

End-to-End Mandarin Tone Classification with Short Term Context Information

Jiyang Tang* and Ming Li*

* Data Science Research Center, Duke Kunshan University, Kunshan, China

E-mail: {jiyang.tang, ming.li369}@dukekunshan.edu.cn

Abstract—In this paper, we propose an end-to-end Mandarin tone classification method from continuous speech utterances utilizing both the spectrogram and the short-term context information as the input. Both spectrograms and context segment features are used to train the tone classifier. We first divide the spectrogram frames into syllable segments using force alignment results produced by an ASR model. Then we extract the short-term segment features to capture the context information across multiple syllables. Feeding both the spectrogram and the short-term context segment features into an end-to-end model could significantly improve the performance. Experiments are performed on a large-scale open-source Mandarin speech dataset to evaluate the proposed method. Results show that this method improves the classification accuracy from 79.5% to 92.6% on the AISHELL3 database.

I. INTRODUCTION

A. Background

In Computer-Assisted Language Learning (CALL) systems for Mandarin learners, tone classification is one of the most important tasks. In Mandarin, there are four lexical tones, denoted as T1 to T4, plus a neutral tone, denoted as T5. Also, some syllables are considered not to have a tone, and they are denoted as T5. Mandarin tones have important lexical meanings. For example, *shi1* (poem), *shi2* (time), *shi3* (let) and *shi4* (matter) have T1, T2, T3, and T4 as their tones. They have completely different meanings and their pronunciation is distinguished solely by the tones.

In terms of pronunciation, T1 is a flat tone with a relatively high frequency, T2 has a rising frequency, T3 first dips and then rises, and T4's frequency falls from a relatively high value [1], [2], [3], [4], [5]. T5 itself doesn't possess a particular pitch contour pattern. Instead, it is a short and unstressed continuation of the previous syllable [1], [6], [7]. Because the first part of a bi-syllabic Chinese character, called the initial (shengmu), does not contribute much to the tone of a syllable, it is not considered to have a tone attached, denoted as T0. It is also referred to as "onset course" in the tone nucleus model [1].

Continuous tone classification is a difficult task. The main issue is that the pitch contours of a syllable can vary drastically depending on the surrounding syllables. Because the human glottis cannot shift from one frequency to another instantaneously, there has to be a transition between two different tones in a continuous speech utterance. Thus, the pitch contours of a syllable can be affected by the context

bi-directionally [5]. This phenomenon is called tonal coarticulation [3], [5], [6], [8], [9].

Tonal coarticulation makes it very difficult to recognize the tone of a syllable directly using pitch contours. Figure 1 shows two consecutive T4 syllables. The first syllable ends with a relatively high frequency, while the second one starts with a low frequency. The transition between the two syllables makes the pitch contours very blurry around the syllable boundaries (the vertical black line).

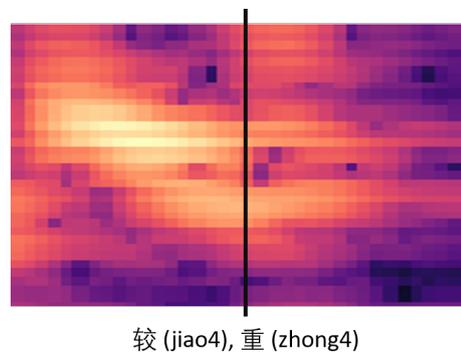


Fig. 1. Example of two consecutive T4 segments, where the second one is very difficult to distinguish

Moreover, "tone sandhi", which is a Mandarin pronunciation rule that changes the pronounced tones in certain situations, makes it more difficult for tone classification [5], [8]. For example, two consecutive T3 are usually pronounced as a T2 followed by a T3 in continuous speech.

In this paper, we propose a continuous Mandarin tone classification framework that utilizes short-term context information to reduce the negative performance impact created by the tonal coarticulation.

B. Related Works

In previous literature, there are two main approaches of tone classification, namely the segment-based approach and the frame-based approach.

