



Data-augmented cross-lingual synthesis in a teacher-student framework

Marcel de Korte, Jaebok Kim, Aki Kunikoshi, Adaeze Adigwe, Esther Klabbers

ReadSpeaker

{marcel.korte, jaebok.kim, aki.kunikoshi, adaeze.adigwe, esther.judd}@readspeaker.com

Abstract

Cross-lingual synthesis can be defined as the task of letting a speaker generate fluent synthetic speech in another language. This is a challenging task, and resulting speech can suffer from reduced naturalness, accented speech, and/or loss of essential voice characteristics. Previous research shows that many models appear to have insufficient generalization capabilities to perform well on every of these cross-lingual aspects. To overcome these generalization problems, we propose to apply the teacher-student paradigm to cross-lingual synthesis. While a teacher model is commonly used to produce teacher forced data, we propose to also use it to produce augmented data of unseen speaker-language pairs, where the aim is to retain essential speaker characteristics. Both sets of data are then used for student model training, which is trained to retain the naturalness and prosodic variation present in the teacher forced data, while learning the speaker identity from the augmented data. Some modifications to the student model are proposed to make the separation of teacher forced and augmented data more straightforward. Results show that the proposed approach improves the retention of speaker characteristics in the speech, while managing to retain high levels of naturalness and prosodic variation.

Index Terms: cross-lingual synthesis, speech synthesis, data augmentation, multilingual synthesis, teacher-student paradigm

1. Introduction

Recent years have brought about vast improvements in the development of TTS systems that are capable of producing synthetic speech that rivals human naturalness [1, 2]. The advent of these models has triggered new research outputs related to cross-lingual synthesis, which is the task of synthesizing speech in a language that a speaker is not proficient in, whilst retaining the speaker's identity.

Although considerable advancements have been made in this domain [3, 4, 5], the presented results often highlight a trade-off between maintaining naturalness, preventing accented speech, and retaining key speaker characteristics. This is possibly caused by the fact that models, particularly non-autoregressive models, have insufficient generalization capabilities to separate speaker and language characteristics from each other. As a result, most cross-lingual models are not fully able to generate unseen speaker-language pairs that satisfy all prerequisites. This problem can be overcome to some extent by training models on multiple speakers per language [5, 6], as the presence of multiple speakers can help to implicitly separate speaker and language information. However, not only is this an infeasible approach for many of the world's languages, it also does not fundamentally provide models with the capability to generalize to unseen speaker-language contexts, meaning that it is often not a fully satisfactory solution.

In this paper, we propose to address this generalization problem by applying the teacher-student framework to the

cross-lingual task. Concretely, the teacher model is not just used to generate teacher forced data through knowledge distillation such as in [7], but also as a tool to generate free-running augmented data for unseen speaker-language combinations. This augmented data is generated to have access to cross-lingual speech in which speaker characteristics in the unseen language remain present. Including this data into student model training lowers the model's generalization burden, as the desired speaker-language combinations are now no longer unseen. As the augmented data typically has lower naturalness and less prosodic variation than teacher forced data, we made several modifications to the student model that allow it to distinguish between teacher forced data and augmented data.

We show that the proposed approach enables us to generate speech that retains speaker characteristics learned from the augmented data, while maintaining naturalness and prosodic variation levels of models trained on teacher forced data only, and that this improvement can be achieved with negligible effects on latency.

2. Related research

Various approaches have been explored in order to resolve the challenges observed in cross-lingual synthesis. For example, [8] proposed to use separate phoneme encoders for each language, while [5] instead chose to use a language embedding that was then added to a large-scale multilingual model. A study on code-switching [3] evaluated both approaches, and showed that separate encoders can generate cross-lingual speech with higher naturalness and less accent, provided that sufficient data is available to train each encoder. If not, phonological features common across languages can be derived [9, 10]. Most of the time, cross-lingual results do highlight a trade-off between producing natural and unaccented speech, and retention of voice characteristics. This is likely caused by insufficient model capabilities to separate speaker and language characteristics. Several papers [5, 6, 11, 12] tried to circumvent this issue by including multiple speakers per language, but the trade-off remained visible [5], and (partial) loss of speaker characteristics [11] or remainders of source language accent [6] were reported in several cases. A recent work [13] tried to address this issue and managed to improve speaker similarity ratings while maintaining naturalness levels, but model training became slower and multiple speakers per language were used for the approach.

