

# Cross-lingual Multispeaker Text-to-Speech under Limited-Data Scenario

Zexin Cai\*, Yaogen Yang†, Ming Li\*†

\*Electrical & Computer Engineering, Duke University, Durham, NC, United States

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

ming.li369@duke.edu

## Abstract

Modeling voices for multiple speakers and multiple languages in one text-to-speech system has been a challenge for a long time. This paper presents an extension on Tacotron2 to achieve bilingual multispeaker speech synthesis when there are limited data for each language. We achieve cross-lingual synthesis, including code-switching cases, between English and Mandarin for monolingual speakers. The two languages share the same phonemic representations for input, while the language attribute and the speaker identity are independently controlled by language tokens and speaker embeddings, respectively. In addition, we investigate the model’s performance on the cross-lingual synthesis, with and without a bilingual dataset during training. With the bilingual dataset, not only can the model generate high-fidelity speech for all speakers concerning the language they speak, but also can generate accented, yet fluent and intelligible speech for monolingual speakers regarding non-native language. For example, the Mandarin speaker can speak English fluently. Furthermore, the model trained with bilingual dataset is robust for code-switching text-to-speech, as shown in our results and provided samples.<sup>1</sup>.

**Index Terms:** Text-to-speech, cross-lingual text-to-speech, multi-speaker text-to-speech

## 1. Introduction

Recently, the combination of an encoder-decoder based text-to-spectrogram network and a neural vocoder has allowed machines to synthesize high-fidelity speech that is as natural as human. This technique can well equip text-to-speech (TTS) applications (e.g., audiobook reader, virtual assistants, navigation systems, etc.) in our daily life. However, these models, like Tacotron2 [1], keep a certain level of limitations in controllability regarding latent speech attributes. Thus the models’ robustness is limited and may be incapable of synthesizing speech with various speech characteristics. Then extensions on Tacotron2 have been proposed to address these problems: Yuxuan Wang et al. modeled the latent speech attributes by global style tokens (GSTs) while there are no explicit labels provided [2]. Ye Jia et al. extend the Tacotron2 with conditioned features extracted from a speaker verification system to achieve speaker identity cloning and multispeaker TTS. [3].

However, as bilinguals and multilinguals are commonly seen in today’s world, the speech communication scenario becomes complicated. It is essential for speech analysis tools, including speech recognition and speech synthesis, to adapt this change for maintaining their current performance. The challenge is that languages, mostly, have different grapheme set

and pronunciations between each other. This challenge motivates researchers to find and investigate shared representations between languages for speech analysis [4, 5, 6].

Even with appropriate representations for multiple languages, the model architecture needs to be upgraded in order to achieve multilingual processing for all speech analysis systems. For TTS, approaches are proposed for multilingual synthesis, even cross-lingual synthesis, based on classical statistical parametric speech synthesis (SPSS) [7, 8]. Since the end-to-end TTS models can generate speech with higher quality compared with classical methods, extensions on the end-to-end TTS frameworks also have been explored for multilingual modeling [9, 10, 11, 12]. Normally, the voices of the multilingual TTS training datasets are different. Therefore, most TTS multilingual systems also support multispeaker synthesis. But the cross-lingual synthesis, where we can generate speech with foreign text for a monolingual speaker, is challenging. Yu Zhang et al. had achieved high-quality cross-lingual synthesis in a sufficient-data scenario [10]. Zhaoyu Liu et al. investigated cross-lingual synthesis with limited data for each speaker, But the synthesized speech has moderate quality due to the data sparsity issue [13].

Motivated by the aforementioned works, in this paper, our focus is to achieve cross-lingual multispeaker TTS with limited data from two languages, English and Mandarin. We propose a model that incorporates speaker embedding and language embedding as the conditioned features for multilingual multispeaker TTS. The proposed model can generate high-quality speech for all speakers with respect to their own language. In addition, we investigate cross-lingual synthesis with the same model in a limited-data scenario by involving a bilingual TTS dataset. Results show that language-related knowledge can be transferred from the bilingual speaker to monolingual speakers, which enables us to generate fluent, high-fidelity, and intelligible speech in both Mandarin and English using monolingual speakers’ voices.

