



# DRSpeech: Degradation-Robust Text-to-Speech Synthesis with Frame-Level and Utterance-Level Acoustic Representation Learning

Takaaki Saeki<sup>1,2</sup>, Kentaro Tachibana<sup>1</sup>, and Ryuichi Yamamoto<sup>1</sup>

<sup>1</sup>LINE Corporation, Japan. <sup>2</sup>The University of Tokyo, Japan.

takaaki.saeki@ipc.i.u-tokyo.ac.jp, kentaro.tachibana@linecorp.com

## Abstract

Most text-to-speech (TTS) methods use high-quality speech corpora recorded in a well-designed environment, incurring a high cost for data collection. To solve this problem, existing *noise-robust TTS* methods are intended to use noisy speech corpora as training data. However, they only address either time-invariant or time-variant noises. We propose a *degradation-robust TTS* method, which can be trained on speech corpora that contain both additive noises and environmental distortions. It jointly represents the time-variant additive noises with a frame-level encoder and the time-invariant environmental distortions with an utterance-level encoder. We also propose a regularization method to attain clean environmental embedding that is disentangled from the utterance-dependent information such as linguistic contents and speaker characteristics. Evaluation results show that our method achieved significantly higher-quality synthetic speech than previous methods in the condition including both additive noise and reverberation.

**Index Terms:** text-to-speech synthesis, degradation-robust training, acoustic representation learning

## 1. Introduction

Text-to-speech (TTS) synthesis is intended to artificially synthesize human speech utterances. Because TTS models are generally trained on high-quality speech corpora recorded in a well-developed environment, the data collection is time-consuming and financially expensive. The cost of data collection can be greatly reduced by using data recorded with smartphones [1] or data downloaded from the Internet [2]. However, those speech data contains distortions such as noise and reverberation; it is challenging to train high-quality TTS models using such data. Some studies have been conducted on *noise-robust TTS*: training high-quality TTS models to be robust to time-variant additive noises [3] or time-invariant background noises [4]. However, real speech data includes both additive noises, which exist independently of the speech signal, and environmental distortions, which is multiplied on the speech signal (e.g., reverberation). Therefore, there is a need for a TTS framework that jointly represents these types of distortion.

We propose a *degradation-robust TTS* method, which can be trained on degraded speech that contains both time-variant additive noises and time-invariant environmental distortions. Our method uses a frame-level noise encoder to represent time-variant additive noises and an utterance-level environment encoder to address time-invariant environmental distortions; the two encoders are jointly trained with the TTS model. For the modeling of environmental distortions, we also propose a regularization method to attain clean environmental embedding that is disentangled from the utterance-dependent information such as linguistic contents and speaker characteristics. The main contributions of this work are as follows:

- We propose a degradation-robust TTS method, which jointly addresses additive noises and environmental distortions.
- We propose a regularization method to acquire clean environmental embedding that is disentangled from linguistic contents and speaker information.
- Evaluation results show that our method outperformed previous methods with respect to objective and subjective metrics in the condition including both noise and reverberation.

## 2. Related work

Noise-robust TTS methods, which learn from noisy speech, have been studied. Some previous studies on noise-robust TTS have used a speech enhancement model; for example, a TTS model has been trained with enhanced speech or features [5, 6], and a model trained with clean speakers has been adapted to noisy new speakers [7]. The existing method that is most similar to ours is one in which a TTS model is trained with frame-level noise representation [3]. The aim of our work is to novelly extend that method to represent both time-variant additive noise and time-invariant distortions of degraded speech. Some previous studies have modeled background noise or distortion [4, 8, 9] without an external speech enhancement model. However, they only focus on time-invariant noise, limiting the applicability. Our method uses a style token layer [10] for the utterance-level acoustic encoder. The original paper [10] also reports experiments conducted in the reverberation condition, with the manually selected token representing clean environmental conditions. In contrast, our method introduces a regularization method to automatically obtain clean environmental embedding.

Neural sequence-to-sequence TTS methods [11–14] that use clean speech corpora have been proposed, including FastSpeech 2 [15], which is used in our work. Our method can also be applied to other sequence-to-sequence TTS models [12, 16] by conditioning the frame-level and utterance-level embedding on the intermediate features after the duration prediction.