A data augmentation approach has recently been applied to various text-to-speech contexts [14, 15, 16, 17]. Most related to our work are [18], where data augmentation is applied in a multilingual low-resource setting for the purpose of knowledge distillation, and [19], where a trained teacher model was used to generate a large corpus of augmented speech with which a student model was trained. Our setup proposes to extend the latter idea into the cross-lingual domain, where the teacher model is instead used to generate augmented data for speaker-language pairs that are unseen during training of the model.

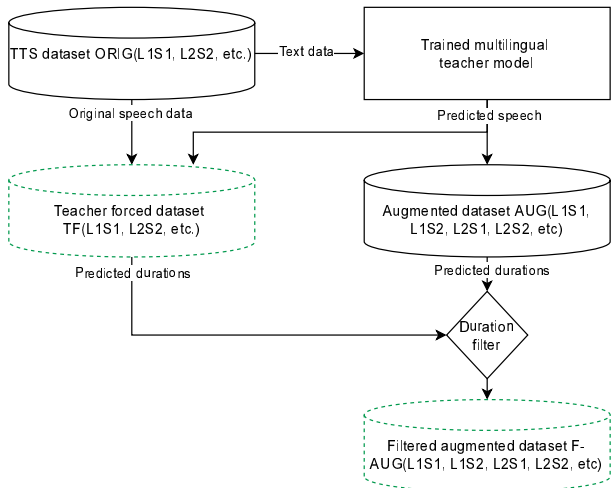


Figure 1: Schematic representation of the proposed data generation approach, where *L1* refers to language one, *S1* refers to speaker one, etc. Green dotted databases indicate databases that are used for student model training.

3. System architecture

The proposed setup of the data generation process is illustrated in Figure 1. We first trained a multilingual teacher model and used it to generate 1) a multilingual teacher forced data set obtained through knowledge distillation, and 2) a multilingual augmented data set obtained through data augmentation that contains both in-lingual and cross-lingual data. The data sets are filtered with a duration filter so that sentences with inaccurate alignments are removed. The proposed student model is then trained on both the teacher forced and the augmented data set and a vocoder is used to generate waveforms. The following subsections will describe these steps in more detail.

3.1. Teacher model

As a teacher model, we used a modified cross-lingual version of the VoiceLoop [20] model with dedicated phoneme encoders for each language [8]. We refer to [21] for a more detailed description of the model’s components.

3.2. Knowledge distillation and data augmentation

As displayed in Figure 1, we used the above-mentioned pre-trained teacher model to generate speech from text. The data set that was used for generating speech was the same that was used to train the teacher model. We refer to this data set as ORIG(L1S1, L2S2, etc.), where *L1* refers to language one, *S1* to speaker one, etc. We generated two sets of data with the teacher model: 1) multilingual teacher forced data (abbreviated as TF(L1S1, L2S2, etc.)) obtained through knowledge distillation such as in [7], and 2) free-running augmented data akin to e.g. [22]. The augmented data is a combination of both seen, in-lingual (AUG(L1S1, L2S2, etc.)), and unseen, cross-lingual (AUG(L1S2, L2S1, etc.)) speaker-language combinations. For each language, the same textual content is used to generate teacher forced, augmented in-lingual and augmented cross-lingual speech.

For the teacher forced data, we used the original output features to condition our current prediction on to ensure equal

length and prosodic variation. For the augmented data sets, there is a risk that prediction errors can compound, which can lead to sentences with unstable alignments that can negatively impact the robustness. We therefore introduced a duration filter that first converted the soft alignments into hard alignments, and then removed sentences with unstable alignments or sentences that were both more than 25% and 30 frames longer or shorter than the original rendering. Applying the duration filter to the data set AUG resulted in a filtered augmented data set F-AUG. This data was then combined with the teacher forced data set TF and used to train the student model with.