## 2. Related works

Developing a multilingual multispeaker (MLMS) TTS model can relieve the efforts of training multiple TTS models used for several voices with different languages. While the voice can be controlled by a text-independent speaker embedding in a multispeaker TTS system [3, 14], TTS regarding multiple languages is more complicated due to different grapheme representations across languages. However, similar pronunciations between different languages can help reduce the gap of cross-lingual text-to-speech. Previously, Huaiping Ming et al. presents a light-weighted bilingual synthesis system that adopts concatenated vectors in the linguistic-feature level to manage two languages in one model. [8]. Bo Li et al. proposed an MLMS TTS approach based on conventional statistical parametric speech

<sup>1</sup><https://caizexin.github.io/mlms-syn-samples/index.html>

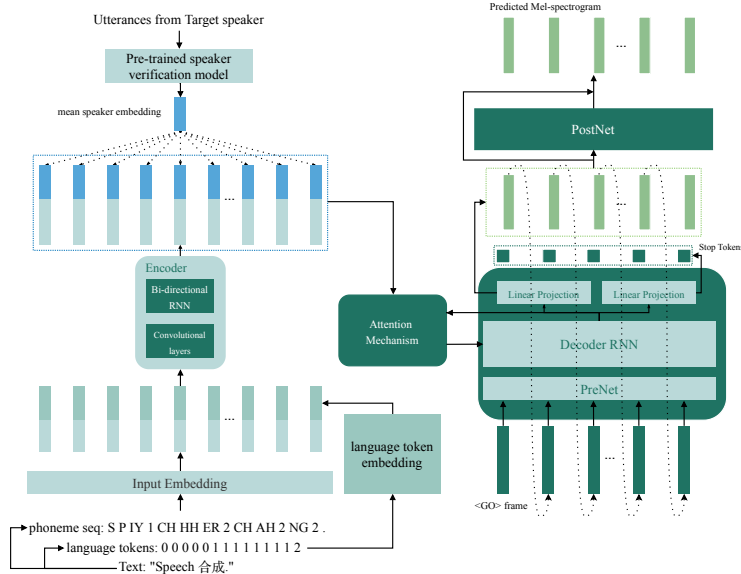


Figure 1: Proposed multilingual multispeaker TTS model

synthesis (SPSS) [7]. They used the international pronunciation Alphabet (IPA) as the input representation and applied cluster adaptive language networks for generating the language-dependent linguistic features, followed by speaker-dependent output layers for different voices.

Then in 2018, Bo Li et al. proposed a novel representation for all languages [6]. The representation, called Bytes, allows speech recognition models and speech synthesis models to manage multilingual processing. The performance of using Bytes in TTS is conducted and evaluated by another group of researchers [10]. Experimental results in [10] showed that phoneme inputs can achieve better performance than Bytes when used as the input for the MLMS TTS model. With sufficient training data (more than 500 hours), their proposed model is able to achieve cross-lingual synthesis with a high naturalness rate. The shared phoneme input is one of the keys to the cross-lingual synthesis, which is also stated in [9]. Similar pronunciations across languages result in close linguistic embedding vectors.

Zhaoyu Liu et al. also used shared phoneme representation and extended the Tacotron2 by incorporating conditional embeddings for MLMS TTS [13], which has a similar structure as our proposed model. However, we have the language-dependent Tacotron encoder designed for allowing the TTS model to synthesized code-switching text. Furthermore, we investigate MLMS TTS with limited data for each language and the performance in cross-lingual synthesis, while [13] investigate multilingual synthesis with limited data concerning each speaker. Xuehao Zhou et al. present a novel method to merge context information between languages by adopting word embedding from a pre-trained language model. Nevertheless, The cross-lingual synthesized speech has moderate quality, as shown in the figures from [11].

### 3. Method

#### 3.1. Input representation

Code-switching is defined as more than one language occurring in one sentence or between sentences, either orally or in

written form. With the world’s globalization, code-switching patterns in speech become a common case in many countries and regions. [15]. The language environment in the globalization inspires more and more bilinguals and multilinguals, which motivates researchers to develop speech processing systems that can handle multilingual challenges. Furthermore, code-switching corpora have been collected and released for research related to speech communication in the recent decade [16, 17], followed with various approaches proposed to address complicated speech analysis, including multilingual automatic speech recognition (ASR), language identification and language diarization with respect to multilingual scenario [18, 19, 20, 21]. Likewise, TTS systems need to be improved for synthesizing natural speech for code-switching sentences [11].

