

Practical 3

Jumping Rivers

Training Neural Networks

The fashion MNIST data is similar to the MNIST digits but instead of images of digits it is made up of images of 10 fashion items. Otherwise it's format is similar. Each image consists of a 28 by 28 pixel, single colour channel image. For the purposes of our example we will treat each of the 784 pixels as independent inputs. There are 60,000 training images and 10,000 test images. The *jrpytensorflow* package provides a function for downloading the data (the first time) and pre-processing it into an appropriate numpy array structure as below

```
import jrpytensorflow

X_train, y_train, X_test, y_test, labels =
    jrpytensorflow.datasets.load_fashion_mnist()
```

For this example, to begin with, we will consider optimiser performance. Since we want to get a feel for how the different optimisers effect training we will focus on training loss.

- Create a model similar to that we used in the notes for the MNIST digits example. (Hint: think about the shape of the data...)
- Try different optimisers and varying the hyperparameters. How do these affect the model fitting process? Remember to recreate your model for each run as this stores the weights as internal state.
- Using TensorBoard can you visualise the architecture of the model you have built and its training progress?