3.3. Student model

An implementation of FastSpeech2 [23] was used as the basis for our experiments. For the baseline model, we included a dedicated phoneme encoder per language for cross-lingual modeling, and both an utterance-level encoder for acoustic condition modeling as well as conditional layer normalization to learn speaker information with few extra parameters [24]. This model is able to produce speech with high naturalness but lacks retention of essential speaker characteristics and therefore does not satisfy all demands for cross-lingual synthesis.

As will be shown in Section 4.3.1, the augmented data tends to have less prosodic variation than teacher forced data, which can negatively affect speech generated with the student model. Hence, we made some changes to our proposed model as compared to the baseline model. Concretely, we added an additional embedding that was tasked with learning the difference between teacher forced and augmented data. The embedding output was then included in the duration model and also used to produce an additional scale vector for conditional layer normalization. In addition, batches were reorganized so that sentences from the teacher forced, in-lingual and cross-lingual data with the same phonetic input were part of the same batch.

3.4. Vocoder

Waveforms were reconstructed from predicted mel-spectrogram features with a variant of MelGan [25]. For the training of the universal vocoder, we used samples of all voices present in the experimental setup.

4. Experimental setup

Our experiments are divided into two parts. The first part compares the baseline approach trained on teacher forced data with the proposed approach that was trained on teacher forced and augmented data. We investigate whether the proposed model is able to retain speaker characteristics better than the baseline model, while also maintaining high naturalness levels like the baseline approach. The second part analyzes the detrimental effects that augmented data can have on the naturalness and prosodic variation of student model speech, and analyzes how such negative effects can be overcome. Audio samples related to the experiments can be found online¹.

4.1. Data

We used data from four different speakers, each from a different language. First, a British English female voice was included so that English could function as target language. The remaining

¹<https://readspeaker-ai-research.github.io/AI-Research/publications/cross-lingual-data-augmentation/>

Table 1: F_0 standard deviation of voices in Hz for the English and Polish voice for in-lingual generation.

Model	English	Polish
Original data	29.73	33.35
Teacher - Teacher forced generation	29.90	33.76
Teacher - Augmented generation	21.23	26.22
Student - Baseline TF	27.88	32.15
Student - Baseline TF + F-AUG	25.70	28.57

three voices were each chosen to have one distinctive characteristic from this voice. We included: a) a female Polish voice, where the language family was different (Slavic vs. Germanic language), b) a male German voice, which differs in sex (male vs. female), and c) a female Swedish voice, where the language accentuation was different (pitch-accent vs. stress-accent). All data was recorded by professional speakers in a studio environment, and then down-sampled to 22 kHz.

4.2. Training procedure

For the teacher model training, we randomly selected 10,000 sentences for training and 1,000 sentences for validation for each voice. The model was first pre-trained for 10 epochs and then finetuned for 100 epochs, and trained as described in [21]. With this model, we generated teacher forced data for all languages, in-lingual augmented data of all voices in their source language, and cross-lingual data of the Polish, German and Swedish voice in British English. Duration filtering was applied to the augmented speech following the procedure described in Section 3.2, and about 0.58% of the data was discarded as a result. We used the teacher forced data set to train a baseline student model as described in section 3.3 for 500 epochs, and trained the proposed models with all the filtered augmented and teacher forced data for 400 epochs.

4.3. Evaluation methods

The evaluation focused on three different aspects of the cross-lingual speech: speaker similarity, naturalness, and prosodic variation through a new metric that we called naturalness of intonation (see Section 4.3.1). Speaker similarity was evaluated with an ABX test design. Listeners were presented with a vocoded reference sample of a speaker in its native language, a sample from the baseline model, and a sample from the proposed model. They were asked to choose which of the latter two samples sounded more like the speaker in the reference sample, or if they had no preference. Naturalness was evaluated with a MOS design, using a scale from 1 to 5 with whole point intervals, ranging from 'very unnatural' to 'very natural'.

4.3.1. Evaluating naturalness of intonation

Prosodic variation can be a hard metric to define to participants in a listening experiment. Instead, we decided to focus on intonation, which is one of prosody's most salient aspects and more straightforward to evaluate. Intonation can be defined as the manipulation of the fundamental frequency (F_0) for communicative or linguistic purposes [26]. We evaluated intonation both subjectively and objectively. The naturalness of the intonation was subjectively evaluated with an AB test design, where listeners were asked which of the two presented samples they thought had more natural intonation, or whether they had no

Table 2: Speaker similarity preference ratings of cross-lingual speech for the target voices in English, synthesized with the Polish, German and Swedish voice.

