



Few-Shot Cross-Lingual TTS Using Transferable Phoneme Embedding

Wei-Ping Huang¹, Po-Chun Chen², Sung-Feng Huang¹, Hung-yi Lee¹

¹Graduate Institute of Communication Engineering, National Taiwan University

²College of Electrical Engineering and Computer Science, National Taiwan University

{b07201054, b07901062, f06942045, hungyilee}@ntu.edu.tw

Abstract

This paper studies a transferable phoneme embedding framework that aims to deal with the cross-lingual text-to-speech (TTS) problem under the few-shot setting. Transfer learning is a common approach when it comes to few-shot learning since training from scratch on few-shot training data is bound to overfit. Still, we find that the naive transfer learning approach fails to adapt to unseen languages under extremely few-shot settings, where less than 8 minutes of data is provided. We deal with the problem by proposing a framework that consists of a phoneme-based TTS model and a *codebook* module to project phonemes from different languages into a learned latent space. Furthermore, by utilizing phoneme-level averaged self-supervised learned features, we effectively improve the quality of synthesized speeches. Experiments show that using 4 utterances, which is about 30 seconds of data, is enough to synthesize intelligible speech when adapting to an unseen language using our framework.

Index Terms: few-shot, speech synthesis, transfer learning, cross-lingual, low-resource language, self-supervised features

1. Introduction

Recently proposed TTS models based on deep learning techniques [1, 2, 3, 4] are capable of synthesizing natural, human-like speech. However, training end-to-end TTS systems requires large quantities of text-audio paired data and high-quality recordings; thus, data efficiency is a major challenge when developing advanced TTS systems for low-resource languages.

Transfer learning is a common approach when dealing with data efficiency problems. By training on rich-resource languages then fine-tuning on the unseen target language, the model benefits from cross-lingual information and improved data efficiency when adapting to the new language. A recent work from Microsoft [5] shows that multilingual training on 50 languages with a balanced data sampling strategy enables the model to adapt to a new language using only 6 minutes of paired data. Input linguistic features also play an important role in cross-lingual TTS. Many researches leverage common linguistic features across multiple languages to achieve cross-lingual knowledge-sharing. [6, 7, 8, 9] use the combination of phonological features (PFs) and phonemes as inputs for multilingual neural TTS models and show improvements in intelligibility across both seen and unseen languages. [10] utilizes International Phonetic Alphabet (IPA)[11] and an automatic speech recognition (ASR) system to build a unified symbol space and a cross-lingual phoneme mapping, respectively. Unicode-bytes-based models have also been proposed [12, 13], which create a unified input space by encoding input texts in UTF-8 and treating each byte as a token. Some works introduce techniques beyond basic transfer learning. For example, Lrspeech[14] uses a high-resource language to bootstrap ASR and TTS systems

for the low-resource language and leverages additional unpaired text and speech data to boost the performances of both systems.

While a variety of approaches have been proposed, methods that try to achieve cross-lingual knowledge-sharing through a unified input space still have some underlying issues. Using phonological features or IPA allows the low-resource language to leverage meaningful knowledge from rich-resource languages. However, these approaches require linguistic expertise, which can be quite unfeasible for dialects or minority languages. While using Unicode-bytes-based input methods bypasses the need for extra data and knowledge, transfer learning of phonological information cannot be achieved as bytes only encode typographical relations. To reduce the need for extra phonological knowledge, we propose learning a meaningful unified latent space for phonemes as a solution.

In this paper, we introduce few-shot cross-lingual adaptation using transferable phoneme embedding. We first pre-train a TTS model by leveraging data from high-resource (source) languages and then try to adapt it to low-resource (target) languages. To better transfer the cross-lingual knowledge from the source languages and tackle input space mismatch across languages, we propose a *codebook* module as shown in Fig 1b. The module utilizes speech representations extracted from self-supervised learned models to generate a better phoneme initialization for the target language. Compared to the baseline, our method generates more intelligible and natural speeches. Furthermore, experiments show that our method can produce intelligible speech on unseen language even with only 4 utterances, which is about 30 seconds of utterances from the low-resource target language. In comparison, the baseline requires more than 10 times the amount of data to generate intelligible speech.