## 3. Method

Fig. 1 shows an overview of our method. It jointly uses a frame-level noise representation and an utterance-level environmental representation. Our model is based on FastSpeech 2 [15], which consists of a phoneme encoder that encodes the input phoneme embedding into a hidden sequence, a length regulator that predicts and extends the length of the encoded representation, a pitch and energy predictor, and a decoder that outputs acoustic features from an intermediate representation. The model is a multispeaker model, in which the embedded speaker ID is added to the output of the phoneme encoder.

### 3.1. Frame-level noise representation learning

To train the TTS model in a manner that takes time-variant noise into account, we introduce a frame-level noise representation with a noise extractor, as in a previous study [3]. The

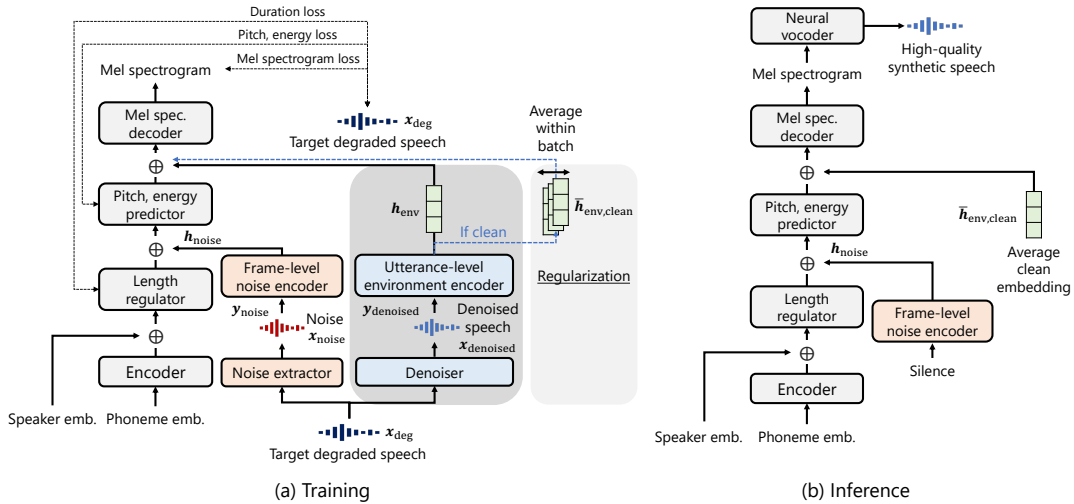


Figure 1: (a) Proposed model is trained on speech corpora that include additive noise and environmental distortions by jointly using frame-level noise representation and utterance-level environmental representation. (b) During inference, model generates high-quality speech by using silence as the input of the noise encoder and average clean environmental embedding as utterance-level representation.

previous study showed that the TTS model using an intermediate frame-level noise representation achieves higher synthetic speech quality than simply using denoised speech feature as the training target. Whereas they use a simple U-Net [17] model as a noise extractor, we use an improved noise extractor in the waveform domain, Conv-TasNet [18], which generalizes to a wider variety of types of additive noise. The noise extractor extracts only additive noise from the target degraded speech, and a noise encoder outputs a hidden noise representation.

As shown in Fig. 1(a), the target degraded waveform  $x_{\text{deg}}$  is input to the noise extractor, which is pre-trained to output a noise waveform from a noisy speech waveform with a large dataset. The output noise waveform  $x_{\text{noise}}$  is then converted into a mel spectrogram  $y_{\text{noise}}$  and input to the frame-level noise encoder. The noise encoder outputs a noise representation  $h_{\text{noise}}$ , which has the same number of frames as the target mel spectrogram and is added to the output of the length regulator. In inference, we use the silence defined as  $x_{\text{noise}}(n) = 0$  (for all  $n$ ) to generate output speech without frame-level additive noise, as shown in Fig. 1(b).