Preference	Polish	German	Swedish
Prefers baseline	25.0%	1.7%	27.5%
No preference	30.0%	0.00%	34.2%
Prefers proposed	45.0%	98.3%	38.3%

Table 3: MOS naturalness ratings of cross-lingual speech (standard deviation between brackets) in English, synthesized with the Polish, German and Swedish voice.

Preference	Polish	German	Swedish
Baseline	3.88 (1.00)	1.74 (1.02)	3.93 (1.02)
Proposed	3.59 (1.14)	2.35 (1.03)	3.77 (1.08)

preference. For objective evaluation, we used F_0 standard deviation as a metric and assumed that a larger standard deviation indicates more natural intonation. To ensure that the computed F_0 standard deviation values were representing intonation accurately, F_0 outliers and unvoiced regions were discarded.

We first evaluated the intonation variation metric on in-lingual data. From Table 1, it can be observed that the augmented data has substantially lower F_0 variation than both original and teacher forced data. While a baseline student model trained only on teacher forced data (TF) could reasonably maintain F_0 variation, scores dropped when filtered augmented data (F-AUG) is also added. The second experiment analyzes whether these effects also occur in cross-lingual synthesis, and evaluates how these effects can be minimized.

4.4. Listening test design

We aimed to make our listening test design compliant with the recommendations made in [27]. We recruited 40 participants, who were randomly assigned one of six batches. Each batch consisted of 18 MOS naturalness screens, 9 ABX speaker similarity screens and 12 AB naturalness of intonation screens, which were randomized within each experiment. The evaluated sentences were between 3 and 5.5 seconds long so that speech fragments would be long enough to evaluate the question asked, but not too long to introduce listener fatigue. For each experiment, three sentences per condition were used per batch. Sentences were reused in other batches, but then generated in one of the other conditions. This way, the text content across batches was similar, minimizing its effects on the results. As such, 18 different sentences were evaluated per experiment. Each part of the listening test was preceded by a demo screen to familiarize participants with the setup. Furthermore, participants were asked about their age group, language background, gender, familiarity with speech technology, and what audio output device they used to allow analysis of inter-group differences.

5. Experimental results

5.1. Evaluating speaker similarity and naturalness

We first evaluated whether the proposed model can better retain target speaker characteristics than the baseline model. Results are shown in Table 2. As can be observed, the proposed model was preferred over the baseline model for all voices with re-

Table 4: Preference ratings of naturalness of intonation. The header names refer to the subsections where the experiment is described.

Preference	5.2.1	5.2.2	5.2.3	5.2.4
Proposed approach	65.0%	48.3%	48.3%	50.8%
No preference	14.2%	25.0%	31.7%	31.7%
Alternative approach	20.8%	26.7%	20.0%	17.5%

gards to speaker similarity. The difference was most marked for the German male voice, where the baseline version did not adequately succeed to retain core speaker characteristics. For the female voices, the difference was not as pronounced, but the proposed model was preferred for these voices as well.

Naturalness results are specified in Table 3. An improvement of 0.61 MOS was found for the German male voice, while for the Swedish and Polish voices naturalness slightly dropped with 0.16 and 0.29 MOS respectively. Across voices, a MOS improvement from 3.18 to 3.24 was found. The higher standard deviation spread of the proposed model for the female voices indicates that its samples have more varied levels of robustness, possibly because of the added augmented data. Indeed, the rendering from the proposed model was rated higher for 39% of sentences for both female voices. Given that the naturalness for both these voices was also above 3.5 MOS, the naturalness scores that were attained can still be considered quite good.

5.2. Evaluating naturalness of intonation

For the second experiment, we evaluated the naturalness of intonation in various contexts. We did not consider speaker similarity in these experiments, as speaker information is treated the same in all models. The Polish voice was chosen to conduct evaluations with, as the Polish language is the most different from the English target language out of the three voices.