2. Few-Shot Learning

Under the few-shot learning setting, we use *tasks* to evaluate a model at test time. A *task* contains k training data and q testing data. The model learns from the k training data and is later evaluated on the q testing data. We call this a k -shot task with q queries. Our work deals with few-shot learning on cross-lingual TTS, where the testing tasks consist of data from unseen languages, and the goal is to learn a TTS model for the unseen language using only few-shot training data. We focus on learning one language at a time, meaning all data in a task belong to the same language.

3. Method

We propose a novel transfer learning framework for few-shot cross-lingual adaptation on TTS by utilizing a better phoneme embedding initialization. We illustrate the overall framework in Fig 1. We first provide basic settings in Sec 3.1, then introduce the baseline method in Sec 3.2, and finally describe our proposed method in Sec 3.3.

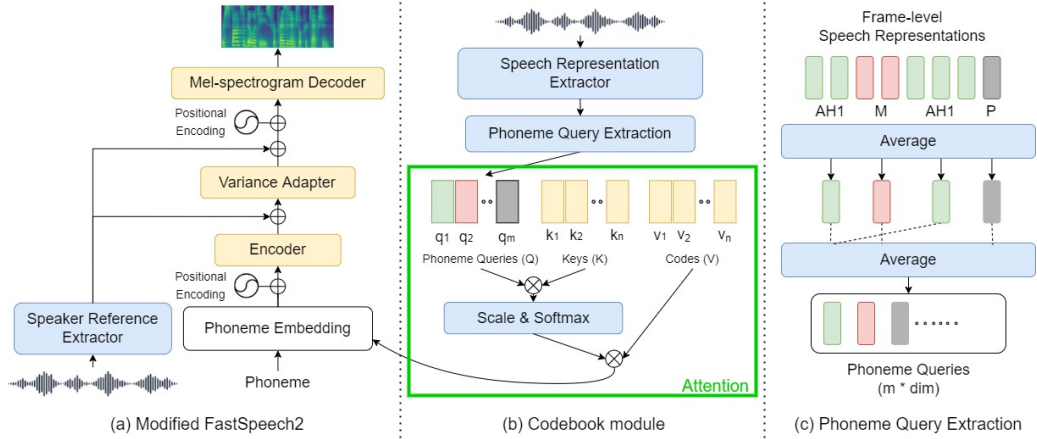


Figure 1: Illustration of the overall framework. (a) Modified multispeaker TTS architecture, (b) the codebook module and (c) the phoneme query extraction. The yellow blocks indicate trainable parts, while the blue blocks indicate fixed pre-trained models or non-trainable parts. n : the size of codebook, m : the size of phoneme set of each language, dim : the dimension of speech representations, \otimes : matrix multiplication.

3.1. Basic Settings

We choose a modified version of FastSpeech 2[15] from a previous paper[16] as our TTS architecture. However, the framework is not restricted to FastSpeech 2. For the multispeaker part, we extract speaker embeddings from a pre-trained speaker verification model. The speaker embeddings are then added to the inputs of both the variance adaptor and the decoder of FastSpeech 2. Fig 1a shows the architecture of the modified FastSpeech 2. In addition to text-speech paired data, our method requires phoneme boundaries, which are already provided since they are needed for training FastSpeech 2. Note that under the few-shot setting, phoneme boundaries can still be manually aligned or forced aligned without too much labeling cost, so our framework can still be easily applied to other TTS architectures.

3.2. Baseline: Multilingual Training

For the baseline method, we train the TTS model on multiple languages together, then fine-tune the model on data from unseen languages. Since the TTS model is trained on multiple different languages, each with its own unique phoneme set, we use a randomly initialized language-dependent phoneme embedding table for each language. When testing, we also have to randomly initialize a phoneme embedding table for the unseen target language.

3.3. Proposed: Phoneme Embedding Transfer

When adapting to an unseen language using the baseline method, the encoder, the variance adaptor, and the decoder part of the model can be transferred. However, the embedding layer part has to be re-initialized, as mentioned in Section 3.2. We introduce a novel *codebook* module to project phonemes from different languages into a unified latent space and generate embedding tables. Our design is therefore capable of achieving phoneme-level knowledge transfer across multiple languages by sharing a common latent space.