In the segment-based approach, speech signal frames are cut into segments representing each syllable in a sentence. Then segment features are fed to the classifier to produce the tone label of the segments. There are many ways of segmenting the frames, such as manual annotation [4], using force alignment [8], [10], [11], using the tone nucleus model

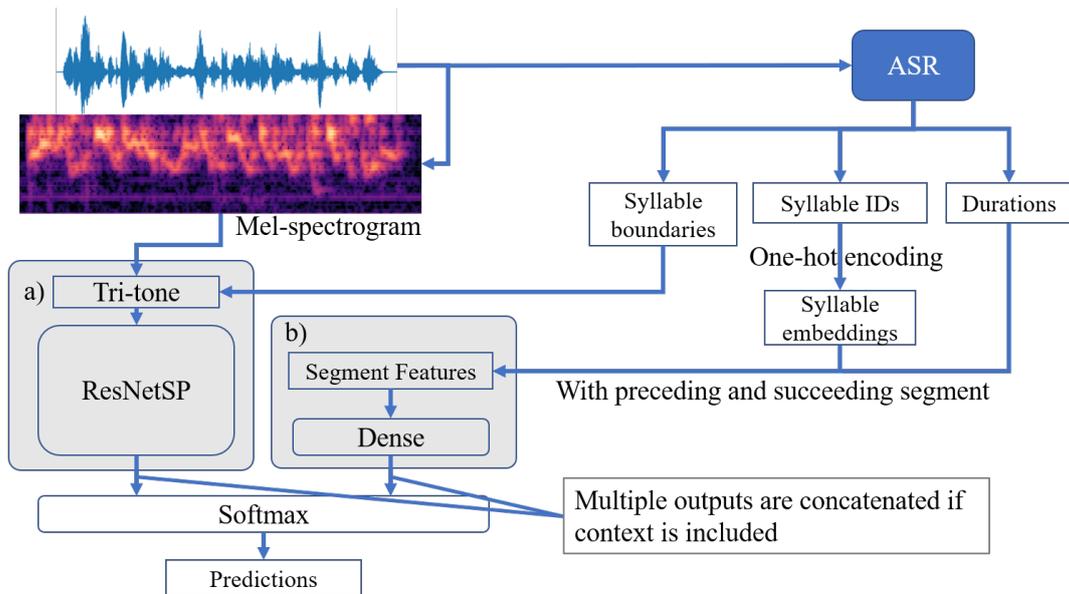


Fig. 2. Overview of the proposed framework

[1], or using unsupervised methods [2]. Also, the context information can be extracted from each segment and used to boost the performance of the model as well [8], [10], [12]. Because the segment-based model relies heavily on force alignment, there might be a performance drop if the force alignment produces inaccurate results. In addition, this approach requires two passes, so the training process is usually more complex than a frame-based approach.

In the frame-based approach, a sequence of frames is directly fed to the classifier, which outputs a sequence of labels. Techniques such as pooling [12], attention [3], or Connectionist Temporal Classification (CTC) [13] are often used to either group the frame-level output into segments or align the frames with phone boundaries so that the model can output a series of tone labels. To capture context information, layers like RNN and CNN can be used [3], [13]. This approach uses a single pass to cover both the frame-segment conversion and the classification task and does not require a pre-trained ASR model. The downside of this approach is that the model is usually much more computational intensive due to the use of RNN, and it is difficult to evaluate the model on a per-tone basis, which is very important for mispronunciation detection.

Besides, if the tone classifier is deployed in a CALL system, which naturally requires force alignment to calculate goodness-of-pronunciation (GOP) [14], a segment-based tone classification method can reuse the existing force alignment results, while the frame-based one cannot.

For feature extraction, there are some popular choices in previous studies. Some researchers used spectrograms of audio signals in their experiments [4], [12], while others used pitch tracking to extract the fundamental frequency contours of a speech utterance [1], [10], [11]. MFCC is also a popular choice [2], [3], [8], [13].

II. PROPOSED FRAMEWORK

Our proposed framework uses a segment-based tone classification approach.

The first part, called the embedding model, is a variant of Residual Neural Network (ResNet) [15] which takes in a sequence of Mel-spectrograms and outputs tone embeddings. The second part, called the contextual model, is a multilayer perceptron model which takes in context segment features extracted from the force alignment results. The outputs of the two parts are concatenated and fed to the final classifier to produce the predictions. The overall architecture can be seen in Figure 2. The detailed configurations of the models will be explained in IV-B.

A. Feature Extraction of the Embedding Model

The Mel-spectrograms are extracted using the same method as ToneNet [4]. More specifically, the waveforms are first resampled to 16000 Hz, and the Mel-spectrograms are then extracted using 64 Mel-filters, [50, 350] Hz frequency range, and 13 ms window length.

