_train), (x\_test, y\_test) = mnist.load\_data()**, add the following four lines:

```
x_train = x_train[0:1000]
y_train = y_train[0:1000]
x_test = x_test[0:1000]
y_test = y_test[0:1000]
```

This will reduce the size of your dataset considerably so your code will run much faster.

3. Run the code for the grader to see. You can run this one in *either* PyCharm console or the terminal window. Jupyter and CoLab users will need to re-run all cells.