3.3.1. Phoneme Query Extraction

To project phonemes from different languages into a unified latent space, we first extract *phoneme queries* from audios as

shown in Fig 1c. For every distinct phoneme that appeared in an audio, we average frame-level speech representation vectors belonging to that phoneme according to its boundaries and take the result as a temporary representation of that phoneme. For multiple audios, we average all temporary representations of the same phonemes again to get the final *phoneme queries*. For phonemes not appeared at all, we set them to zero vectors. Note that the process works regardless the phoneme has or hasn't been seen before. Since this is a general framework, it is possible to incorporate any kind of speech representation. Thus, we consider several different kinds of representations in this paper, such as Mel-spectrograms, or representations extracted from self-supervised learning (SSL) models such as hubert[17], wav2vec 2.0[18], and XLSR-53[19].

3.3.2. Codebook Module

The *codebook* module consists of phoneme query extraction and an attention module with learnable keys and values as shown in Fig 1b. We denote learnable keys and values as *Keys* and *Codes*. *Phoneme queries* produced in the previous step are inputted as queries of the attention module. We apply scaled dot-product attention[20]. We expect *Keys* to capture the patterns in phoneme queries, and *Codes* to construct a unified latent space shared across different languages. *Codes* can be viewed as a phoneme embedding basis of the latent space for all phonemes.

3.3.3. Training and Fine-tuning

In a single training step, every batch contains data from a single language. We then split the batch into two groups. Speech representations from one group are passed through the *codebook* module to generate a phoneme embedding table. The training loss is calculated using the embedding table and the data from the other group.

When adapting to an unseen language, the only difference between our method and the baseline method is the phoneme embedding initialization stage. We generate a phoneme embedding table using the few-shot training data, and take the result as the initialization instead of randomly initializing it. The *codebook* module is used to generate the phoneme embedding table

at the beginning and is ignored at the fine-tuning stage.

4. Dataset

We use Mandarin, French, German, Japanese, and Korean in our experiments. Japanese and German are chosen for testing and the other 4 languages for training. We use LibriTTS[21] train-clean-100 subset for English (*en*), AISHELL-3[22] for Mandarin (*zh*), GlobalPhone[23] for French (*fr*), CSS10[24] for German (*de*) and Korean (*ko*), and JSUT[25] for Japanese (*jp*). We use Montreal Forced Aligner[26] to generate phoneme boundaries, and every language is independently aligned with their own phoneme set.

5. Experiments

5.1. Training setup

We split 10% of the data of each language for validation. During the training stage, we extract 40 samples from the same language for each batch. 32 samples are used to generate the phoneme embedding table, while the remaining 8 samples are used to calculate the loss. We split the samples in such a way that any phoneme that appears in the remaining 8 samples also appears at least once in the 32 samples. This phoneme coverage condition ensures that embeddings used at the training step are meaningful since the *phoneme queries* are only meaningful for phonemes that appeared in those 32 samples. The parameters are trained using Adam optimizer[27] with an initial learning rate of 0.001 for 50K iterations. Warmup for 4K iterations and exponential learning rate decay is applied. For the *codebook* module, the codebook size is set to 128, and we use multi-head attention[20] with 4 attention heads, each with 64-dim *Codes*. The dimension of the embedding table of FastSpeech 2 is set to 256 to match the output from the *codebook* module. We use MelGAN[28] vocoder from a pre-trained checkpoint released by the MelGAN’s authors to synthesize audios from Mel-spectrograms for all experiments.

5.2. Few-Shot Language Adaptation

For few-shot language adaptation, we test on two dissimilar languages: Japanese and German. We first find the minimum amount of data required for FastSpeech 2 baseline adaptation. We find that 64 shots, which is about 8 minutes of audio data, is the minimum amount of data required to fine-tune the multilingual baseline model to produce intelligible speech. Since we expect to push our model to the limit, we test our model under 4, 16, and 64-shot settings. The extreme 4-shot setting is very challenging since it contains only about 30 seconds of data. We compare models trained with features from the 24th layer of hubert-large, Mel-spectrogram features, and the baseline method. The ground truth audios are also evaluated for reference. Ground truth audios are first transformed into Mel-spectrograms and then resynthesized back to audios to show the influence of the vocoder. During testing, the speaker embeddings are extracted from the few-shot training data of each task.