Meanwhile, a pre-trained DNN-HMM ASR model is used to perform force alignment in order to obtain syllable boundaries and syllable IDs. The syllable boundaries contain the start time and the duration of the syllables, either being an initial (*shengmu*) or a final (*yunmu*). The syllable IDs are unique identifiers of the syllables in Mandarin pronunciation, such as “*ai*”, “*h*”, “*ue*”, etc. They are configured to be consecutive integers starting from 0, as shown in Table I.

TABLE I
EXAMPLE SYLLABLE IDS

sil	a	ai	an	...
0	1	2	3	...

After that, the Mel-spectrograms are cut into syllable segments using the syllable boundaries. Since the syllable boundaries can have errors, we propose a segmentation method called tri-tone to improve the robustness of the boundaries. In each tri-tone segment, we include half of the frames from the previous segment and half of the frames from the next segment. Figure 3 illustrates how tri-tone segments are more tolerant to inaccurate force alignment (dashed line) than syllable segments (solid line). In addition, models trained with tri-tone segments are more robust against tonal coarticulation, because tri-tone segments include some information about the surrounding segments.

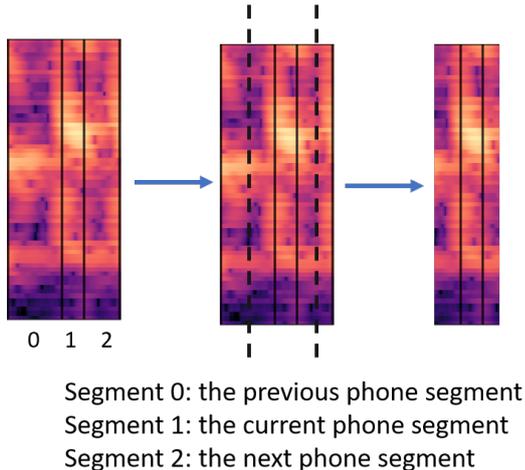


Fig. 3. Example of tri-tone segments being more robust when the force alignment results are inaccurate

Note that not all of the frames from neighboring segments are included in a tri-tone segment. This is because including the neighboring segments can cause ambiguities to occur in the overall pitch curve. The spectrogram on the right part of Figure 4 includes all frames in neighboring segments, where the pitch contours suggest T3, while in reality, it is T1. Meanwhile, a tri-tone segment accurately reflects a T1, shown on the left part of the figure.

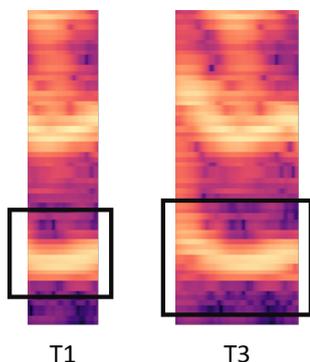


Fig. 4. An example of frame segmentation that causes ambiguities because it includes too many frames in neighbor segments

In contrast to the di-tone method used in [11], our method preserves an additional section before the tone segment itself, because of the bi-directional effect of tone co-articulation [5].

In addition, segment features are extracted from force alignment results. A segment feature consists of the duration and the syllable embedding of the current syllable. The duration has already been proved to be useful in previous literature [1], [8], [10]. Since different finals can have the same tone in Mandarin, adapting the model to different types of finals can boost the model’s performance. Similar to speaker adaptation methods [16], we can give the model a vector that uniquely identifies individual syllables, and then the tone classifier can adapt to different syllables to have a better performance. Since the alignment results from the ASR model are in syllable-level transcripts, we can reuse this information to create syllable embeddings, which is a one-hot encoding of the syllable index, shown in Table I.

Note that this method assumes that the annotated syllable IDs match the actual pronunciation if there are any pronunciation errors made by the speakers. If the syllable embeddings do not indicate the actual syllable pronunciation, the syllable embeddings provide false information to the model and lead to performance degradation. However, we do not need to consider this case because all of AISHELL-3’s labels represent the actual pronunciation.

B. Feature Extraction of the Contextual Model

Since context information is important for continuous tone recognition, we propose context tri-tone and context segment features to cover such information. A context tri-tone consists of the tri-tone segments of n preceding syllables, the current syllable, and n succeeding syllables, where n is the context size. Similarly, the context segment feature includes the segment features of n preceding syllables, the current syllable, and the succeeding syllables.

