
Machine Learning HW4

ML TAs

- mlta-2023-spring@googlegroups.com
-
-

Outline

- Task Description
- Dataset
- Data Segmentation
- Model Architecture
- Baselines
- Report
- Guidelines

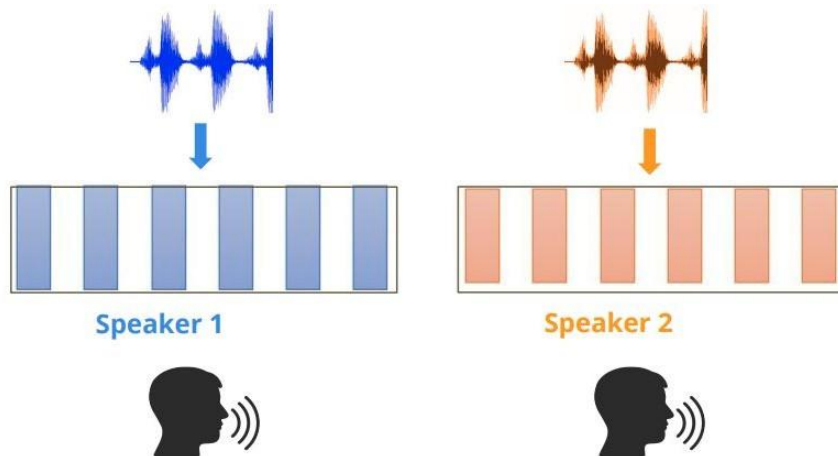
All links

- [Kaggle Competition](#)
- [Sample Code on Colab](#)
- [Sample Code on Kaggle](#)
- [Gradescope](#)
- [Data](#)

Speaker Identification

Task: Multiclass Classification

Predict speaker class from given speech.



Dataset - VoxCeleb2

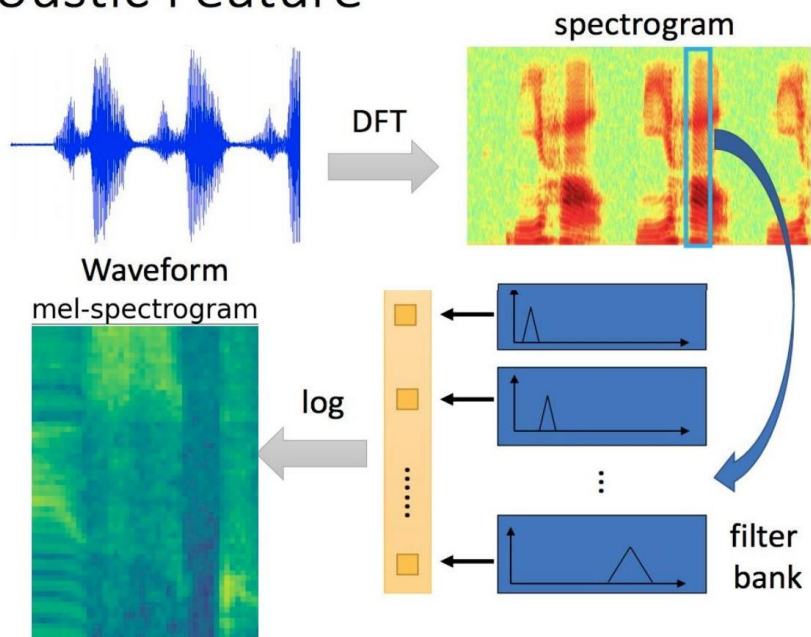
- Training: 56666 processed audio features with labels.
- Testing: 4000 processed audio features (public & private) without labels.
- Label: 600 classes in total, each class represents a speaker.