### 3.2. Utterance-level environmental representation learning

In addition to a frame-level noise representation, the proposed method also learns an utterance-level environmental representation. To extract only the environmental conditions from speech—including additive noise and environmental distortion—we introduce a denoiser to remove additive noise. We use a denoiser based on Conv-TasNet, which is pretrained to output a denoised speech waveform from a noisy speech waveform. The target degraded waveform  $x_{\text{deg}}$  is input to the denoiser, which outputs the denoised speech waveform  $x_{\text{denoised}}$ . The mel-spectrogram  $y_{\text{denoised}}$  is obtained from  $x_{\text{denoised}}$  and input to the utterance-level environment encoder to obtain the environmental representation  $h_{\text{env}}$ . We use a style token layer [10] for the utterance-level environment encoder. The environmental embedding  $h_{\text{env}}$  is conditioned on the TTS model as shown in Fig. 1(a).

In the original paper proposing the style token layer [10], they report evaluations of the TTS model trained in reverberation condition. During inference, they manually select a style token corresponding to the clean condition, which requires substantial time and effort. Therefore, during inference of our

method, the average clean embedding  $\bar{h}_{\text{env}, \text{clean}}$  is conditioned on the TTS model, where  $\bar{h}_{\text{env}, \text{clean}}$  is obtained by averaging  $h_{\text{env}}$  for all the training data recorded in clean environmental conditions.

However, simply providing the TTS model with  $h_{\text{env}}$  from each utterance during training and the average clean embedding  $\bar{h}_{\text{env}, \text{clean}}$  during inference does not guarantee that the desired degradation-robust training can be achieved. This is because  $h_{\text{env}}$  is utterance-dependent and entangled with speaker characteristics. Therefore,  $\bar{h}_{\text{env}, \text{clean}}$  does not necessarily represent the clean conditions, and using this embedding can lead to low-quality synthetic speech. We thus propose the regularization method described in Section 3.3 to obtain a disentangled clean environmental embedding  $h_{\text{env}}$ .

### 3.3. Training objective with regularization term

The training objective of our model consists of mean squared error (MSE) loss for duration, pitch, and energy in addition to L1 loss for mel spectrogram. The sum of the above loss functions is denoted by  $\mathcal{L}_{\text{main}}$ , which is the same as the objective used by the original FastSpeech 2 [15].

Furthermore, we propose a regularization method that uses the average clean environmental embedding during training. For the regularization, we introduce subtask learning that uses only speech data with clean environmental conditions. In the subtask learning, the environmental embedding  $h_{\text{env}, \text{clean}}$  estimated from target speech is averaged within the batch. The averaged embedding  $\bar{h}_{\text{env}, \text{clean}}$  is conditioned on the TTS model as described in Section 3.2 and the loss function is calculated in the same manner as in the main task. Let  $\mathcal{L}_{\text{average}}$  be the regularization term defined by the subtask learning. The overall training objective is then defined as:

$$\mathcal{L} = \mathcal{L}_{\text{main}} + \alpha \mathcal{L}_{\text{average}}, \quad (1)$$

where  $\alpha$  is a weighting term and set to 1.0 in our evaluations. With this regularization, the utterance-level encoder can extract only the acoustic aspects shared across multiple utterances (i.e., time-invariant environmental conditions), which are disentangled from the utterance-dependent information such as linguistic contents and speaker characteristics. The dark gray area in Fig. 1(a) depicts the main-task learning to define  $\mathcal{L}_{\text{main}}$ , and the light gray area depicts the subtask learning to define  $\mathcal{L}_{\text{average}}$ .