We use Character Error Rate (CER) and Mean of Opinion Score test (MOS) as our evaluation metrics. Synthesized recordings are sent to Azure Speech-To-Text to get CER as a large-scale intelligibility metric. CER is averaged over 20 tasks to reduce the performance variance since only 30 seconds of training data is used in the most extreme case. Each task contains 64 queries. We construct the tasks in such a way that any phoneme that appears in the 64 queries also appears

Table 1: CERs[%] of recordings produced by different models under few-shot settings.

Lang.	Shots	GT	Baseline	Mel	Hubert
<i>jp</i>	4		81.27	44.51	25.26
	16	16.52	48.61	32.57	21.19
	64		28.41	25.30	18.86
<i>de</i>	4		76.19	51.14	27.16
	16	6.79	56.52	39.87	18.94
	64		32.55	30.08	15.11

Table 2: MOS of recordings produced by different models under few-shot settings.

Lang.	Shots	GT	Baseline	Mel	Hubert
<i>jp</i>	4		-	2.21	2.80
	16	4.32	-	2.79	3.01
	64		2.97	3.28	3.42
<i>de</i>	4		-	1.83	2.37
	16	3.55	-	2.24	2.61
	64		2.32	2.46	2.72

at least once in the few-shot training set. This way, we can ensure all phonemes used in the queries have a corresponding trained phoneme embedding. For the Mean of Opinion Score (MOS) test, we randomly sample 20 different sentences from all queries. Each sentence is rated by at least 5 individuals, and over 40 individuals are invited for both languages. The baseline method fails to produce intelligible speech under the 4-shot and 16-shot settings, so we exclude them from the MOS test. CER and MOS results are summarized in Table 1 and Table 2. **Hubert** represents the model trained with features from the 24th layer of hubert-large, and **GT** represents ground truth audios resynthesized by the vocoder.

Although our framework consistently achieves a higher naturalness score than other baselines, the overall score is low due to the few-shot setting. Some Japanese audios have been reported to sound like Korean, and some audios have unnatural intonation. Nevertheless, by using Mel-spectrogram features and the *codebook* module, our method outperforms the baseline under few-shot settings, especially in the 4-shot and 16-shot categories, where the baseline model fails to produce intelligible speech. Furthermore, by using SSL features, the performance can be further improved. This suggests that transfer learning on phoneme embeddings improves the model’s ability to generalize to different languages, and the model does indeed capture more shared knowledge between multiple languages.

5.3. Feature Selection

Since what we proposed is a general framework, it is possible to incorporate any kind of speech representation. Thus, we test our framework on several different extractors. We choose 4 different feature extractors: Mel-spectrogram converter, hubert-large, wav2vec 2.0-large, and XLSR-53. For SSL models with multiple layers, we test on all 24 layers of each model. Both hubert-large and XLSR-53 are SSL models built on the wav2vec 2.0 architecture. Hubert-large and wav2vec 2.0-large are both trained on English, but with different training objectives. XLSR-53 has the same training objective as wav2vec 2.0-large, but is trained




Figure 2: CERs of recordings produced by different feature extractors. Top: Japanese. Bottom: German.

on multiple languages.

Since there are dozens of models to evaluate, we only use CER as the evaluation metric. Each reported CER is averaged over 5 tasks, where each task contains 64 queries for CER testing. We use 16-shots for each task to show the transferability gap between the baseline method and our proposed framework since the baseline fails to synthesize intelligible speech under this setting. Results are summarized in Figure 2. We have also conducted the same experiment under the 64-shot setting. While the performance gap between the baseline method and our framework is narrower under the 64-shot setting, it still shows a similar trend as the 16-shot experiment. However, the result is not shown due to the space limit.

Most layers of SSL features outperform Mel-spectrogram features, and using Mel-spectrogram features is better than the baseline method. In general, XLSR-53 performs slightly better than wav2vec2.0-large, indicating that multilingual training is beneficial to multilingual tasks since XLSR-53 is built on wav2vec2.0 and trained on multiple languages. Hubert-large performs on par with XLSR-53 in most layers; however, its 24th layer shows an outstanding ability to extract features adaptable to different languages. We also observe that performances tend to heavily deteriorate at the last few layers for all three models, except the 24th layer of hubert-large. The exact reason is left unexplored.