VoxCeleb2: [Link](#)

Data Preprocessing

Acoustic Feature



Ref. Prof. Hung-Yi Lee
[\[2020Spring DLHLP\] Speech Recognition](#)

Data Format


- Data Directory
 - metadata.json
 - testdata.json
 - mapping.json
 - uttr-{audioID}.pt
- The information in metadata
 - "n_mels": The dimension of mel-spectrogram.
 - "speakers": A dictionary.
 - Key: speaker ids
 - Value: "feature_path" and "mel_len"

```
mapping.json
metadata.json
sox_effects.pt
testdata.json
uttr-0002067f80214182ab863378bdccd68a.pt
uttr-00020e01ed0549928fb3dc8fa50e9cbb.pt
uttr-00031f1879ae44dea75db5aa4ed72b87.pt
uttr-0003464b61764fe98d7992fac6cb66dd.pt
```


```
{
  "n_mels": 40,
  "speakers": {
    "id03074": [
      {
        "feature_path": "uttr-18e375195dc146fd8d14b8a322c29b90.pt",
        "mel_len": 435
      },
      {
        "feature_path": "uttr-da9917d5853049178487c065c9e8b718.pt",
        "mel_len": 490
      }
    ]
  }
}
```

Data Segmentation During Training

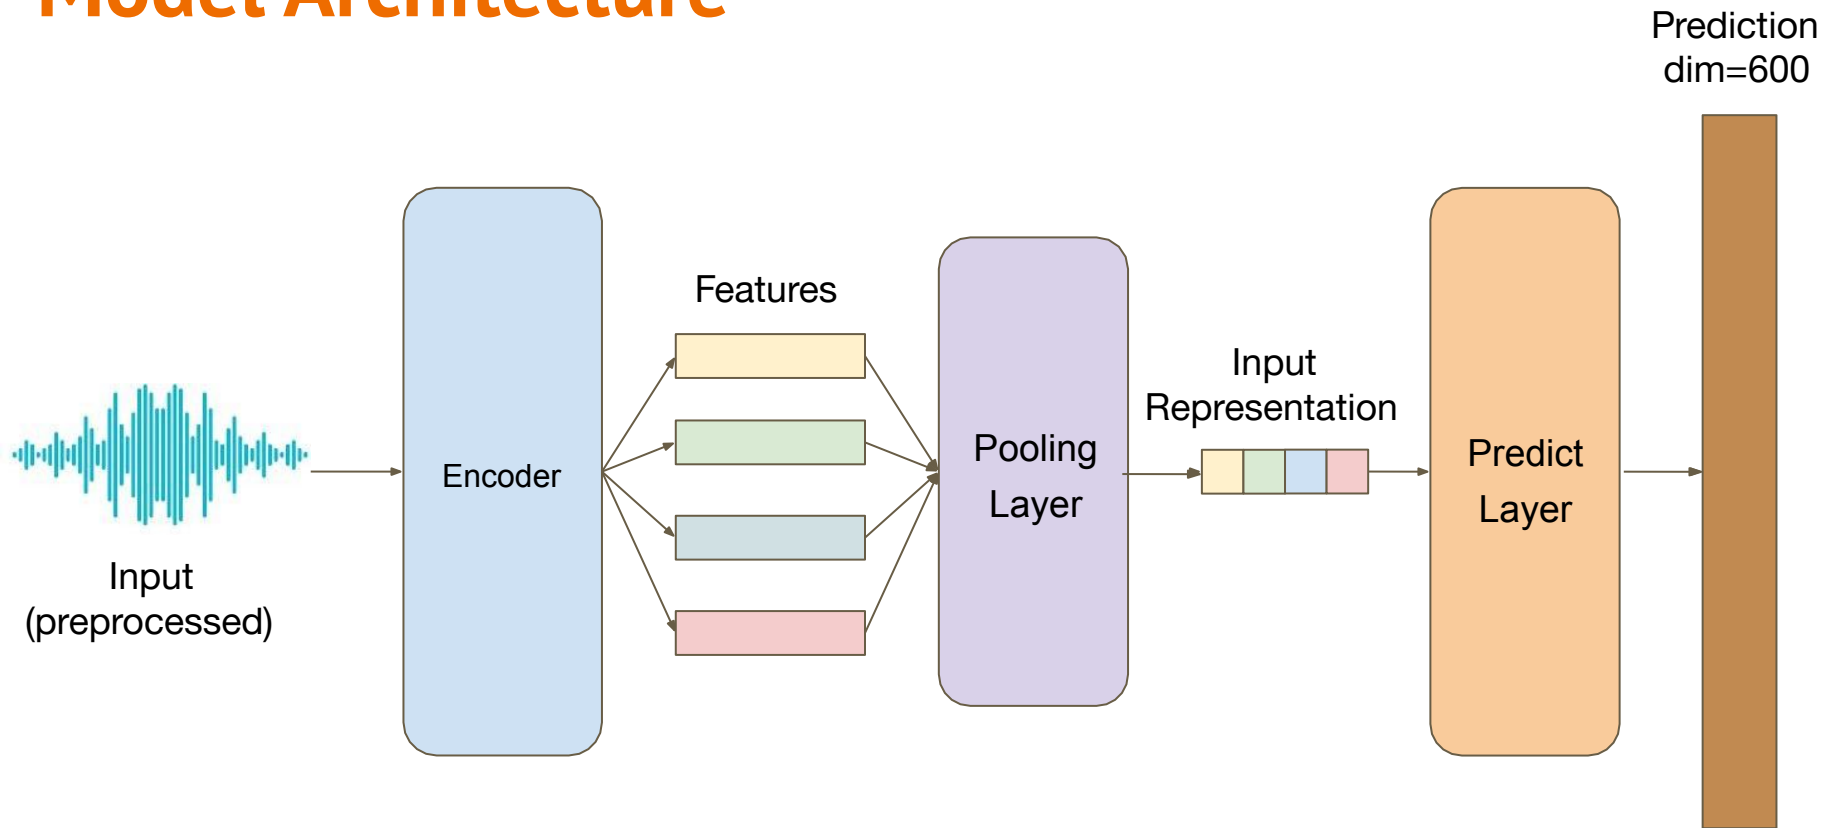
Different length:



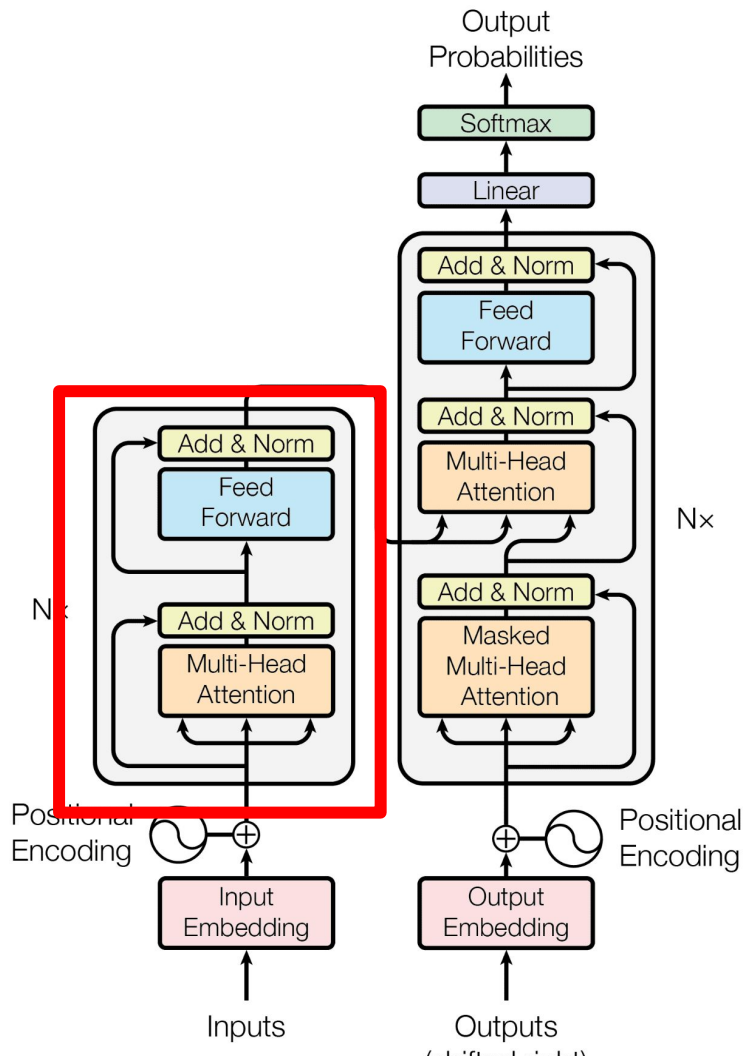
Data Segmentation During Training



Model Architecture



Transformer



Sample Code

- Baseline Methods
 - Simple: Run sample code & know how to use Transformer.
 - Medium: Adjust hyper-parameters of Transformer and prediction layer.
 - Strong: Construct [Conformer](#), which is a variety of Transformer.
 - Boss: Implement [Self-Attention Pooling](#) & [Additive Margin Softmax](#) to further boost the performance.