III. DATA

AISHELL-3 [17] is used in the following experiments. AISHELL-3 is a large-scale Mandarin speech corpus with Chinese pinyin transcripts. It contains approximately 85 hours of mono-channel recordings from 218 native Mandarin speakers with a sample rate of 44.1 kHz.

The transcripts are labeled manually to reflect the actual pronunciation perceived by a human and to overcome the difficulties caused by Homography, Tone Sandhi, Erization, and more [17].

For our experiments, the data is split into a training set and a testing set with 174 and 44 speakers, respectively. The train set contains 65.6 hours of speech, and the test set contains 9.4 hours of speech. There is no speaker overlap between the two subsets.

The number of data samples in the training set per tone is shown in Table II. Since there is a significantly larger amount of initials in the training data than other tones, a portion of initials is randomly removed during the training process to avoid performance degradation caused by data imbalance. The

TABLE II
NUMBER OF SAMPLES OF EACH TONE IN TRAINING DATA

Tone	T0	T1	T2	T3	T4	T5
Count	751150	160885	179606	122707	253441	46609
Percentage	49.6	10.6	11.9	8.10	16.7	3.08

probability of removing an initial is equal to the inverse of the proportion of the number of initials to the total number of samples in the training set.

IV. EXPERIMENTS

A. Data Processing

A time delay neural network (TDNN) [18] model [19] pre-trained using the data from DataTang [20] is used to perform force alignments. The model is deliberately chosen to have a different training dataset than the one used in our experiment to train the tone classifier, in order to include imperfect force alignment results. The syllable error rate of this model tested on AISHELL-3 is 6.29%.

Then we obtain the tri-tone Mel-spectrogram segments as describe above. When the tri-tone segments are fed to the neural network in mini-batches, samples are zero-padded to the maximum length found in the batch.

B. Configurations

In all of the following experiments, a variant of ResNet is used, shown in Figure 2(a). This model is the same as the one used for far-field speaker recognition [21], except the embedding layer contains only 128 units. For convenience’s sake, this model will be referred to as ResNetSP.

Cross entropy is used as the loss function, and Stochastic Gradient Descent (SGD) with 0.001 as the learning rate and 0.9 as the momentum is used as the optimizer. Also, early stopping and a ReduceOnPlateau learning rate scheduler are used.

In the first experiment, the ResNetSP model is trained without any segment features or context information. This experiment is referred to as **ResNetSP** in the results.

The second experiment is to train a classifier with tone embedding and segment features without context, using the architecture shown in Figure 2(a) plus Figure 2(b). In this configuration, the segment features are passed to a dense layer first, and the output of ResNetSP is concatenated to the output of the dense layer. The concatenated vector is fed to a softmax output layer, which has 128 units. This experiment is referred to as **+SF** in the results.

The third experiment is performed in the same way except the context size is set to 1. Compared to the previous configuration, all tri-tone segment features in a context are fed to the ResNetSP model one at a time, and the outputs of the segments are concatenated into a context tone embedding. Moreover, the context segment features are used, so the weight matrix of the dense layer mentioned in the second experiment is larger. This experiment is referred to as **+SF+C1** in the results.

Other experiments are performed with different context sizes. They are referred to as **+SF+C n** in the results, where n is the context size.

V. RESULTS AND DISCUSSIONS

The results of the experiments are shown in Table III.

In the first experiment, ResNetSP achieves an overall accuracy of 79.5%. Meanwhile, the f1-scores of T3 and T5 are the lowest among all tones. This matches the results in previous studies [1], [3], [8], [11]. Note that the f1-score of T0 is very high. This is likely because the model can easily recognize the fact that there is usually a fast pitch transition before and a smooth pitch transition after an initial.

Starting from here, we use **ResNetSP** as a baseline and compare our proposed framework to it. As discussed in previous sections, we expect that by adding segment features and context information, the model will have a better performance. In particular, i) the f1-score of T5 should be boosted significantly by segment features, and ii) by adding context information, tonal coarticulation should have a much less negative impact on the performance.