Note that using features from some layers results in a worse error rate than the baseline method. We find that models using features extracted from those particular layers failed to learn the training languages during the training phase. From our experiments, we can see that the performance varies from layer to layer; therefore, layer selection should also be considered. Experiments show that under our proposed framework, features extracted from the 24th layer of hubert-large perform the best among all layers.

5.4. Phoneme Mapping Discovery

Inspired by [10], we study what the *codebook* module has learned by analyzing its attention. First, we choose 256 sentences from every language in such a way that all phonemes appear at least once. Then, we generate the phoneme embedding table from them and extract the attention weights from the *codebook* module. For any two distinct phonemes (which may come from different languages), cosine similarity is calculated

Table 3: Top-5 phoneme mapping discoveries of 10 *en* phonemes. A phoneme is marked blue if it shares the same IPA with the source phoneme.

Source	IPA	Top-5 closest phonemes
<i>en</i> -AA0	ɑ	<i>de</i> -a1, <i>de</i> -&1, <i>zh</i> -a1, <i>zh</i> -a4, <i>de</i> -O0
<i>en</i> -AO1	ɔ	<i>de</i> -O0, <i>de</i> -O1, <i>jp</i> -w, <i>ko</i> -wv, <i>fr</i> -O
<i>en</i> -CH	tʃ	<i>de</i> -J, <i>zh</i> -ch, <i>zh</i> -c, <i>zh</i> -q, <i>ko</i> -ch
<i>en</i> -D	d	<i>de</i> -d, <i>de</i> -t, <i>jp</i> -d, <i>ko</i> -t0, <i>ko</i> -c0
<i>en</i> -EY2	eɪ	<i>ko</i> -we, <i>zh</i> -ei3, <i>zh</i> -ei2, <i>de</i> -e0, <i>jp</i> -e
<i>en</i> -M	m	<i>de</i> -m, <i>jp</i> -m, <i>fr</i> -M, <i>zh</i> -m, <i>ko</i> -mm
<i>en</i> -OW2	oʊ	<i>zh</i> -ou2, <i>zh</i> -ou3, <i>de</i> -U0, <i>de</i> -U1, <i>zh</i> -ou4
<i>en</i> -S	s	<i>jp</i> -sh, <i>de</i> -s, <i>ko</i> -ss, <i>zh</i> -s, <i>fr</i> -S
<i>en</i> -TH	θ	<i>jp</i> -z, <i>zh</i> -z, <i>zh</i> -f, <i>de</i> -f, <i>jp</i> -u
<i>en</i> -Y	j	<i>de</i> -j, <i>jp</i> -ry, <i>ko</i> -ii, <i>de</i> -I0, <i>jp</i> -i

between their attention weights from 4 attention heads and averaged as the final *phoneme mapping score*. If two phonemes share similar attention weights among the attention heads, the score will be high, and vice versa.

For each phoneme, we list the top-5 phonemes from different languages with the highest *phoneme mapping scores* and manually map them onto the International Phonetic Alphabets (IPA). Partial results of *en* phonemes are shown in Table 3. Even in the mismatched cases, their corresponding IPA symbols still sound very alike. For example, *ko*-ii, *de*-I0, and *jp*-i all sound like *ee*, as in *meet*. There are also some failed cases, such as *en*-TH. We find that the other five languages lack phonemes that correspond to 'θ' in IPA, and therefore don't share the same IPA as *en*-TH. Intuitive mappings indicate that our framework can indeed learn a unified phoneme latent space successfully.

6. Conclusions

In this paper, we propose a framework for few-shot cross-lingual adaptation on text-to-speech. By learning a unified phoneme latent space using SSL representations, the model can adapt to an unseen language with only 30 seconds of data. Different experiment settings show that our method outperforms the naive transfer learning baseline by a large margin in terms of performance and data efficiency.

