

Homework 3 Part 2

UTTERANCE TO PHONEME MAPPING

11-785: Introduction to Deep Learning (Spring 2020)

DUE: November 8th, 2020, 11:59 PM

1 Introduction

In this homework you will again be working with speech data. We are going to be using unaligned labels in this contest, which means the correlation between the features and labels is not given explicitly and your model will have to figure this out by itself. Hence your data will have a list of phonemes for each utterance, but not which frames correspond to which phonemes.

Your main task for this assignment will be to predict the phonemes contained in utterances in the test set. You are not given aligned phonemes in the training data, and you are not asked to produce alignment for the test data.

2 Dataset

Similar to HW1P2, you will be provided with mel-spectrograms that have 13 band frequencies for each time step of the speech data. However in this assignment, the labels will not have a direct mapping to each time step of your feature, instead they are simply the list of phonemes in the utterance [0-40]. There are 41 phoneme labels. The phoneme array will be as long as however many phonemes are in the utterance. We provide a look-up, mapping each phoneme to a single character for the purposes of this competition. However, don't forget about the importance of the blank symbol in CTC Loss. Therefore, your output in the model should be **42 classes**.

The feature data is an array of utterances, whose dimensions are (Utterances, time step, 13), and the labels will be of the dimension (Utterances, frequencies). The second dimension, viz., frequencies will have variable length which has no correlation to the time step dimension in feature data.

2.1 File Structure

In total, we have 7 files for this assignment, including 1 '.csv' file, 1 '.py' file and 5 '.npy' files. Their functions and structure of organization are explained below.

- `train.npy`: This file contains your feature data for training the model. It will be of the shape (Utterances, time step, 13), where the second dimension will be variable as in HW1P2.
- `dev.npy`: This file is similar to `train.npy`, but should be used to calculate your validation losses and accuracy.
- `test.npy`: This file is similar to `train.npy`, but should be used to predict the phoneme labels for the final Kaggle submission.
- `train_labels.npy`: This file contains the labels or phoneme list for each utterance of the `train.npy` file. The dimensions of the data in this file will be of the form (Utterances, frequencies) where the second dimension is variable.
- `dev_labels.npy`: This file is similar to the one above, but instead will map the labels to the `dev.npy` file. You can use this for predicting validation losses and accuracy.
- `sample_submission.csv`: This is an empty submission file that contains the headers in the first row, followed by the test utterance Id and predictions for each utterance of test data.
- `phoneme_list.py`: This file contains the phoneme list and the mapping of each phoneme in the list to their respective sounds. Your submission file should contain these sounds as output and not the phoneme or their corresponding integer.

3 Getting Started

3.1 CTC Loss

As described above, there is no alignment between utterances and their corresponding phonemes. Thus, train your network using CTC loss. Decode your predictions, preferably using beam search. Use the list of phonemes provided on the data page to make each prediction into a text string.

Tensor flow has built-in CTC loss function. Also for Pytorch, you can use `nn.CTCLoss` as mentioned in the recitation.

3.2 CTC Decoding

If you are using PyTorch you can manually install the library, `ctcdecode`, here. They have an implementation of beam search, use it in your code to decode the output of your model. If you are using Tensorflow, it has it's own implementation of beam search.

3.3 Using Greedy Search/Beam Search for CTC Decode

However, you may have issues installing the CTC Decode. You may be able to provide some support on piazza. If the issue remains, then we recommend writing a greedy search implementation or beam search implementation. Additionally, you could always fist submit a test submission using greedy search as it is very trivial to implement.

4 Evaluation & Submission

You will be evaluated using Kaggle's character-level string edit distance. Since we mapped each phoneme to a single character, that means you are being evaluated on phoneme edit-distance.

We are using Levenshtein distance, which counts how many additions, deletions and modifications are required to make one sequence into another.

Your submission should be a CSV file. The headers should be "Id", and "Predicted" - Id refers to the 0-based index of utterance in the test set and Predicted is the phoneme string. Please note that the headers are case sensitive.

See sample submission for details.

5 Conclusion

That's all. As always, feel free to ask on Piazza if you have any questions.

Glhf!