Machine Learning Workshop 2021
April 5th 2021

Due date: Monday April 5th 2021 23.59h.

Introduction

Deep Neural Networks can be used for Audio Classifications in many different ways. During this workshop we will explore some examples of Convolutional Neural Networks (CNN's) used for image classification and see how they can be adapted and used for audio classification tasks.

Preparations:

1. Download the .gz file containing all data, scripts and notebooks:
   wget <ask for the link if you did not attend the lecture>
   tar -xvf <file_name>.tar.gz
   Do not distribute the data and scripts any further! It is for this workshop only.

2. Make your virtual environment:
   virtualenv ml --python=python3
   source ./ml/bin/activate

   Note: to exit the virtual environment issue the command: deactivate

3. Install some further required packages:
   python3 -m pip install tensorflow
   python3 -m pip install scikit-image
   python3 -m pip install jupyter
   python3 -m pip install matplotlib
   python3 -m pip install keras

4. Start the jupyter notebook:
   jupyter notebook
   Start firefox with page: localhost:8888

Assignment 1:

In the Jupyter Notebook in Firefox browse to the cats_and_dogs-directory and select the Notebook 5.1. This is an example of a neural network trained on the NIST dataset of written characters.
Click the <run>-button for each of the pieces of python script, while reading the explanations that are given. Report the final accuracy.

Do this for each of the other notebook. Note that training can take a very long time on a CPU. In Notebook 5.4 you may have to change some of the paths that were used. Report the accuracies of each of the different versions of the CNN's.

Assignment 2 (Optional):
This assignment is optional as you have to use your own system to run the scripts. Browse to ml_workshop/audio_classification-directory after installing the following components on your system.

```
sudo apt-get install libasound-dev
sudo apt-get install portaudio19-dev python-pyaudio
sudo apt-get install python-pyaudio
sudo apt-get install python3-pyaudio
```

Execute:

```
source ./ml/bin/activate
python3 -m pip install PyAudio
python3 -m pip install scipy=1.0.0
python3 fetch.py
```

Report the performance. What do you notice?

Assignment 3 (Optional)
Study the script from assignment 2. Note that a big part is used to get images from the .wav files. In the images_digits-sub-directory the images of the spectrograms are given.
Make code such that you produce a directory structure similar to the cats and dogs data set. Adapt one of the scripts of the Jupiter Notebook 5.3 for classifying the audio images in the 10 audio classes. Report your code and results.

NB Send your answers in a single .zip file to erwin@liacs.nl before Monday April 5th 2021 23.59h.

References and Links:

[1] The Jupyter Notebooks are from Chapter 5 from the book by Francois Chollet. Deep Learning with Python, Manning Publication, December 2017. (Highly recommended book.)

[2] The Cats and Dogs small data set is extracted from the original data set: https://www.kaggle.com/c/dogs-vs-cats/data , please adhere to the Kaggle-rules, i.e., use the data for this workshop only, and do not distribute any further.

[3] Assignment 2 uses adapted code from Rishi Sidhu, (22-3 2019) and processed data from Zohar Jackson (2019).