



Generalizable Zero-Shot Speaker Adaptive Speech Synthesis with Disentangled Representations

Wenbin Wang¹, Yang Song¹, Sanjay Jha¹

¹School of Computer Science and Engineering, University of New South Wales, Australia

wenbin.wang@unsw.edu.au, yang.song1@unsw.edu.au, sanjay.jha@unsw.edu.au

Abstract

While most research into speech synthesis has focused on synthesizing high-quality speech for in-dataset speakers, an equally essential yet unsolved problem is synthesizing speech for unseen speakers who are out-of-dataset with limited reference data, i.e., speaker adaptive speech synthesis. Many studies have proposed zero-shot speaker adaptive text-to-speech and voice conversion approaches aimed at this task. However, most current approaches suffer from the degradation of naturalness and speaker similarity when synthesizing speech for unseen speakers (i.e., speakers not in the training dataset) due to the poor generalizability of the model in out-of-distribution data. To address this problem, we propose GZS-TV, a generalizable zero-shot speaker adaptive text-to-speech and voice conversion model. GZS-TV introduces disentangled representation learning for both speaker embedding extraction and timbre transformation to improve model generalization and leverages the representation learning capability of the variational autoencoder to enhance the speaker encoder. Our experiments demonstrate that GZS-TV reduces performance degradation on unseen speakers and outperforms all baseline models in multiple datasets. The audio samples, detailed network architecture and pre-trained model are available at <https://gzs-tv.github.io/>.

Index Terms: text-to-speech, voice conversion, zero-shot

1. Introduction

With the development of speech synthesis methods, applications based on speech synthesis have become increasingly prevalent [1]. Currently, most speech synthesis models learn to synthesize speech for speakers within the training dataset by memorizing their unique rhythm (speaking rate) and timbre (voice characteristics) during training [2, 3, 4, 5, 6]. For example, VITS [6] can synthesize high-quality speech for in-dataset speakers by utilizing the uncertainty modeling over latent variables and adversarial training. However, these approaches cannot synthesize speech for unseen speakers, hence are unable to meet the increasing demand for personalized speech synthesis.

Recent research has proposed speaker adaptive speech synthesis to address this problem. These approaches can be divided into two categories: few-shot speaker adaptation based and zero-shot speaker adaptation based. Few-shot speaker adaptation approaches [7, 8, 9, 10, 11, 12] typically pre-train a speech synthesis model on multi-speaker datasets, then fine-tune it with a few speech samples from the unseen speaker. Although these approaches can achieve a good quality of synthesized speech, as demonstrated in [13, 14, 15, 16], the requirements of computational resources and transcriptions for fine-tuning limit their application scenarios. In contrast, zero-shot speaker adaptation based approaches [17, 18, 19, 20, 21] jointly train a speaker

encoder with a speech synthesis model. For instance, the VITS-based zero-shot approach YourTTS [22], utilizes a speaker encoder to extract speaker embeddings from reference speech and then utilizes these embeddings as input for YourTTS's speech synthesis model to generate speech for unseen speakers. Similarly, StyleSpeech [23] uses the same method to achieve zero-shot adaptation, and Meta-StyleSpeech [23] introduces meta-learning for StyleSpeech to improve the quality of synthesized speech. These approaches offer a broad range of application scenarios. However, the limited capability of speaker embedding and distribution shift between seen and unseen speakers present challenges for the models' generalizability, resulting in performance gaps between seen and unseen speakers and poor model performance in zero-shot scenarios [24, 11].

To tackle these issues, inspired by recent research on disentangled representation learning and out-of-distribution generalization problems [25, 26, 27], we propose GZS-TV, a generalizable zero-shot speaker adaptation text-to-speech (TTS) and voice conversion (VC) model. Since speech contains highly entangled phoneme and timbre information, we aim to disentangle these two types of information to enhance the generalizability of zero-shot speaker-adaptive speech synthesis. Specifically, GZS-TV employs disentangled learning for both speaker information extraction and timbre transformation to avoid phoneme information leakage into the speaker embedding and to better align the timbre of synthesized speech with the ground truth (GT) speech. Additionally, we leverage the representation learning capability of the variational autoencoder (VAE) [28] to enhance the capacity of the speaker encoder and enrich speaker embedding. The TTS and VC experiments conducted on LibriTTS [29] and VCTK [30] datasets demonstrate that GZS-TV is able to synthesize speech that is more natural and more similar to reference speaker's voice than recent state-of-the-art methods.

