Machine Learning Workshop 2022
May 30th 2022

Due date: Monday June 20st 2022 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 .zip file containing all data, scripts and notebooks:
    wget <link will be given on Brightspace>
    unzip <file_name>.zip
    Do not distribute the data and scripts any further! It is for this workshop only.

2. Make your virtual environment:
    virtualenv ml --python=python3.9
    source ./ml/bin/activate

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

3. Install some further required packages:
    pip install --upgrade pip
    pip install tensorflow
    pip install jupyter
    pip install pillow
    pip install numpy scipy matplotlib ipython
    pip install scikit-image

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.

NB This is an exploratory workshop and should not take you more than an hour or two in total! If the code in any of the notebooks fails on your computer, instead of mentioning the results, mention the error (1 sentence per error).

Assignment 2 (Optional, no extra credits):
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
        sudo apt-get install python-pyaudio
        sudo apt-get install python3-pyaudio

Execute Step 2 and 3 as before to make a new virtual environment ml2 and install:
        pip install opencv-python
        pip install --upgrade setuptools
        pip install pyaudio
        pip install soundfile

        python3 fetch.py

Report the performance. What do you notice?
NB If you have to execute fetch.py a second time, be sure to first delete the generated .jpg images.

NB   Submit your answers in a single .zip file using Brightspace before Monday June 20st 2022 23.59h.

References and Links:
[1]   The Jupyter Notebooks are slightly adapted from Chapter 5 from the book by Francois Chollet. Deep Learning with Python, Manning Publication, December 2017. (Highly recommended book. See also 2nd Edition.)
[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 provided 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).