However, one of the main challenges of code-switching TTS is that the grapheme set or the phoneme set between languages are different. Regarding that some phonetic pronunciations between different languages are close. Thus exploring a multilingual TTS model with minimum data requirement, including textual and vocal data, is possible and essential. Previous approaches, which are proposed for addressing multilingual issues in TTS, indicate that shared input representation across languages is one of the keys to realize cross-lingual synthesis [6, 7, 9]. The shared representations include shared phoneme set, international pronunciation alphabet (IPA), and the Bytes coding [6], where the phoneme representation can obtain better performance [10].

In our work, we choose to use a shared phoneme set from CMU dictionary [22] to investigate bilingual multispeaker TTS and cross-lingual synthesis between Mandarin and English. As for Mandarin, the pronunciation representation called pinyin can be converted to CMU phoneme by the pinyin-to-cmu mapping table [23]. Since Mandarin is a tone-language, digits 1 to 6 are used to denote different tones, while ‘0’, ‘1’, ‘2’ are used to mark the lexical stress for English. Although the tone and stress share the same annotations in our input, which may cause ambiguity, we have language identification tokens as another input stream. Moreover, language identification tokens are used to

Table 1: Phonemes (without tone and stress) and their corresponding frequencies in LJ-Speech, DB-1 and DB-4

Phoneme	LJS	DB-1	DB-4	Phoneme	LJS	DB-1	DB-4	Phoneme	LJS	DB-1	DB-4
J		10088	12499	X		8050	11895	Q		5435	7489
IY	28587	54859	85601	EH	26397	3598	11791	AA	16976	11173	23205
L	32893	9420	23510	AY	12079	7479	15619	UW	15345	30630	44593
SH	7957	11456	17804	OW	10201	6921	13698	Y	4426	16540	27793
N	68392	33006	56359	T	65657	8698	26504	JH	4824	8994	13821
AE	21502	27640	42203	NG	7229	25895	36286	AH	102042	12558	33953
G	5901	6960	12298	AW	4248	9654	15397	Z	27845	5749	14135
M	23778	5967	14833	AO	16035	6970	14496	S	43700	5485	17965
UH	2856	7576	11253	W	20352	7151	15411	CH	4751	5118	7940
D	43601	14192	30390	ER	23525	15131	30264	B	15608	7577	15252
F	17018	4111	8890	R	40428	5025	16386	K	27866	3325	12650
HH	13785	7915	14745	EY	14695	4891	10838	P	20212	2496	8607
V	19628		4089	DH	29311		4716	IH	53904		11368
TH	3604		1250	OY	831		595	ZH	607		237
AX	156		418								

generate language-dependent encoding features while preserving the shared information between languages, like close pronunciations. Similarly, ‘0’, ‘1’, ‘2’ are used for language identification in our input representations, where ‘0’ represents the corresponding phoneme or stress annotation is from English, ‘1’ is for Mandarin and ‘2’ for language-unrelated symbols like punctuation marks. Take the phrase ‘speech 合成.’ (speech synthesis.) as an example, two input sequences are obtained after the front-end text processing. One is the phoneme sequence ‘S P IY 1 CH HH ER 2 CH AH 2 NG 2 .’, and the other is the corresponding language identification tokens ‘0 0 0 0 1 1 1 1 1 1 1 1 2’ which has the same length as the phoneme sequence. We break up phonemes with its corresponding tones, e.g., ‘AH2’ is converted to ‘AH 2’, to allow our proposed model to share close pronunciations between Mandarin and English.

### 3.2. Proposed model

Our proposed bilingual multispeaker TTS model is illustrated in figure 1. The input text is converted into phoneme sequence and language token sequence, as introduced in section 3.1. The phoneme sequence is converted to phoneme embedding sequence by a learnable lookup table. Correspondingly, the language tokens are converted to a 64-dimensional language embedding sequence through another learnable embedding table. Two embedding sequences are concatenated together as the input of the Tacotron encoder, which accumulates the linguistic and context characteristics of the input vector sequence with layers of convolutional layers and a bi-directional long short-term memory (BLSTM) layer.