5.2.1. Comparing speech from proposed model with augmented speech

We first evaluated how speech from the proposed model compares with augmented speech in terms of intonation to evaluate whether the student model does not just learn intonation from the augmented speech. From Table 4, it can be observed that the naturalness of intonation from the proposed model was preferred in 65.0% of the cases over that of the augmented speech from the teacher model, which was preferred in 20.8% of the cases. The objective evaluation in Table 5 shows that the F_0 standard deviation of the proposed student model was 30.08 Hz, while the F_0 variation of the augmented teacher data was only 24.04 Hz. This decrease in intonation variation is comparable to that of the in-lingual speech in section 4.3.1. We conclude that speech from the proposed model is able to retain the level of intonation variation present in teacher forced data rather than the lower intonation variation from the augmented speech.

5.2.2. Evaluating the effect of the proposed model changes

We then evaluated how the exclusion of the proposed generation mode embedding into the architecture impacted intonation. In Table 4, we can see that the version with the generation mode embedding is preferred in 48.3% of the cases over the version without this embedding. The objective evaluation of F_0 variation from Table 5 shows a decrease to 25.75 Hz without the

Table 5: F_0 standard deviation in Hz of the Polish voice in English for cross-lingual generation.

Model	F_0 std
5.2.1-5.2.4: Student - Proposed	30.08
5.2.1: Teacher - Augmented data	24.04
5.2.2: Student - Baseline, TF + F-AUG	25.75
5.2.3: Student - No in-lingual augmented data	27.81
5.2.4: Student - Proposed, augmented mode	26.00

embedding, compared to 30.08 Hz when the embedding is included. Thus, the embedding plays an important role in retaining the intonation variation present in the teacher forced speech.

5.2.3. Evaluating the effect of in-lingual augmented data

We also evaluated whether including in-lingual augmented data makes distinguishing teacher forced and augmented data more straightforward. As Table 4 shows, the model without in-lingual augmented data is preferred in 20.0% of the cases, while the proposed model with such data is preferred in 48.3% of the cases. The speech without in-lingual data attains a F_0 standard deviation of 27.81 Hz, which is lower than the 30.08 Hz from the proposed model. The results thus suggest that the inclusion of in-lingual data can be a helpful means to better distinguish between teacher forced and augmented data.

5.2.4. Comparing teacher forced and augmented generation

Finally, we evaluated whether speech generated with an augmented embedding differs from speech generated with a teacher forced embedding, using the same model. The results from Table 4 show that the speech with the teacher forced embedding is preferred in 50.8% over the speech with the augmented embedding, while the latter was preferred in 17.5% of the cases. The objective evaluation from Table 5 show a similar pattern, where the F_0 variation decreases from 30.08 Hz to 26.00 Hz. These results show that the embedding and its integration in the model gives the model the capability to learn the distinction between the two data types and should therefore be included.

6. Conclusions and Discussion

In this paper, we have proposed to use a teacher model to not just generate teacher forced data, but also augmented data of unseen speaker-language combinations to reduce generalization problems in the student model. Our results show that the proposed approach helps to improve the retention of speaker characteristics, and that the architectural changes that were made in the student model are important to retain naturalness and prosodic variation levels observed in teacher forced data. These results are achieved using just one speaker per language, making the approach feasible for low-resource languages.

We see several opportunities for further research. First, we anticipate that the benefits of distinguishing between teacher forced and augmented data are not limited to the cross-lingual domain. It would also be interesting to experiment with smaller quantities of augmented data to learn speaker characteristics, as training time would decrease. Research by [28] suggests that reducing the amount of augmented data may be possible. Finally, we would like to evaluate whether we could increase the naturalness of the cross-lingual speech by using a more powerful teacher model or a voice conversion approach.