## 4. Experimental evaluation

### 4.1. Experimental condition

#### 4.1.1. Database and preprocessing

We used VCTK [19], a multi-speaker English speech corpus, for the degraded speech dataset. The sampling frequency was set to 22.05 kHz. As in the previous study [3], we used PNL 100 Nonspeech Sounds [20] for the noise dataset to generate degraded speech data. We chose reverberation in this study as the type of environmental distortion for degraded speech. In previous studies on noise-robust TTS [3, 8], VCTK speakers were divided into two subsets: one for the clean condition and the other for the noise condition. In this study, in order to consider both noise and environmental conditions, the speakers were randomly divided into four sets: 1) Clean, 2) Noise, 3) Reverb, and 4) Noise+Reverb conditions. For the Clean set, we used the original VCTK speech data. In the Noise set, a noise clip was randomly selected from the noise dataset and added to each speech utterance, where the lengths of the noise and speech waveform were matched by repeating or truncating the noise clip. We set the power of the noise signal according to the loudness units relative to full scale (LUFS) [21], with the LUFS sampled randomly according to a uniform distribution between -40 and -32. In the Reverb set, we simulated reverberation using an open-source simulator [22]. We created a room with size  $(l, w, h) = (10 \text{ m}, 7.5 \text{ m}, 3.5 \text{ m})$  and placed a speech source at  $(5 \text{ m}, 3 \text{ m}, 1.6 \text{ m})$  and a monaural microphone at  $(0.5 \text{ m}, 4.0 \text{ m}, 0.5 \text{ m})$ , assuming that the speaker was standing in the room. The  $T_{60}$  was set to 0.2 s, which was between the values corresponding to the weak and strong reverberation conditions in the previous study [16]. In the Noise+Reverb set, we placed a noise source at  $(3 \text{ m}, 7 \text{ m}, 0.2 \text{ m})$  and set the other room configuration parameters as for the Reverb set. We refer to the above dataset as VCTK-degraded.

We randomly selected 512 utterances from VCTK-degraded as the validation set, 512 utterances as the test set, and used the remainder as the training set. When analyzing speech features, the 80-dimensional mel spectrogram was extracted with a frame size of 1024 and a frame shift of 256. The pitch was calculated by the Harvest [23] algorithm, and the duration of the phoneme sequence was extracted using Montreal Forced Aligner [24].

We also created a large-scale noisy speech dataset to train the denoiser and noise extractor. Our dataset was inspired by LibriMix [25], a large-scale dataset for speech separation based on LibriSpeech [26]. LibriSpeech itself is not suitable for TTS because of its low sampling frequency and relatively low speech quality. Therefore, we created a new corpus named LibriTSMix, which is based on LibriTTS [27] and the wham! noise dataset [28] used in LibriMix. We randomly selected a noise clip from the noise dataset and added it to each speech utterance of train-clean-360 and train-clean-100, setting the noise power was set according to a uniform distribution between -38 and -30 LUFS. The sampling frequency of the datasets was set to 22.05 kHz. We trained the denoiser to estimate the clean speech from the noisy speech on the dataset, and we trained the noise extractor to estimate the noise signal.

#### 4.1.2. Model details

We used TTS models based on FastSpeech 2 [15], which has four and six Transformer blocks in the encoder and decoder, respectively. In each Transformer block, the hidden size of the self-attention and feedforward layers was set to 256, with 256-dimension phoneme embedding. The dimension of speaker em-

bedding was set to 256, and the number of speakers was 109 when training the model with VCTK-degraded. For the neural vocoder to generate synthetic speech waveforms from mel spectrograms, we used a pretrained multi-speaker HiFi-GAN<sup>1</sup>.

We used Conv-TasNet for the noise extractor and denoiser. We set the number of repeats to two and we followed the published recipe to set the remaining parameters<sup>2</sup>. Because the model architecture of the frame-level noise encoder was not described in the previous paper [3], we set up a new model based on 1D convolution layers. The model consists of four residual convolution blocks, each of which consists of two 1D-convolution layers with a kernel size of 3, two batch normalization [29] layers, and a skip connection [30]. We used a style token layer [10] for the model architecture of the utterance-level environment encoder. We set the number of heads, the number of tokens, and the dimension of tokens to 8, 10, and 256, respectively.

#### 4.1.3. Training settings

We firstly trained the noise extractor and noise encoder for 300k steps with a batch size of 8. We used the Adam optimizer [31] with a learning rate of 0.001, and used the loss function described in the original paper [18]. After pretraining FastSpeech 2 on LibriTTS, we fixed the weights of only the phoneme encoder and used the weights for the initialization of the proposed model and the other baseline models described in Section 4.2.1. We trained DRSpeech and the other baseline models on VCTK-degraded with a batchsize of 16, which was the maximum batchsize available with a NVIDIA V100 GPU (16 GiB). During training, the silence was provided as the input of the noise encoder in the Clean and Reverb conditions. We used the model with the lowest validation mel spectrogram loss for evaluations. In contrast to the previous study [3], we did not use the ground-truth noise signal, to assume a more realistic situation.

