

Artificial Intelligence

Exercise 7

Marc Toussaint

Machine Learning & Robotics lab, U Stuttgart
Universitätsstraße 38, 70569 Stuttgart, Germany

13. Februar 2019

These is an “extra” exercise for students that need to collect extra points.

1 Keras for image classification

Keras is a really easy to use layer on top of Tensor Flow for neural networks in Python. Overview (details are found below):

- (a) Use the pre-trained VGG16 neural network (that is included in Keras) to train and test on the CIFAR10 data set (that is included in Keras). To this end, use only the fully convolutional parts of the VGG16 (which can be applied to images of any size), fix the weights of this VGG16, add a fully connected output layer that maps to 10 classes, and train only this output layer. There are many projects online that do exactly the same. We give pointers and a starting point below.
- (b) Test the change in performance when you change the output layer’s activation function, or add two output layers.
- (c) Remove again the output layer and use the VGG16 to map all images in the database to a feature vector. This creates a new data base with tuples (feature-vector, label).
- (d) Use python’s scikit-learn to train a logistic regression on this dataset. Compare the results with the original VGG16 results.

Abgabe: Please check in your python code in your group’s git repository. Send me an Email. In this email, tell me your group number and report on two numbers: the classification error you get for (i) and for (iii). Dearline: Please Feb 20 latest.

Details:

- Read <https://keras.io/>, perform the installation using

```
pip install keras --user
```

and test the “30 seconds” example
- Follow <https://keras.io/datasets/> to test loading the cifar10 dataset
- On <https://github.com/mjiansun/cifar10-vgg16> there is an example for training the VGG16 architecture on cifar10. Here are is one version that extracted the essential parts from that:

```
1 # load the pre-trained VGG16 without top layers
2 from keras.applications.vgg16 import VGG16
3 model = VGG16(include_top=False, weights='imagenet', input_tensor=None, input_shape=(32,32,3), pooling=None, classes=10)
4
5 # freeze all these layers
6 for l in model.layers:
7     l.trainable = False
8
9 # add a fully connected top layer
10 from keras.layers import Flatten, Dense
11 from keras.engine import Model
12 x = Flatten()(model.layers[-1].output) #a trivial 'reshape' from image to vector
```

```
13 pred = Dense(10, activation='sigmoid')(x) #a dense layer with 10 outputs
14
15 # report on the architecture
16 full_model = Model(model.input, pred)
17 full_model.summary()
```

- Prepare the data a bit:

```
1 from keras.datasets import cifar10
2 from keras.utils import np_utils
3
4 (X_train, y_train), (X_test, y_test) = cifar10.load_data()
5
6 X_train = X_train.astype('float32')
7 X_test = X_test.astype('float32')
8 X_train /= 255
9 X_test /= 255
10
11 Y_train = np_utils.to_categorical(y_train, 10)
12 Y_test = np_utils.to_categorical(y_test, 10)
13
14
15 N_TRAIN_SAMPLES = 1000
16 N_TEST_SAMPLES = 100
17 X_train = X_train[:N_TRAIN_SAMPLES, :, :, :]
18 Y_train = Y_train[:N_TRAIN_SAMPLES, :]
19 y_train = y_train[:N_TRAIN_SAMPLES]
20
21 X_test = X_test[:N_TEST_SAMPLES, :, :, :]
22 Y_test = Y_test[:N_TEST_SAMPLES, :]
23 y_test = y_test[:N_TEST_SAMPLES, :]
24
25 # check class balance
26 import numpy as np
27 print(np.sum(Y_train, axis=0))
28 print(np.sum(Y_test, axis=0))
```

- And finally train:

```
1 from keras import optimizers
2
3 full_model.compile(loss=categorical_crossentropy',
4                   optimizer=optimizers.SGD(lr=1e-3, momentum=0.9),
5                   metrics=['accuracy'])
6
7 full_model.fit(X_train, Y_train,
8              batch_size=64,
9              epochs=10,
10             validation_split=0.1,
11             shuffle=True)
12
13 scores = full_model.evaluate(X_test, Y_test, verbose=1)
14 print('Test loss:', scores[0])
15 print('Test accuracy:', scores[1])
```

- Continue with (b)-(c) on your own. Information on extracting features is here: <https://keras.io/applications/>, and information on logistic regression of scikit learn is here: https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html.