2. Methods

The training, TTS, and VC inference procedures of our GZS-TV model are depicted in Figure 1. The subscripts *gt*, *ref* and *src* indicate that the corresponding sequences or embeddings come from the GT, reference, and source speech, respectively. The GT speech is used only during training and is obtained from the seen speakers on the training set, providing phoneme-speech pairs for the model training procedure. The reference speech is also obtained from the same speakers as the GT speech during training, but it can come from any speaker (seen or unseen) during TTS and VC inference procedures. It provides timbre and rhythm information to guide speech synthesis, thereby achieving zero-shot speaker adaptation. The source speech is used only during VC inference and can be from any speaker, providing phoneme and rhythm information for the VC procedure.

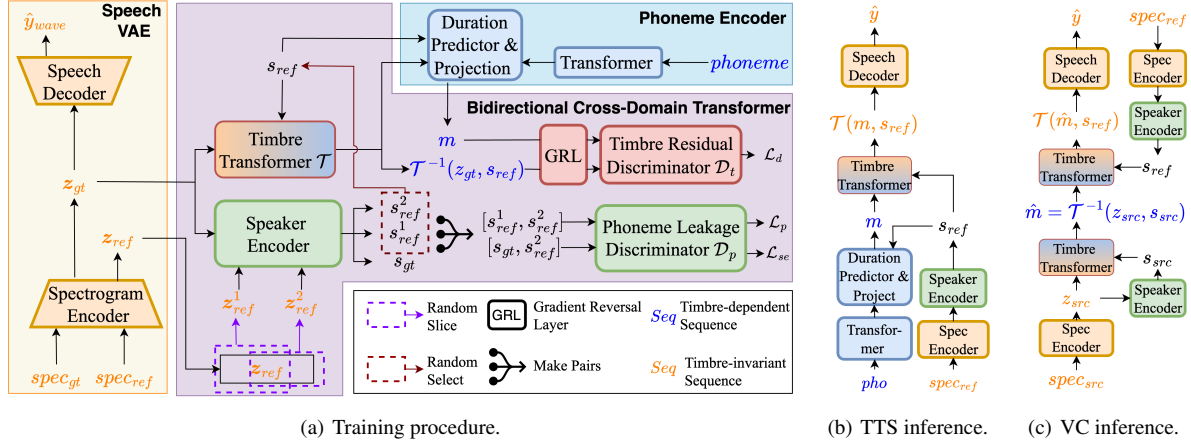


Figure 1: Model diagram depicting (a) training procedure, (b) TTS inference procedure and (c) VC inference procedure.

As shown in Figure 1a, GZS-TV comprises three parts: a speech VAE, a phoneme encoder and a bidirectional cross-domain transformer. During training, the *speech VAE* samples frame-level speech representations z_{gt} and z_{ref} from linear-spectrogram $spec_{gt}$ and $spec_{ref}$, which contain highly entangled phoneme, timbre and rhythm information. It then reconstructs waveform audio \hat{y}_{wave} from z_{gt} . The *phoneme encoder* encodes and projects the phoneme to a frame-level (same as z_{gt}) timbre-invariant phoneme representation m . The *bidirectional cross-domain transformer* bridges the timbre-dependent domain (where z_{gt} belongs) and timbre-invariant domain (where m belongs) with disentangled representations learning, which is the main component in our model. Specifically, the *speaker encoder* can extract speaker embeddings that represent the speaker’s timbre and rhythm characteristics from frame-level speech representations, i.e., extracting speaker embedding s_{gt} from z_{gt} . The *timbre transformer* can perform lossless bidirectional transformations. The forward transformation fuses the speaker embedding with the timbre-invariant phoneme representation to synthesize a timbre-dependent speech representation, e.g., transforming m into $\mathcal{T}(m, s_{ref})$ in Figure 1b. This transformation is used exclusively by the TTS and VC inference procedures. The reverse transformation disentangles and removes the timbre information from the timbre-dependent speech representation based on the speaker embedding and produces a timbre-invariant phoneme representation, e.g., transforming z_{gt} into $\mathcal{T}^{-1}(z_{gt}, s_{ref})$ in Figure 1a. This transformation is only used during training and VC inference procedures. Subsequently, two discriminators are designed to improve the model’s generalizability on unseen speakers. Details of the speaker encoder and phoneme leakage discriminator can be found in Sections 2.1 and 2.2. The timbre transformer and timbre residual discriminator are described in Section 2.3. Our GZS-TV is developed on the VITS baseline, which is similar to YourTTS. In contrast to YourTTS, we remove the requirement for speaker embedding as an input for the speech VAE, so that we can extract the speaker embedding from the latent speech representation. Furthermore, design two discriminators to enhance the model’s generalizability to unseen speakers.

