FFT Workshop 2021
March 10th 2021

Due date: Wednesday March 10th 2021, 23.59h.

Introduction

This assignments can be made using the programming language and libraries of your choice (e.g. Python, C, MATLAB®, etc.; librosa, numpy, etc.).

Getting started:

This is just an example using *Python 3.8*, *numpy* and *librosa* that calculates the Fourier transform of an audio file and displays its spectrogram. This and many more examples can be found on https://librosa.org .

1. Make a virtualenv:

```
virtualenv fft --python=python3.8
source ./fft/bin/activate
```

2. Install some packages in this virtual environment:

```
python3.8 -m pip install jupyter
python3.8 -m pip install matplotlib
python3.8 -m pip install librosa
```

3. Start a Jupyter notebook.

```
jupyter notebook
```

4. Enter and run the following code in the notebook to calculate the Fourier transform of an audio file and display a spectrogram:

```
import numpy as np
import matplotlib.pyplot as plt
import librosa
import librosa.display

y, sr = librosa.load(librosa.ex('trumpet'))
D = librosa.stft(y)  # STFT of y
S_db = librosa.amplitude_to_db(np.abs(D), ref=np.max)
plt.figure()
librosa.display.specshow(S_db)
plt.colorbar()
```

Assignment 1 Feature Vectors

In the introduction we showed how we can compute the Fourier transform of a sound signal. Implement a procedure that takes as input a wav file and gives as output for every part of 512 samples the energy of 8 frequency bands (for example: [0Hz,1kHz), [1kHz, 2kHz), … , [7kHz, 8kHz) , please note you may have to use other bands) in the Fourier transformed signal. Calculate these features for the <piano.wav> file which can be found in the <audio_data.zip> file. Use these features to calculate the following code (C_i) for the piano.wav, where

- C_i = **up**, if the current max energy band is of a higher frequency as the previous max energy band
- C_i = **down**, if the current max energy band is of a lower frequency as the previous max energy band
- C_i = **repeat**, otherwise

This codes the signal as a sequence of pitch tendencies, which can be used to recognize melodies. In some respect it resembles the so called Parsons code.

Send your coding-routine and your 'up/down/repeat'-code for the <piano.wav> in a single zip file to erwin@liacs.nl before Wednesday March 10th 2021, 23.59h.