256-dimensional speaker embedding is concatenated with the encoder outputs for conditioning the network to synthesize expected voices. For the speaker embedding, we use the mean embedding derived from all embeddings extracted with a pre-trained speaker verification model [24] by feeding all training utterances of each speaker. We believe that it can induce the same performance as using a trainable lookup table yet costs less training time. Mel-spectrogram is used as the predicted acoustic feature in our bilingual multispeaker TTS model. Accordingly, we trained a neural vocoder, WaveRNN [25], for converting the Mel-spectrogram back to audio signals.

## 4. Experiments

### 4.1. Datasets

Our experiments are conducted with three TTS datasets, including the publicly available LJ Speech (LJS) dataset [26] and two Chinese female voice datasets, DB-1 and DB-4, from Data Baker <sup>2</sup> (LJS, DB-1 and DB-4 are used as representations for both speaker identity and dataset in this section). DB-1 is publicly open, and DB-4 is a commercial one. LJS contains approximately 24 hours of audio-transcript English pairs recorded by a female English native speaker. The DB-1 has approximately 12 hours of Mandarin speech synthesis data recorded by a female Mandarin native speaker. The DB-4 is a bilingual dataset, which contains 12 hours of Chinese audio-transcript pairs, 6 hours of English pairs and 6 hours of code-switching data with a female Mandarin speaker.

Table 1 illustrates the frequencies of all phonemes in three datasets. LJS contains only English utterances, while DB-1 only Chinese utterances. Three consonants, ‘J’, ‘X’, and ‘Q’ do not exist in the English dataset when using shared phoneme representations. However, these three phonemes frequently exist in the Mandarin dataset. On the other hand, 7 phonemes are not presented in the Mandarin dataset while frequently existed in the English dataset, as shown in the table. The bilingual dataset DB-4 contains all phonemes. Most phonemes between two languages share the same representation in our experiments. This indicates that the intersecting shared phonemes may be less challenging to learn by a cross-lingual TTS system compared to those phonemes that only exist in one language. Moreover, the cross-lingual synthesis can be achieved when the model catches the pronunciation similarity of these phonemes between English and Mandarin.

### 4.2. Training setup

We trained two bilingual multispeaker TTS systems with different datasets. The first system, notated by **BLMS**, is the bilingual multispeaker TTS model trained with DB-1 and LJS. The other system, notated by **CLMS**, is the system trained with all datasets, including the bi-lingual dataset DB-4. Although the latter system also can be used for bilingual multispeaker synthesis, we focus on its capability of cross-lingual

<sup>2</sup><https://www.data-baker.com/us.html>




Figure 2: Attention alignments when synthesizing code-switching text ‘其实我很难判断 in my heart I think my Chinese is better but people tell me that my English 是比较好’ (Actually, it’s hard for me to tell. In my heart, I think my Chinese is better, but people tell me that my English is better): (a) The alignment from BLMS with speaker DB-1; (b) The alignment from BLMS with speaker LJS; (c) The alignment from CLMS with speaker DB-1; (d) The alignment obtained from CLMS with speaker LJS; (e) The alignment obtained from CLMS with speaker DB-4;

Table 2: The naturalness mean opinion scores (MOS)

Text type	BLMS		CLMS		
	DB-1	LJS	DB-1	LJS	DB-4
Mandarin	4.08	3.07	3.94	3.10	4.11
English	3.48	3.87	3.74	4.08	3.97
Code-switching	3.53	3.18	3.83	3.27	3.99

Table 3: The speaker similarity mean opinion scores (MOS)

Text type	BLMS		CLMS		
	DB-1	LJS	DB-1	LJS	DB-4
Mandarin	4.43	3.53	4.46	3.31	4.37
English	3.61	4.04	4.18	4.20	4.17
Code-switching	3.81	3.52	4.28	3.38	4.22

synthesis here. All training audios are downsampled to 16 kHz. The vocoder WaveRNN is first pre-trained with the ground truth spectrogram-audio pairs from all three datasets. Then we finetune the pre-trained vocoder model with their ground truth alignment spectrograms after TTS training for each system.