During TTS inference procedure (Figure 1b), GZS-TV begins by extracting the speaker embedding s_{ref} from the reference speech. It then encodes the phoneme sequence pho to a timbre-invariant frame-level phoneme representation m , and transforms it into a timbre-dependent speech representation

$\mathcal{T}(m, s_{ref})$. Finally, GZS-TV generates the synthesized speech waveform \hat{y} from this speech representation. During VC inference procedure (Figure 1c), GZS-TV first extracts speaker embeddings s_{src} and s_{ref} from the source and reference speech, respectively. It then transforms the spectrogram $spec_{src}$ to a timbre-invariant phoneme representation \hat{m} conditioned on s_{src} , and further transforms \hat{m} into a timbre-dependent speech representation $\mathcal{T}(\hat{m}, s_{ref})$. Finally, GZS-TV generates the synthesized speech waveform \hat{y} from this speech representation.

2.1. Extracting Speaker Embedding from Latent Speech Representation

A critical challenge in zero-shot speaker adaptation is extracting accurate and generalizable speaker embeddings from a short segment of reference speech. We find that extracting speaker embeddings from high-level speech representations can address this challenge and enhance the generalization of the extracted speaker embeddings, compared to extracting them from raw audio or spectrogram. This improvement can be attributed to two factors. Firstly, the sampling process is a form of data augmentation that adds noise, which can enhance the model’s generalizability, as shown in the previous study [25]. Secondly, the spectrogram encoder’s representation learning capability can simplify the distribution of speaker embeddings and enable the speaker encoder to learn richer embeddings, leading to better generalization for unseen speakers.

In practice, we built the speaker encoder upon ECAPA-TDNN [31], a popular speaker verification model. To obtain the speaker embedding, we replaced the classifier layers of ECAPA-TDNN with feedforward layers. During each training step of GZS-TV, a speech sample is randomly selected from the same speaker as the GT speech and used as the reference speech. The spectrogram encoder then samples speech representation z_{ref} from this reference speech and divides it into two sub-sequences, z_{ref}^1 and z_{ref}^2 , with an overlap of λ_{ol} frames, where λ_{ol} is a hyperparameter for the phoneme leakage discriminator. Subsequently, the speaker encoder extracts two speaker embeddings s_{ref}^1 and s_{ref}^2 from z_{ref}^1 and z_{ref}^2 , respectively. Finally, one of the speaker embeddings is randomly selected as input for downstream modules to provide the speaker’s timbre and rhythm information. However, both embeddings are used as input for the phoneme leakage discriminator. More details are available in Section 2.2.

2.2. Disentangled Representation Learning for The Speaker Encoder to Avoid Information Leakage

To ensure that the speaker embedding remains free from any leaked speaker-irrelevant information, particularly phoneme information, which could hinder the generalizability of the encoder on unseen speakers [26], we propose a *phoneme leakage discriminator*. This discriminator can detect leakage phoneme information and improve the phoneme-speaker information disentanglement ability of the speaker encoder.