Indeed, the results match our expectations. The overall accuracy of **+SF** and **+SF+C1** increase to 84.7% and 89.8%, and f1-scores per-tone also increase significantly. Note that by adding segment features, the f1-score of T0 becomes nearly one. The reason is that syllable IDs themselves can be used to distinguish initials from finals. Additionally, even without the context features, the f1-score of T3 rises by a large margin, as shown in **+SF** and **ResNetSP**. This is likely because T3 has the longest average duration among all tones [5]. Finally, the difference between the f1-scores of T3 and other tones became much less when context information is added. This suggests the importance of context information in helping the model learn about tonal coarticulation.

In **+SF+C4** and **+SF+C5**, there is no significant performance gain or degradation brought by introducing a larger context size. This suggests that the syllables far away from the current syllable do not contribute much to its pitch pattern formation, which matches conclusions from previous prosodic studies that the coarticulation effect caused by neighboring syllables diminishes over time [5]. Therefore, a long-term context is not very helpful in terms of performance, compared to a short-term one.

On top of the overall accuracy and the per-tone f1-scores, some common tone subsequences prone to tonal coarticulation in sentences are selected to test the model. We call these tone subsequences tone patterns for convenience’s sake. Three sequences prone to tonal coarticulation, T4-T4, T4-T3-T4, and T2-T2, are selected based on previous studies [22], [23] and our data analysis. The test accuracy of individual tones in the sequences is shown in the bottom half of Table III.

As expected, as the context size increases, the accuracy of the tone patterns generally increases, suggesting again that the model can handle tonal coarticulation better using segment features and context information. However, **+SF+C5** does not achieve a higher performance in one of the tone patterns,

TABLE III
OVERALL ACCURACY, THE F1-SCORES PER TONE, AND ACCURACY ON TONE PATTERNS

	ResNetSP	+SF	+SF+C1	+SF+C2	+SF+C3	+SF+C4	+SF+C5
Overall Accuracy	0.795	0.847	0.898	0.907	0.924	0.925	0.926
T0	0.939	0.999	0.999	0.999	0.999	0.999	0.999
T1	0.840	0.856	0.903	0.910	0.925	0.927	0.929
T2	0.760	0.802	0.867	0.874	0.894	0.897	0.896
T3	0.729	0.781	0.856	0.874	0.902	0.901	0.902
T4	0.777	0.805	0.854	0.875	0.899	0.899	0.904
T5	0.638	0.820	0.914	0.916	0.928	0.930	0.929
T4,T4	0.690	0.719	0.817	0.830	0.879	0.886	0.893
T4,T3,T4	0.763	0.807	0.891	0.889	0.918	0.928	0.923
T2,T2	0.717	0.761	0.855	0.842	0.864	0.862	0.876

compared to +SF+C4. This confirms that a long-term context does not provide too much useful information.

In summary, our experiment results show that i) segment features are useful, especially for T0 and T5, ii) combined with the context information, the model can deal with tonal coarticulation better, and iii) adding a long-term context does not boost the model's performance by much.

VI. CONCLUSIONS AND FUTURE WORK

In the proposed tone classification framework, by reusing the existing information from the force alignment, segment features and context information are extracted to improve the performance of tone classifiers by enhancing the robustness against tonal coarticulation. In addition, experiment results show that a short-term context is preferred, rather than a long-term one. However, there are some possible improvements to this study. First, the model is both trained and tested on a dataset containing only native speakers. However, to evaluate the performance in terms of CALL applications, the model should be tested on an L2 dataset. Second, our work can be improved by performing data augmentation, such as adding slight noise and applying speed perturbation. The imbalanced data issue can also be handled better by data augmentation, instead of random downsampling. Finally, we should perform experiments using RNN or Transformer models to verify our conclusion about the effect of long-term context.

VII. ACKNOWLEDGEMENTS

This work is supported by the Teaching and Assessment Grant of Duke Kunshan University.