Requirements - Simple

- Build a self-attention network to classify speakers with sample code.
- Simple public baseline: 0.66025
- Estimate training time: 30~40 mins on Colab.

Requirements - Medium

- Modify the parameters of the transformer modules in the sample code.
- Medium public baseline: 0.81750

```
class Classifier(nn.Module):
    def __init__(self, d_model=80, n_spks=600, dropout=0.1):
        super().__init__()
        # Project the dimension of features from that of input into d_model.
        self.prenet = nn.Linear(40, d_model)
        # TODO:
        #   Change Transformer to Conformer.
        #   https://arxiv.org/abs/2005.08100
        self.encoder_layer = nn.TransformerEncoderLayer(
            d_model=d_model, dim_feedforward=256, nhead=2
        )
        # self.encoder = nn.TransformerEncoder(self.encoder_layer, num_layers=2)

        # Project the the dimension of features from d_model into speaker nums.
        self.pred_layer = nn.Sequential(
            nn.Linear(d_model, d_model),
            nn.Sigmoid(),
            nn.Linear(d_model, n_spks),
        )
```

Estimate training time:
1~1.5 hour on Colab

Requirements – Strong

- Construct Conformer, which is a variety of Transformer.
- Strong public baseline: 0.88500

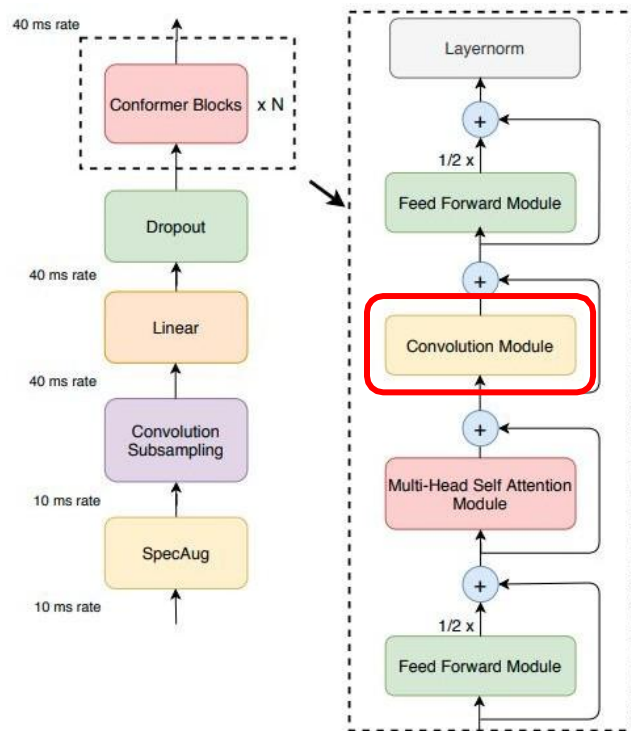
```
class Classifier(nn.Module):
    def __init__(self, d_model=80, n_spks=600, dropout=0.1):
        super().__init__()
        # Project the dimension of features from that of input into d_model.
        self.prenet = nn.Linear(40, d_model)
        # TODO:
        #   Change Transformer to Conformer.
        #   https://arxiv.org/abs/2005.08100
        self.encoder_layer = nn.TransformerEncoderLayer(
            d_model=d_model, dim_feedforward=256, nhead=2
        )
        # self.encoder = nn.TransformerEncoder(self.encoder_layer, num_layers=2)

        # Project the the dimension of features from d_model into speaker nums.
        self.pred_layer = nn.Sequential(
            nn.Linear(d_model, d_model),
            nn.Sigmoid(),
            nn.Linear(d_model, n_spks),
        )
```


Estimate training time:
1~1.5 hour on Colab

Hints

- Conformer

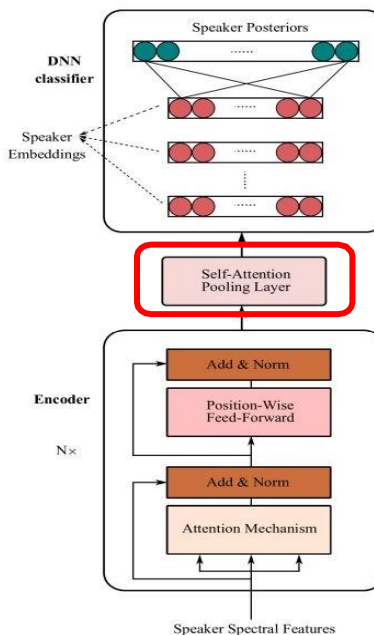


Ref. Prof. Hung-Yi Lee
[\[2021Spring ML\] Network Compression](#)