During the training phase, the speaker encoder extracts an additional speaker embedding s_{gt} from the GT speech, in addition to the s_{ref}^1 and s_{ref}^2 embeddings extracted as described in Section 2.1, as shown in Figure 1a. From these three speaker embeddings, two contrastive embedding pairs are constructed: $[s_{ref}^1, s_{ref}^2]$ and $[s_{gt}, s_{ref}^2]$. The first pair is extracted from overlapped speech representations and would include overlapping phoneme information if there is a phoneme information leakage issue in the speaker encoder. On the other hand, the second pair is extracted from different speech representations and would include rare to no overlapping phoneme information, regardless of the existence of phoneme information leakage. If a well-trained discriminator cannot distinguish which pair of speaker embeddings contains more leaked information, we can assume that the leakage is negligible and can be ignored.

To this end, we introduce the phoneme leakage discriminator \mathcal{D}_p , which aims to detect phoneme information leakage and applies an adversarial penalty to the speaker encoder. The adversarial penalty \mathcal{L}_{se} for speaker encoder and the training loss for phoneme leakage discriminator \mathcal{L}_p are defined as follows:

$$\mathcal{L}_p = \mathbb{E}_{s_{ref}^{1,2}, s_{gt}} (\mathcal{D}_p(s_{gt} \oplus s_{ref}^2) - 1)^2 + (\mathcal{D}_p(s_{ref}^1 \oplus s_{ref}^2))^2 \quad (1)$$

$$\mathcal{L}_{se} = \lambda_{se} \mathbb{E}_{s_{ref}^{1,2}} (\mathcal{D}_p(s_{ref}^1 \oplus s_{ref}^2) - 1)^2 \quad (2)$$

where λ_{se} is a parameter that adjusts the weight of \mathcal{L}_{se} and \mathcal{D}_p refers to a feedforward neural network. By playing this min-max game between the speaker encoder and the phoneme leakage discriminator, the speaker encoder is expected to extract embeddings purely related to the speaker’s information.

2.3. Disentangled Representation Learning for Timbre Transformer to Align Timbre Characteristic

We further design the *timbre transformer* and the *timbre residual discriminator* to align the timbre characteristics of synthesized and GT speech. The timbre transformer is based on the normalizing flow [32] in VITS. It comprises multiple affine coupling layers [33] and provides bidirectional lossless transformations between timbre-dependent and time-invariant sequences to meet different requirements during the training and inference procedures. Since the timbre transformer’s transformation is bidirectional and lossless, we argue that enhancing its ability to disentangle and remove timbre information in speech representation during the reverse transformation is equivalent to improving its ability to align the timbre information of synthesized and GT speech during the forward transformation. Thus we utilize the timbre residual discriminator to detect the presence of residual timbre information in the output of reverse transformation and add a penalty to the timbre transformer to improve its inverse transformation. The timbre residual discriminator mainly consists of multiple Res2Net layers [34], an attentive statistics pooling layer [35] and a classification layer.

Specifically, for the output sequence of the timbre transformer’s reverse transformation, as denoted by $\mathcal{T}^{-1}(z_{gt}, s_{ref})$ in Figure 1a, and the time-invariant sequence generated from the phoneme, as denoted by m in Figure 1a, we use the timbre residual discriminator \mathcal{D}_t to determine which of them do not contain timbre information. Then, we introduce the gradient reversal layer (GRL) [36] to invert the gradient. Therefore, if the timbre transformer can fool a well-trained timbre residual discriminator, we can assume that its reverse transformation can disentangle and remove most of the timbre information, and its forward transformation can align the timbre information with GT as well as possible. The timbre transformer and timbre residual discriminator will be optimized in different ways according to \mathcal{L}_d :

$$\mathcal{L}_d = \mathbb{E}_{m, \mathcal{T}^{-1}(z_{gt}, s_{ref})} (\mathcal{D}_t(m) - 1)^2 + (\mathcal{D}_t(\mathcal{T}^{-1}(z_{gt}, s_{ref})))^2 \quad (3)$$

$$\theta \leftarrow \theta - \epsilon(-\lambda_d \frac{\partial \mathcal{L}_d}{\partial \theta}); \nu \leftarrow \nu - \epsilon(\frac{\partial \mathcal{L}_d}{\partial \nu}) \quad (4)$$

where θ is the parameter of the timbre transformer, ν is the parameter of the timbre residual discriminator, ϵ is the learning rate, λ_d is a weight hyperparameter and \mathcal{T} is timbre transformer.