REFERENCES

- [1] X.-D. Wang, K. Hirose, J.-S. Zhang, and N. Minematsu, "Tone recognition of continuous mandarin speech based on tone nucleus model and neural network," *IEICE - Trans. Inf. Syst.*, vol. E91-D, no. 6, pp. 1748–1755, Jun. 2008.
- [2] C. Chen, R. Bunescu, L. Xu, and C. Liu, "Tone classification in mandarin chinese using convolutional neural networks," in *Interspeech 2016*, 2016, pp. 2150–2154. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2016-528>
- [3] L. Yang, Y. Xie, and J. Zhang, "Improving mandarin tone recognition using convolutional bidirectional long short-term memory with attention," in *Proc. Interspeech 2018*, 2018, pp. 352–356. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-2561>
- [4] Q. Gao, S. Sun, and Y. Yang, "ToneNet: A CNN Model of Tone Classification of Mandarin Chinese," in *Proc. Interspeech 2019*, 2019, pp. 3367–3371.
- [5] Y. Xu, "Contextual tonal variations in mandarin," *Journal of Phonetics*, vol. 25, no. 1, pp. 61–83, 1997.
- [6] Sim-Hong Chen and Yih-Ru Wang, "Tone recognition of continuous mandarin speech based on neural networks," *IEEE Transactions on Speech and Audio Processing*, vol. 3, no. 2, pp. 146–150, 1995.
- [7] Y. Cao, T. H. Shuwu Zhang, and B. Xu, "Tone modeling for continuous mandarin speech recognition," *International Journal of Speech Technology*, vol. 7, pp. 115–128, 2004.
- [8] N. Ryant, J. Yuan, and M. Liberman, "Mandarin tone classification without pitch tracking," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 4868–4872.
- [9] H. HUANG and J. ZHU, "Discriminative tonal feature extraction method in mandarin speech recognition," *The Journal of China Universities of Posts and Telecommunications*, vol. 14, no. 4, pp. 126–130, 2007.
- [10] H. Huang, Y. Hu, and H. Xu, "Mandarin tone modeling using recurrent neural networks," *CoRR*, vol. abs/1711.01946, 2017. [Online]. Available: <http://arxiv.org/abs/1711.01946>
- [11] Y. Qian, T. Lee, and Y. Li, "Overlapped di-tone modeling for tone recognition in continuous cantonese speech," in *INTERSPEECH*, 2003.
- [12] S. wen Li, Y. bang Wang, L. che Sun, and L. shan Lee, "Improved tonal language speech recognition by integrating spectro-temporal evidence and pitch information with properly chosen tonal acoustic units," in *Proc. Interspeech 2011*, 2011.
- [13] L. Lugosch and V. S. Tomar, "Tone recognition using lifters and CTC," *CoRR*, vol. abs/1807.02465, 2018. [Online]. Available: <http://arxiv.org/abs/1807.02465>
- [14] W. Hu, Y. Qian, F. K. Soong, and Y. Wang, "Improved mispronunciation detection with deep neural network trained acoustic models and transfer learning based logistic regression classifiers," *Speech Communication*, vol. 67, pp. 154–166, 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0167639314001010>
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778.
- [16] G. Saon, H. Soltan, D. Nahamoo, and M. Picheny, "Speaker adaptation of neural network acoustic models using i-vectors," in *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, 2013, pp. 55–59.
- [17] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "AISHELL-3: A multi-speaker mandarin TTS corpus and the baselines," *CoRR*, vol. abs/2010.11567, 2020. [Online]. Available: <https://arxiv.org/abs/2010.11567>
- [18] V. Peddinti, D. Povey, and S. Khudanpur, "A time delay neural network architecture for efficient modeling of longtemporal contexts," in *INTERSPEECH-2015*, 2015.
- [19] DataTang, "Datatang mandarin asr system." [Online]. Available: <https://kaldi-asr.org/models/m10>
- [20] —, "1,505 hours-mandarin speech by mobile phone." [Online]. Available: <https://www.datatang.ai/dataset/info/speech/35>
- [21] A. Gusev, V. Volokhov, T. Andzhukaev, S. Novoselov, G. Lavrentyeva, M. Volkova, A. Gazizullina, A. Shulipa, A. Gorlanov, A. Avdeeva, A. Ivanov, A. Kozlov, T. Pekhovsky, and Y. Matveev, "Deep speaker embeddings for far-field speaker recognition on short utterances," 2020.
- [22] Q. Li and Y. Chen, "An acoustic study of contextual tonal variation in tianjin mandarin," *Journal of Phonet-*

ics, vol. 54, pp. 123–150, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0095447015000911>

- [23] J. Zhang and J. Liu, “Tone sandhi and tonal coarticulation in tianjin chinese,” *Phonetica*, vol. 68, no. 3, pp. 161–191, 2011. [Online]. Available: <https://doi.org/10.1159/000333387>