### 4.3. Objective evaluations

The objective evaluation is done by speech synthesis MOS-scale rating, a categorical score from 1 to 5, with 0.5 increments. We ask 16 native Mandarin speakers (all speakers are familiar with English) to rate the synthesized speech concerning naturalness, similarity, and intelligibility. The naturalness is related to the quality of synthesized audios regardless of the content. The speaker similarity score is to measure how close is the synthesized voice to the expected speaker, while the intelligibility evaluates the clarity level of the speech content. We have three types of synthesized text for evaluating the performance, which are Mandarin sentences, English sentences, and code-switching sentences that contain both Mandarin and English content in each sentence. Each type of text has 15 sentences.

The naturalness mean opinion scores (MOS) are shown in table 2. As shown in the table, the quality of synthesized audios reaches around 4.0, While the performance degrades when generating cross-lingual speech for monolingual speakers. For example, DB-1 obtains MOS with 4.12 when synthesizing Mandarin sentences but degrades to 3.64 for English sentences. As shown in table 3, the speech synthesized by our proposed model can well preserve the speaker identity according to the speaker embedding. Most speaker similarity MOS are above 4, while scores lower than 4 can be observed in cross-lingual cases.

Table 4: The intelligibility mean opinion scores (MOS)

Text type	BLMS		CLMS		
	DB-1	LJS	DB-1	LJS	DB-4
Mandarin 4.76	1.90	4.72	<b>3.53</b>	4.81	
English	1.25	4.01	<b>3.86</b>	4.56	4.47
Code-switching	2.60	2.34	4.18	3.84	4.59

The code-switching performance can be clearly observed from table 4. Although BLMS can achieve bilingual multi-speaker synthesis, the cross-lingual synthesis performance is poor, which matches the result in [10]. The cross-lingual synthesized speech is unintelligible as the intelligibility MOS are below 2. However, while involving a bilingual dataset, CLMS is able to generate cross-lingual speech, even in code-switching cases, with intelligible pronunciations. Raters said that the synthesized speech is exactly like a foreign speaker speak another language with the accent from their native language. This indicates that, with our proposed model, using a bilingual dataset can significantly improve cross-lingual speech synthesis, although we only have limited data for each language.

### 4.4. Alignments

In addition, the cross-lingual synthesis performance also can be seen from the attention alignments in Figure 2. The synthesized content is a code-switching sentence. For system BLMS, we can observe clear breaks when the language switches in the sentence for monolingual speakers DB-1 and DB-4 in figure 2 (a) and (b). However, the attention alignments obtained from CLMS are smooth even for monolingual speakers. This also implies that language-related knowledge can be transferred from the bilingual speaker to monolingual speakers with our proposed model.

## 5. Conclusion

We present a bilingual multispeaker TTS approach based on shared phonemic representations. Our proposed model is able to achieve high-fidelity bilingual multispeaker TTS. In addition, results show that, by involving a bilingual dataset, the model is capable of cross-lingual synthesis, even for code-switching synthesis, under the limited-data scenario. We are able to obtain fluent, accented, and intelligible cross-lingual speech as monolingual speakers speak a foreign language.

**Acknowledgments** This research is funded in part by the National Natural Science Foundation of China (61773413) and Duke Kunshan University.