3. Experiments

3.1. Experimental Setup

Training Details. We trained GZS-TV on the clean set of LibriTTS and downsampled all audio samples to 22050 Hz. All modules were trained together in a single training stage. During training, we set both λ_{se} and λ_d to 8, and limited λ_{ol} to a range of 20% to 40% to avoid the discriminator overwhelming the timbre transformer. We used a batch size of 64, employed the AdamW optimizer [37], with $\beta_1 = 0.8$, $\beta_2 = 0.99$, and set the weight decay to 0.01. Additionally, we initialized the learning rate to 2×10^{-4} , with a decay factor of $\gamma = 0.999875$.

Evaluation metrics. We employ the mean opinion score (MOS) [14] to evaluate the naturalness of synthetic speech. To evaluate the speaker similarity of synthetic speech, we use both the speaker embedding cosine similarity (SMCS) [22] and similarity mean opinion score (SMOS). Both MOS and SMOS are rated on a 1-to-5 scale (1 means worst and 5 means best) by 30 native English speakers through the crowdsourcing form, reported with 95% confidence intervals. Following previous research [22, 38], the SMCS is computed by Resemblyzer¹, an off-the-shelf tool for computing speaker embedding. A larger SMCS indicates better speaker similarity. We also provide the word error rate (WER) as the intelligibility metric and a smaller WER indicates more explicitly synthesized speech. We adopt a public pre-trained ASR model² for speech transcription.

Baselines approaches. We compare GZS-TV with several baselines. 1) Ground-Truth: GT speech. 2) Reconstruction: speech reconstructed from z_{gt} through speech decoder. 3) StyleSpeech [23]: a speaker adaptive TTS approach based on style-adaptive layer normalization [23]. 4) Meta-StyleSpeech [23]: another version of StyleSpeech based on meta-learning. Both baselines 3 and 4 are trained on LibriTTS. 5) YourTTS [22]: current state-of-the-art zero-shot speaker adaptive TTS model in English, which is based on VITS like ours. By comparing our model with YourTTS, we can demonstrate that our advancement does not solely result from using a more sophisticated speech synthesis backbone. We use their public check-

¹ <https://github.com/resemble-ai/Resemblyzer>

² <https://huggingface.co/facebook/wav2vec2-large-960h-lv60-self>

Table 1: Zero-Shot TTS evaluation results.

Dataset _{speaker}	VCTK _{unseen}				LibriTTS _{unseen}				LibriTTS _{seen}			
	SMCS	SMOS	MOS	WER(%)	SMCS	SMOS	MOS	WER(%)	SMCS	SMOS	MOS	WER(%)
Ground-Truth Reconstruction	-	-	4.38 ± 0.07	5.1	-	-	4.32 ± 0.09	2.3	-	-	4.31 ± 0.09	2.4
	0.934	4.57 ± 0.07	4.24 ± 0.08	7.4	0.954	4.59 ± 0.08	4.20 ± 0.09	3.0	0.972	4.64 ± 0.08	4.22 ± 0.09	2.8
StyleSpeech [23]	0.730	3.19 ± 0.08	3.19 ± 0.09	13.2	0.762	3.40 ± 0.08	3.25 ± 0.09	6.1	0.780	3.58 ± 0.08	3.43 ± 0.09	3.4
Meta-StyleSpeech [23]	0.738	3.21 ± 0.08	3.18 ± 0.09	13.4	0.762	3.41 ± 0.08	3.21 ± 0.09	6.3	0.781	3.59 ± 0.08	3.42 ± 0.09	3.5
YourTTS[22]	0.783	3.56 ± 0.07	3.64 ± 0.08	11.9	0.788	3.62 ± 0.06	3.65 ± 0.08	6.0	0.806	3.80 ± 0.07	3.78 ± 0.06	5.8
GZS-TV	0.796	3.70 ± 0.05	3.75 ± 0.07	11.2	0.814	3.87 ± 0.06	3.77 ± 0.07	5.9	0.825	3.98 ± 0.08	3.86 ± 0.06	5.8

1. Since YourTTS is trained on volume-normalized GT speech, its results are computed on this type of speech too, in the same way as in its original experiment.




