

CNN Build Instructions

Examination Candidate #Y3898772

8th April 2026

This document describes a suitable method of bootstrapping, building, and testing the *Flowers* convolutional neural network, as submitted by Examination Candidate #Y3898772 in partial fulfilment of the IMLO module assignment. The instructions given here assume a Linux shell on a CUDA-equipped machine, such as `csgpu13.cs.york.ac.uk`.

1. Building the Python virtual environment.

- (a) Using the system Python interpreter, navigate to the `classifier/` project folder and execute `python -m venv .venv`. Activate the virtual environment with `source .venv/bin/activate`.
- (b) Install all third-party library dependencies with `pip install -r requirements.txt`. This will download and install (or recover from the system cache) PyTorch with Torchvision for tensor and neural network support, and SciPy for the parsing of MATLAB files distributed with *Flowers*.

2. Training and evaluating the model.

Still with the virtual environment activated, execute `python ModelDriver.py`; this will train the model for 200 epochs¹, periodically checkpointing the best network parameters to `data/flowers-net.pt`, and then evaluate the trained model. The output of training over the first epoch will appear as follows:

```
3 CUDA GPUs are available; selecting 'cuda:2' (47622258688 bytes free).
Training...
Epoch 0 / 199
  Time since training start: 0:00:29.244856
  Training data accumulated accuracy: 1.47 %
  Validation data accuracy: 4.41 %
  Closing learning rate: 0.001
  Accumulated training loss: 295.82687425613403
  Accumulated validation loss: 294.55209827423096
  Checkpoint? No
```

The model will then self-evaluate, having re-loaded the network parameters of the iteration which incurred the maximal validation accuracy:

```
Testing...
  Loaded model from 'data/flowers-net.pt'
  Testing time: 0:00:29.308905
  Entire network accuracy: 41.05 %
  Finished testing.
```

¹The number of training epochs E can be modified through the sole parameter to `ModelDriver.train_model`, although in the interests of avoiding excessive file-system transactions on early-epoch (continually improving) models, parameters representing maximally accurate networks will only be committed to persistent storage after $\lfloor E/3 \rfloor$ epochs.