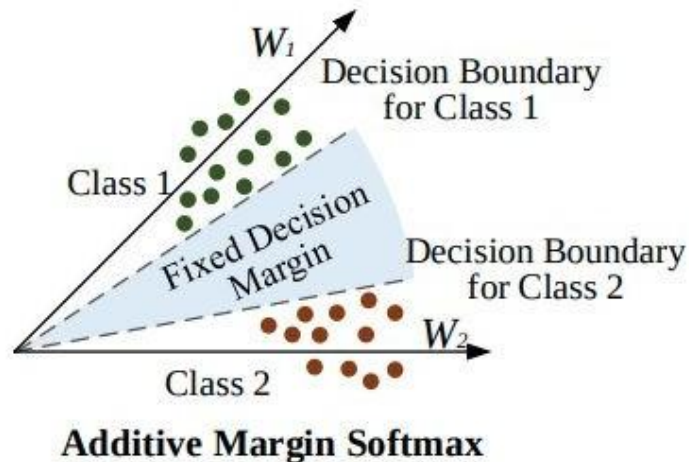

Hints

- Implement Self-Attention Pooling & Additive Margin Softmax to further boost the performance
- Self-Attention Pooling.



Hints

- Additive Margin Softmax



Grading

- Evaluate Metric: @1 Accuracy
- Simple Baseline (Public / Private) +0.5 pt / 0.5 pt
- Medium Baseline (Public / Private) +0.5 pt / 0.5 pt
- Strong Baseline (Public / Private) +0.5 pt / 0.5 pt
- Boss Baseline (Public / Private) +0.5 pt / 0.5 pt
- Code Submission +2 pts
- Gradescope +4 pts

Submission Format

- "Id, Category" split by ',' in the first row.
- Followed by 8000 lines of "filename, speaker name" split by ','.

```
Id,Category  
uttr-434259d99e6a4a34b9a4bf75c9487229.pt,id00000  
uttr-81afc99b0d4c4411ba61ec19c1d6cf40.pt,id00000  
uttr-600d499177b14f669f90ceca4491d817.pt,id00000
```

Code Submission

- Submit your code to **NTU COOL** (2 pts).
 - We can only see your **last submission**.
 - Do **NOT** submit the model or dataset.
 - If your codes are not reasonable, your final grade will be **x 0.9**
 - You should compress your code into a single zip file:
 - ex. b08902126_hw4.zip

<Student ID>_hw4.zip

Gradescope (4 pts)

1. Make a brief introduction about one of [the variant of Transformer](#), and use a image of the structure of the model to help explain (2 pts)

see also: [ref](#)

2. Briefly explain what's the advantages of this variant under certain situations. (2 pts)

Deadline

- Kaggle: **2023/04/07 23:59 (UTC+8)**
- NTU COOL: **2023/04/07 23:59 (UTC+8)**
- Gradescope: **2023/04/07 23:59 (UTC+8)**

Regulation

- You should NOT plagiarize, if you use any other resource, you should cite it in the reference. (*)
- You should NOT modify your prediction files manually.
- Do NOT share codes or prediction files with any living creatures.
- Do NOT use any approaches to submit your results more than 5 times a day.
- **Do NOT search or use additional data or pre-trained models.**
- Your **final grade x 0.9** if you violate any of the above rules.
- Prof. Lee & TAs preserve the rights to change the rules & grades.

(*) [Academic Ethics Guidelines for Researchers by the Ministry of Science and Technology \(MOST\)](#)

If any questions, you can ask us via...

- NTU COOL (Recommended)
 - https://cool.ntu.edu.tw/courses/24108/discussion_topics/184435
- Email
 - mlta-2023-spring@googlegroups.com
 - The title should begin with “[hw4]”
- TA hour
 - Mon 20:00~21:00 (mandarin)
 - Fri 20:00~21:00 (english)

Appendix

- Colab 縮排問題
 - 工具 -> 設定