Figure 2: Effect of reference speech length on SMCS.

point, which is pre-trained on VCTK and fine-tuned on LibriTTS. Since baselines 3, 4 and 5 can only synthesize 16 kHz speech, we downsample the synthesized speech of GZS-TV to 16 kHz when evaluating.

Evaluation datasets. We utilized 37 out of 39 speakers in the LibriTTS test set (two speakers had insufficient samples) and 108 out of 109 speakers in the VCTK dataset (one speaker lost transcriptions) for the unseen speakers’ TTS and VC evaluations. Moreover, we selected 37 random speakers from the LibriTTS training set to evaluate seen speakers. To ensure the diversity of the data, we randomly chose 5 test sentences for each LibriTTS speaker and 2 test sentences for each VCTK speaker. Additionally, we conducted YourTTS’s TTS experiment on LibriTTS for reference. Notably, the reference speech in this experiment is considerably longer than in other zero-shot speech synthesis experiments [39, 38], leading to significantly better outcomes. Furthermore, we examined the impact of reference speech length using SMCS on LibriTTS’s unseen speakers. The results are presented in Figure 2.

3.2. Experimental Results

Table 2: Zero-Shot VC evaluation results.

Dataset	VCTK _{unseen→unseen}		LibriTTS _{unseen→unseen}	
	SMOS	MOS	SMOS	MOS
YourTTS	3.61 ± 0.07	3.67 ± 0.07	3.68 ± 0.06	3.68 ± 0.07
GZS-TV	3.74 ± 0.09	3.77 ± 0.05	3.91 ± 0.04	3.79 ± 0.06

Table 3: YourTTS’s experimental results on LibriTTS.

Metric	SMCS	SMOS	MOS
YourTTS	0.856	4.17 ± 0.08	3.79 ± 0.10
GZS-TV	0.885	4.31 ± 0.07	3.90 ± 0.09

The evaluation results in Tables 1 and 2 demonstrate that GZS-TV outperforms StyleSpeech, Meta-StyleSpeech, and YourTTS in almost all metrics. Firstly, GZS-TV exhibits good generalizability by effectively reducing the performance gaps between seen and unseen speakers. These results suggest that GZS-TV can handle out-of-dataset speakers better than other methods. It is noteworthy that YourTTS performs well on VCTK due to its pre-training on this dataset. However, de-

spite using the same speech synthesis backbone as YourTTS, GZS-TV outperforms YourTTS regarding speaker similarity on both seen and unseen speakers, thanks to its disentanglement learning of speaker encoder and timbre transformer. Additionally, cross-dataset speaker adaptation remains challenging, as evidenced by the significant disparities in results on the unseen speakers of LibriTTS and VCTK datasets for all approaches.

Table 4: TTS ablation studies on LibriTTS_{unseen}.

	Setting	SMCS	CSMOS	CMOS
	GZS-TV	0.814	0	0
#1	w/o \mathcal{L}_d	0.799	-0.13	-0.01
#2	w/o \mathcal{L}_{se}	0.801	-0.12	-0.11
#3	extract s from $spec$	0.807	-0.07	-0.05




Figure 3: Visualization of speaker encoder embeddings after PCA dimensionality reduction.

To verify the effectiveness of each module, we conducted ablation studies. As shown in Table 4, removing \mathcal{L}_d (#1) led to a drop in speaker similarity, while removing \mathcal{L}_{se} (#2) and using direct extraction (#3) resulted in reduced speaker similarity and naturalness. We also conducted an ablation study on the VCTK dataset to investigate the effect of (#2) and (#3) on the speaker encoder. The visualization results in Figure 3 demonstrate that both (#2) and (#3) enhance the generalizability of the speaker encoder, reducing confusion and outliers when extracting embeddings for previously unseen speakers.

