Exercise: A CNN with TensorFlow (ca. 3h)

1. Learning goal

In the lecture the Convolution Neural Network model was introduced. You should have a rough understanding of its architecture and what happens in the convolutional and the pooling layers and the MLP at the end of the model. Now the goal is to deepen your understanding of CNNs and to learn how they are implemented within a Deep Learning framework as TensorFlow.

2. Download training and test data

For experimenting with CNNs we first need data. The ImageNet database (http://www.image-net.org/) e.g. provides several hundreds to several thousands of samples for different object categories. Use this database or another image database in order to download about 1000 example images of cars and 1000 example images of bikes. Then divide the data into subfolders “train\car”, “train\bike” and “test\car”, “test\bike” and use about 80% of the images you could collect for training and the rest of the images for testing.

3. Find an easy to understand CNN implementation

The internet is full of TensorFlow examples (see e.g. https://github.com/) that show how to build up a CNN model. The difficulty is to find example TensorFlow code that (i) is easy to understand and (ii) allows you to adopt it such that you can feed the model during training and testing with your own images.

4. First steps with image classification

After you have selected a CNN example implementation for image classification that you get running and allows you to feed in your own images, experiment with it it:

- How many training images can you feed into the CNN within 1 minute vs. 1 hour on your PC?
- What is the resulting classification for the model that you get after having it trained for 1 minute vs. 1 hour?
- Slightly adopt the model: use fewer convolutional layers or more convolutional layers and see how the classification error changes.