7. References

- [1] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, “Tacotron: Towards End-to-End Speech Synthesis,” *Proc. Interspeech 2017*, pp. 4006–4010, 2017.
- [2] 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 (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [3] W. Ping, K. Peng, A. Gibiansky, S. O. Arik, A. Kannan, S. Narang, J. Raiman, and J. Miller, “Deep Voice 3: Scaling Text-to-Speech with Convolutional Sequence Learning,” in *International Conference on Learning Representations*, 2018.
- [4] J. Sotelo, S. Mehri, K. Kumar, J. F. Santos, K. Kastner, A. C. Courville, and Y. Bengio, “Char2Wav: End-to-End Speech Synthesis,” in *ICLR (Workshop)*, 2017.
- [5] J. Yang and L. He, “Towards Universal Text-to-Speech.” in *INTERSPEECH*, 2020, pp. 3171–3175.
- [6] A. Gutkin, “Uniform Multilingual Multi-Speaker Acoustic Model for Statistical Parametric Speech Synthesis of Low-Resourced Languages,” *Proc. Interspeech 2017*, pp. 2183–2187, 2017.
- [7] A. Gutkin and R. Sproat, “Areal and Phylogenetic Features for Multilingual Speech Synthesis,” *Proc. Interspeech 2017*, pp. 2078–2082, 2017.
- [8] I. Demirsahin, M. Jansche, and A. Gutkin, “A Unified Phonological Representation of South Asian Languages for Multilingual Text-to-Speech,” in *Proc. The 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages*, pp. 80–84.
- [9] G. Maniati, N. Ellinas, K. Markopoulos, G. Vamvoukakis, J. S. Sung, H. Park, A. Chalamandaris, and P. Tsiakoulis, “Cross-lingual Low Resource Speaker Adaptation Using Phonological Features,” *arXiv preprint arXiv:2111.09075*, 2021.
- [10] Y.-J. Chen, T. Tu, C.-c. Yeh, and H.-Y. Lee, “End-to-End Text-to-Speech for Low-Resource Languages by Cross-Lingual Transfer Learning,” *Proc. Interspeech 2019*, pp. 2075–2079, 2019.
- [11] I. P. Association, I. P. A. Staff *et al.*, *Handbook of the International Phonetic Association: A guide to the use of the International Phonetic Alphabet*. Cambridge University Press, 1999.
- [12] M. He, J. Yang, L. He, and F. K. Soong, “Multilingual byte2speech models for scalable low-resource speech synthesis,” *arXiv preprint arXiv:2103.03541*, 2021.
- [13] 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 *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5621–5625.
- [14] J. Xu, X. Tan, Y. Ren, T. Qin, J. Li, S. Zhao, and T.-Y. Liu, “Lrspeech: Extremely low-resource speech synthesis and recognition,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 2802–2812.
- [15] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech 2: Fast and high-quality end-to-end text to speech,” *arXiv preprint arXiv:2006.04558*, 2020.
- [16] S.-F. Huang, C.-J. Lin, and H.-y. Lee, “Meta-TTS: Meta-Learning for Few-Shot Speaker Adaptive Text-to-Speech,” *arXiv preprint arXiv:2111.04040*, 2021.
- [17] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [18] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [19] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, “Unsupervised cross-lingual representation learning for speech recognition,” *arXiv preprint arXiv:2006.13979*, 2020.
- [20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [21] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, “LibriTTS: A corpus derived from LibriSpeech for text-to-speech,” *arXiv preprint arXiv:1904.02882*, 2019.
- [22] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, “Aishell-3: A multi-speaker mandarin tts corpus and the baselines,” *arXiv preprint arXiv:2010.11567*, 2020.
- [23] T. Schultz, “Globalphone: a multilingual speech and text database developed at Karlsruhe university,” in *Seventh International Conference on Spoken Language Processing*, 2002.
- [24] K. Park and T. Mulc, “Css10: A collection of single speaker speech datasets for 10 languages,” *arXiv preprint arXiv:1903.11269*, 2019.
- [25] R. Sonobe, S. Takamichi, and H. Saruwatari, “JSUT corpus: free large-scale Japanese speech corpus for end-to-end speech synthesis,” *arXiv preprint arXiv:1711.00354*, 2017.
- [26] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi.” in *Interspeech*, vol. 2017, 2017, pp. 498–502.
- [27] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [28] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, “Melgan: Generative adversarial networks for conditional waveform synthesis,” *Advances in neural information processing systems*, vol. 32, 2019.