4. Conclusions

We proposed GZS-TV, a generalizable zero-shot speaker adaptive TTS and VC model. GZS-TV incorporates feature disentanglement in both the speaker encoder and timbre transformer to enhance their generalizability. Additionally, it leverages the representation learning capability of VAE to improve the speaker encoder. Our experiments demonstrate that GZS-TV can synthesize high-quality speech in zero-shot scenarios and significantly reduce the quality gap between seen and unseen speakers. For future research, we plan to explore more effective timbre transformation techniques and investigate strategies to mitigate mispronunciations of TTS in zero-shot scenarios.

5. References

- [1] X. Tan, T. Qin, F. K. Soong, and T. Liu, "A survey on neural speech synthesis," *CoRR*, vol. abs/2106.15561, 2021.
- [2] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions," in *Proc. of ICASSP*, 2018, pp. 4779–4783.
- [3] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *Proc. of ICLR*, 2021.
- [4] A. Gibiansky, S. Ö. Arik, G. F. Diamos, J. Miller, K. Peng, W. Ping, J. Raiman, and Y. Zhou, "Deep voice 2: Multi-speaker neural text-to-speech," in *Proc. of NeurIPS*, 2017, pp. 2962–2970.
- [5] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. A. Kudinov, "Grad-tts: A diffusion probabilistic model for text-to-speech," in *Proc. of ICML*, vol. 139, 2021, pp. 8599–8608.
- [6] J. Kim, J. Kong, and J. Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *Proc. of ICML*, vol. 139, 2021, pp. 5530–5540.
- [7] M. Tamura, T. Masuko, K. Tokuda, and T. Kobayashi, "Speaker adaptation for hmm-based speech synthesis system using MLLR," in *The Third ESCA/COCOSDA Workshop on Speech Synthesis*, 1998, pp. 273–276.
- [8] Y. Yan, X. Tan, B. Li, T. Qin, S. Zhao, Y. Shen, and T. Liu, "Adaspeech 2: Adaptive text to speech with untranscribed data," in *Proc. of ICASSP*, 2021, pp. 6613–6617.
- [9] D. Xin, Y. Saito, S. Takamichi, T. Koriyama, and H. Saruwatari, "Cross-lingual text-to-speech synthesis via domain adaptation and perceptual similarity regression in speaker space," in *Proc. of INTERSPEECH*, 2020, pp. 2947–2951.
- [10] T. Wang, R. Fu, J. Yi, J. Tao, Z. Wen, C. Qiang, and S. Wang, "Prosody and voice factorization for few-shot speaker adaptation in the challenge m2voc 2021," in *Proc. of ICASSP*, 2021, pp. 8603–8607.
- [11] S. Choi, S. Han, D. Kim, and S. Ha, "Attentron: Few-shot text-to-speech utilizing attention-based variable-length embedding," in *Proc. of INTERSPEECH*, 2020, pp. 2007–2011.
- [12] Y. Huang, L. He, W. Wei, W. Gale, J. Li, and Y. Gong, "Using personalized speech synthesis and neural language generator for rapid speaker adaptation," in *Proc. of ICASSP*, 2020, pp. 7399–7403.
- [13] W. Wang, Y. Song, and S. Jha, "Autolv: Automatic lecture video generator," in *Proc. of ICIP*, 2022, pp. 1086–1090.
- [14] M. Chen, X. Tan, B. Li, Y. Liu, T. Qin, S. Zhao, and T. Liu, "Adaspeech: Adaptive text to speech for custom voice," in *Proc. of ICLR*, 2021.
- [15] S. Ö. Arik, J. Chen, K. Peng, W. Ping, and Y. Zhou, "Neural voice cloning with a few samples," in *Proc. of NeurIPS*, 2018, pp. 10 040–10 050.
- [16] Y. Wu, X. Tan, B. Li, L. He, S. Zhao, R. Song, T. Qin, and T. Liu, "Adaspeech 4: Adaptive text to speech in zero-shot scenarios," in *Proc. of INTERSPEECH*, 2022, pp. 2568–2572.
- [17] A. Tjandra, R. Pang, Y. Zhang, and S. Karita, "Unsupervised learning of disentangled speech content and style representation," in *Proc. of INTERSPEECH*, 2021, pp. 4089–4093.