7. 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 (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [2] A. Łańcucki, “Fastpitch: Parallel text-to-speech with pitch prediction,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6588–6592.
- [3] Y. Cao, X. Wu, S. Liu, J. Yu, X. Li, Z. Wu, X. Liu, and H. Meng, “End-to-end code-switched tts with mix of monolingual recordings,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6935–6939.
- [4] L. Xue, W. Song, G. Xu, L. Xie, and Z. Wu, “Building a mixed-lingual neural tts system with only monolingual data,” *arXiv preprint arXiv:1904.06063*, 2019.
- [5] 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,” *arXiv preprint arXiv:1907.04448*, 2019.
- [6] J. Latorre, C. Bailleul, T. Morrill, A. Conkie, and Y. Stylianou, “Combining speakers of multiple languages to improve quality of neural voices,” *arXiv preprint arXiv:2108.07737*, 2021.
- [7] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech: Fast, robust and controllable text to speech,” in *Advances in Neural Information Processing Systems*, 2019, pp. 3165–3174.
- [8] E. Nachmani and L. Wolf, “Unsupervised polyglot text-to-speech,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7055–7059.
- [9] M. Staib, T. H. Teh, A. Torresquintero, D. S. R. Mohan, L. Foglianti, R. Lenain, and J. Gao, “Phonological features for 0-shot multilingual speech synthesis,” *arXiv preprint arXiv:2008.04107*, 2020.
- [10] 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.
- [11] J. Yang and L. He, “Towards universal text-to-speech,” in *INTER-SPEECH*, 2020, pp. 3171–3175.
- [12] T. Nekvinda and O. Dušek, “One model, many languages: Meta-learning for multilingual text-to-speech,” *arXiv preprint arXiv:2008.00768*, 2020.
- [13] J. Yang and L. He, “Cross-lingual text-to-speech using multi-task learning and speaker classifier joint training,” *arXiv preprint arXiv:2201.08124*, 2022.
- [14] G. Huybrechts, T. Merritt, G. Comini, B. Perz, R. Shah, and J. Lorenzo-Trueba, “Low-resource expressive text-to-speech using data augmentation,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6593–6597.
- [15] Z. Byambadorj, R. Nishimura, A. Ayush, K. Ohta, and N. Kitaoka, “Text-to-speech system for low-resource language using cross-lingual transfer learning and data augmentation,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2021, no. 1, pp. 1–20, 2021.
- [16] W.-N. Hsu, Y. Zhang, R. J. Weiss, Y.-A. Chung, Y. Wang, Y. Wu, and J. Glass, “Disentangling correlated speaker and noise for speech synthesis via data augmentation and adversarial factorization,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5901–5905.
- [17] E. Cooper, C.-I. Lai, Y. Yasuda, and J. Yamagishi, “Can speaker augmentation improve multi-speaker end-to-end tts?” *arXiv preprint arXiv:2005.01245*, 2020.
- [18] 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.
- [19] M.-J. Hwang, R. Yamamoto, E. Song, and J.-M. Kim, “Tts-by-tts: Tts-driven data augmentation for fast and high-quality speech synthesis,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6598–6602.
- [20] Y. Taigman, L. Wolf, A. Polyak, and E. Nachmani, “Voiceloop: Voice fitting and synthesis via a phonological loop,” *arXiv preprint arXiv:1707.06588*, 2017.
- [21] M. de Korte, J. Kim, and E. Klabbbers, “Efficient neural speech synthesis for low-resource languages through multilingual modeling,” *arXiv preprint arXiv:2008.09659*, 2020.
- [22] H. Guo, F. K. Soong, L. He, and L. Xie, “A new gan-based end-to-end tts training algorithm,” *arXiv preprint arXiv:1904.04775*, 2019.
- [23] 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.
- [24] M. Chen, X. Tan, B. Li, Y. Liu, T. Qin, S. Zhao, and T.-Y. Liu, “Adaspeech: Adaptive text to speech for custom voice,” *arXiv preprint arXiv:2103.00993*, 2021.
- [25] K. Kumar, R. Kumar, T. de Boissiere, L. Gestein, 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.
- [26] P. Taylor, *Text-to-speech synthesis*, 2009, p. 227.
- [27] M. Wester, C. Valentini-Botinhao, and G. E. Henter, “Are we using enough listeners? no!—an empirically-supported critique of interspeech 2014 tts evaluations,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [28] M. S. Ribeiro, J. Roth, G. Comini, G. Huybrechts, A. Gabrys, and J. Lorenzo-Trueba, “Cross-speaker style transfer for text-to-speech using data augmentation,” *arXiv preprint arXiv:2202.05083*, 2022.