## 4.2. Results

### 4.2.1. Evaluation of DRSpeech and baseline methods

We conducted objective and subjective evaluations to compare DRSpeech with several baseline methods. We prepared a baseline method named Enhancement TTS, in which VCTK-degraded was processed with the speech enhancement model based on Conv-TasNet and then the FastSpeech 2 was trained on the enhanced data. This method is a straightforward approach that corresponds to the previous work [6]. We also prepared a model that considers only the frame-level noise representation described in Section 3.1. This method corresponds to the previous study [3] and is named Noise-robust TTS. It should be noted that this method is not identical to the previous method [3]: Noise-robust TTS uses a more powerful noise extractor, does not use an adversarial connectionist temporal classification loss, and fixes the model parameter of the noise extractor. For a fair evaluation of the proposed method, we kept all conditions the same as those of the proposed method, with the exception of the utterance-level environment encoder. We refer to the proposed method—using all the methods described in Section 3—as DRSpeech. We also compared the natural speech of VCTK (Clean GT) and the target degraded speech of VCTK-degraded (Degraded GT), used as training data for DRSpeech. All the TTS methods were trained on the entire training set of

<sup>1</sup><https://github.com/jik876/hifi-gan>

<sup>2</sup><https://github.com/asteroid-team/asteroid/blob/v0.5.3/egs/librimix/ConvTasNet/local/conf.yml>

Table 1: Results of objective evaluation for all methods

	Clean		Noise		Reverb		Noise+Reverb	
	MCD	Log F0 RMSE	MCD	Log F0 RMSE	MCD	Log F0 RMSE	MCD	Log F0 RMSE
Enhancement TTS [6]	9.04	10.12	9.16	9.86	9.41	10.34	10.39	10.33
Noise-robust TTS [3]	<b>7.74</b>	<b>8.00</b>	<b>7.99</b>	<b>8.56</b>	8.64	8.51	9.65	9.58
DRSpeech	7.94	8.24	8.19	8.79	<b>8.25</b>	<b>8.09</b>	<b>9.18</b>	<b>9.11</b>

Table 2: Results of MOS evaluation test for all methods

(a) Without reverberation		
	MOS	
	Clean	Noise
Clean GT	4.14 ± 0.14	4.77 ± 0.08
Degraded GT	-	2.78 ± 0.18
Enhancement TTS [6]	3.62 ± 0.16	2.85 ± 0.16
Noise-robust TTS [3]	3.53 ± 0.16	2.88 ± 0.15
DRSpeech	<b>3.77 ± 0.15</b>	<b>3.30 ± 0.15</b>

(b) With reverberation		
	MOS	
	Reverb	Noise+Reverb
Clean GT	4.62 ± 0.11	4.71 ± 0.09
Degraded GT	3.24 ± 0.17	2.03 ± 0.15
Enhancement TTS [6]	2.39 ± 0.15	2.05 ± 0.15
Noise-robust TTS [3]	2.47 ± 0.16	2.23 ± 0.16
DRSpeech	<b>2.80 ± 0.16</b>	<b>2.56 ± 0.16</b>

VCTK-degraded and evaluated separately for the four conditions (Clean, Noise, Reverb, and Noise+Reverb) in the test set. Audio samples are publicly available<sup>3</sup>.

We first conducted objective evaluations. Two metrics—mel cepstral distortion (MCD) [32] and log F0 root mean squared error (RMSE)—were calculated for all the TTS methods. Table 1 lists the results. We observe that Noise-robust TTS and DRSpeech, which learn to take acoustic representation into account, performed better than Enhancement TTS in all conditions. In addition, the MCD and log F0 RMSE of Noise-robust TTS are both about 0.2 less than those of DRSpeech in the Clean and Noise conditions. In the Reverb and Noise+Reverb conditions, DRSpeech achieved scores of 0.4–0.5 better than those of Noise-robust TTS.

Second, we conducted a mean opinion score (MOS) test to evaluate the naturalness of synthetic speech. Nine utterances, all by different speakers, were sampled from each of the four degradation conditions in the test set. Twenty evaluators rated the naturalness on a scale of 1 to 5. Table 2 lists the results. DRSpeech achieved the highest MOS in all the conditions (Clean, Noise, Reverb, and Noise+Reverb). In particular, in the Noise+Reverb condition, DRSpeech achieved a MOS of 0.5 greater than that of Enhancement TTS. In the Noise condition, the average MOS of Noise-robust TTS was greater than that of Degraded GT and Enhancement TTS; this is consistent with the result of the previous study [3]. Moreover, DRSpeech outperformed Noise-robust TTS in the Noise condition. Because a noise is added to the entire speech utterance, DRSpeech, which considers global conditions, may be expected to achieve better results than Noise-robust TTS.