- [18] C. Chien, J. Lin, C. Huang, P. Hsu, and H. Lee, "Investigating on incorporating pretrained and learnable speaker representations for multi-speaker multi-style text-to-speech," in *Proc. of ICASSP*, 2021, pp. 8588–8592.
- [19] 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 *Proc. of ICASSP*, 2020, pp. 6184–6188.
- [20] Y. Zhang, S. Pan, L. He, and Z. Ling, "Learning latent representations for style control and transfer in end-to-end speech synthesis," in *Proc. of ICASSP*, 2019, pp. 6945–6949.
- [21] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. Lopez-Moreno, and Y. Wu, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," in *Proc. of NeurIPS*, 2018, pp. 4485–4495.
- [22] E. Casanova, J. Weber, C. D. Shulby, A. C. Júnior, E. Gölge, and M. A. Ponti, "Yourtts: Towards zero-shot multi-speaker TTS and zero-shot voice conversion for everyone," in *Proc. of ICML*, vol. 162, 2022, pp. 2709–2720.
- [23] D. Min, D. B. Lee, E. Yang, and S. J. Hwang, "Meta-stylespeech: Multi-speaker adaptive text-to-speech generation," in *Proc. of ICML*, vol. 139, 2021, pp. 7748–7759.
- [24] A. Mathur, F. Kawsar, N. Berthouze, and N. D. Lane, "Libri-adapt: a new speech dataset for unsupervised domain adaptation," in *Proc. of ICASSP*, 2020, pp. 7439–7443.
- [25] M. Yi, L. Hou, J. Sun, L. Shang, X. Jiang, Q. Liu, and Z. Ma, "Improved OOD generalization via adversarial training and pre-training," in *Proc. of ICML*, vol. 139, 2021, pp. 11 987–11 997.
- [26] J. Wang, C. Lan, C. Liu, Y. Ouyang, and T. Qin, "Generalizing to unseen domains: A survey on domain generalization," in *Proc. of IJCAI*, 2021, pp. 4627–4635.
- [27] S. Lee, J. Kim, H. Chung, and S. Lee, "Voicemixer: Adversarial voice style mixup," in *Proc. of NeurIPS*, 2021, pp. 294–308.
- [28] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *Proc. of ICLR*, 2014.
- [29] 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," in *Proc. of INTERSPEECH*, 2019, pp. 1526–1530.
- [30] V. Christophe, Y. Junichi, M. Kirsten, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," in *University of Edinburgh. The Centre for Speech Technology Research (CSTR)*, 2016.
- [31] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Proc. of INTERSPEECH*, 2020, pp. 3830–3834.
- [32] D. J. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *Proc. of ICML*, vol. 37, 2015, pp. 1530–1538.
- [33] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using real NVP," in *Proc. of ICLR*, 2017.
- [34] S. Gao, M. Cheng, K. Zhao, X. Zhang, M. Yang, and P. H. S. Torr, "Res2net: A new multi-scale backbone architecture," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 2, pp. 652–662, 2021.
- [35] K. Okabe, T. Koshinaka, and K. Shinoda, "Attentive statistics pooling for deep speaker embedding," in *Proc. of INTERSPEECH*, 2018, pp. 2252–2256.
- [36] Y. Ganin and V. S. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *Proc. of ICML*, vol. 37, 2015, pp. 1180–1189.
- [37] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. of ICLR*, 2019.
- [38] Y. Lei, S. Yang, J. Cong, L. Xie, and D. Su, "Glow-wavegan 2: High-quality zero-shot text-to-speech synthesis and any-to-any voice conversion," in *Proc. of INTERSPEECH*, 2022, pp. 2563–2567.
- [39] J. Cong, S. Yang, L. Xie, and D. Su, "Glow-wavegan: Learning speech representations from gan-based variational auto-encoder for high fidelity flow-based speech synthesis," in *Proc. of INTERSPEECH*, 2021, pp. 2182–2186.