## 6. References

- [1] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan *et al.*, “Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 4779–4783.
- [2] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 5180–5189.
- [3] Y. Jia, Y. Zhang, R. Weiss, Q. Wang, J. Shen, F. Ren, P. Nguyen, R. Pang, I. L. Moreno, Y. Wu *et al.*, “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” in *Advances in neural information processing systems*, 2018, pp. 4480–4490.
- [4] H. of the International Phonetic Association *et al.*, “A guide to the use of the international phonetic alphabet,(1999),” *The Press Syndicate of the University of Cambridge, Cambridge*.
- [5] M. J. Gales, K. M. Knill, and A. Ragni, “Unicode-based graphemic systems for limited resource languages,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2015, pp. 5186–5190.
- [6] B. Li, Y. Zhang, T. Sainath, Y. Wu, and W. Chan, “Bytes are all you need: End-to-end multilingual speech recognition and synthesis with bytes,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 5621–5625.
- [7] B. Li and H. Zen, “multi-language multi-speaker acoustic modeling for lstm-rnn based statistical parametric speech synthesis.”
- [8] H. Ming, Y. Lu, Z. Zhang, and M. Dong, “A light-weight method of building an LSTM-RNN-based bilingual TTS system,” in *2017 International Conference on Asian Language Processing*, 2017, pp. 201–205.
- [9] Y. Lee, S. Shon, and T. Kim, “Learning pronunciation from a foreign language in speech synthesis networks,” *arXiv preprint arXiv:1811.09364*, 2018.
- [10] Y. Zhang, R. J. Weiss, H. Zen, Y. Wu, Z. Chen, R. Skerry-Ryan, Y. Jia, A. Rosenberg, and B. Ramabhadran, “Learning to Speak Fluently in a Foreign Language: Multilingual Speech Synthesis and Cross-Language Voice Cloning,” in *Proc. Interspeech 2019*, 2019, pp. 2080–2084.
- [11] X. Zhou, X. Tian, G. Lee, R. K. Das, and H. Li, “End-to-end code-switching tts with cross-lingual language model,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7614–7618.
- [12] M. Chen, M. Chen, S. Liang, J. Ma, L. Chen, S. Wang, and J. Xiao, “Cross-lingual, multi-speaker text-to-speech synthesis using neural speaker embedding,” *Proc. Interspeech 2019*, pp. 2105–2109, 2019.
- [13] Z. Liu and B. Mak, “Cross-lingual multi-speaker text-to-speech synthesis for voice cloning without using parallel corpus for unseen speakers,” *arXiv preprint arXiv:1911.11601*, 2019.
- [14] E. Cooper, C. Lai, Y. Yasuda, F. Fang, X. Wang, N. Chen, and J. Yamagishi, “Zero-shot multi-speaker text-to-speech with state-of-the-art neural speaker embeddings,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020, pp. 6184–6188.
- [15] A. B. Bernardo, “Bilingual code-switching as a resource for learning and teaching: Alternative reflections on the language and education issue in the philippines,” *Linguistics and language education in the Philippines and beyond: A Festschrift in honor of Ma. Lourdes S. Bautista*, pp. 151–169, 2005.
- [16] D.-C. Lyu, T.-P. Tan, E. S. Chng, and H. Li, “Seame: a mandarin-english code-switching speech corpus in south-east asia,” in *Eleventh Annual Conference of the International Speech Communication Association*, 2010.
- [17] H.-P. Shen, C.-H. Wu, Y.-T. Yang, and C.-S. Hsu, “Cecos: A chinese-english code-switching speech database,” in *2011 International Conference on Speech Database and Assessments*, 2011, pp. 120–123.
- [18] B. H. Ahmed and T.-P. Tan, “Automatic speech recognition of code switching speech using 1-best rescoring,” in *2012 International Conference on Asian Language Processing*, 2012, pp. 137–140.
- [19] N. T. Vu, D.-C. Lyu, J. Weiner, D. Telaar, T. Schlippe, F. Blaicher, E.-S. Chng, T. Schultz, and H. Li, “A first speech recognition system for mandarin-english code-switch conversational speech,” in *2012 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2012, pp. 4889–4892.
- [20] D.-C. Lyu and R.-Y. Lyu, “Language identification on code-switching utterances using multiple cues,” in *Ninth Annual Conference of the International Speech Communication Association*, 2008.
- [21] D.-C. Lyu, E.-S. Chng, and H. Li, “Language diarization for code-switch conversational speech,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7314–7318.
- [22] “The Carnegie Mellon Pronouncing Dictionary,” <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>.
- [23] “Mandarin Pinyin to CMU Dictionary Phoneme Set,” <https://github.com/kaldi-asr/kaldi/blob/master/egs/hkust/s5/conf/pinyin2cmu>.
- [24] W. Cai, J. Chen, J. Zhang, and M. Li, “On-the-fly data loader and utterance-level aggregation for speaker and language recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1038–1051, 2020.
- [25] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, “Efficient neural audio synthesis,” in *Proceedings of the 35th International Conference on Machine Learning*, 2018, pp. 2410–2419.
- [26] K. Ito, “The LJ Speech Dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.