<sup>3</sup>[https://takaaki-saeki.github.io/drspeech\\_demo](https://takaaki-saeki.github.io/drspeech_demo)

Table 3: CMOS of DRSpeech with and without regularization

	CMOS			
	Clean	Noise	Reverb	Noise+Reverb
DRSpeech	0	0	0	0
w/o Regularization	-0.207	-0.177	-0.205	0.005

#### 4.2.2. Evaluation of regularization

We conducted a CMOS (Comparison MOS) test to evaluate the effectiveness of the regularization method separately from the utterance-level environment encoder. In this evaluation, two synthetic speech samples from DRSpeech—with and without the regularization method—were played in the random order. Twenty evaluators rated the naturalness of the second speech utterance relative to the first one on a 7-point scale from -3 to +3. Ten pairs with the same utterance text were sampled from the test sets, all with different speakers. Table 3 lists the results.

DRSpeech with the regularization outperformed DRSpeech without the regularization by about 0.2 in the Clean, Noise, and Reverb conditions. However, there was only 0.005 difference between the two methods in the Noise+Reverb condition.

### 4.3. Analysis

In the objective evaluation, Noise-robust TTS outperformed DRSpeech under the Clean and Noise conditions. In the subjective evaluation of Section 4.2.1, however, DRSpeech outperformed Noise-robust TTS in both conditions. A discrepancy between objective and subjective evaluation results has often been observed in other studies [33, 34]. From the standpoint that subjective metrics are more reliable than objective metrics, we can conclude that DRSpeech achieves higher-quality synthetic speech than the Noise-robust TTS. Furthermore, DRSpeech outperformed the baseline methods in objective and subjective evaluations in the Reverb and Noise+Reverb conditions, indicating the effectiveness of DRSpeech in the reverberant conditions.

The evaluation of Section 4.2.2 confirms the effectiveness of the regularization method in the Clean, Noise, and Reverb conditions, although there was no significant difference in the Noise+Reverb condition. This means that the improvement achieved by the utterance-level encoder was more dominant than that achieved by the regularization method in the presence of both noise and reverberation. Further improvement of the regularization method is required in future research.

## 5. Conclusions

We proposed a degradation-robust TTS method with a frame-level noise representation and an utterance-level environmental representation. Our method also uses a regularization method to obtain utterance-independent clean environmental embedding during training. Experimental results showed that our method outperformed existing noise-robust TTS methods with respect to both objective and subjective measures in conditions including reverberation. Our future work includes conducting experiments with a wider variety of types of distortion.

## 6. References

- [1] A. A. Raza, A. Athar, S. Randhawa, Z. Tariq, M. B. Saleem, H. B. Zia, U. Saif, and R. Rosenfeld, "Rapid collection of spontaneous speech corpora using telephonic community forums," in *Proc. INTERSPEECH*, 2018, pp. 1021–1025.
- [2] S. Takamichi, L. Kurzinger, T. Saeki, S. Shiota, and S. Watanabe, "Jtubespeech: corpus of japanese speech collected from youtube for speech recognition and speaker verification," *arXiv preprint arXiv:2112.09323*, 2021.
- [3] C. Zhang, Y. Ren, X. Tan, J. Liu, K. j. Zhang, T. Qin, S. Zhao, and T.-Y. Liu, "Denoispeech: Denoising text to speech with frame-level noise modeling," *Proc. ICASSP*, pp. 7063–7067, 2021.
- [4] W.-N. Hsu, Y. Zhang, R. Weiss, H. Zen, Y. Wu, Y. Cao, and Y. Wang, "Hierarchical generative modeling for controllable speech synthesis," in *Proc. ICLR*, 2019.
- [5] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Investigating rnn-based speech enhancement methods for noise-robust text-to-speech," in *Proc. SSW*, 2016, pp. 159–165.
- [6] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks," in *Proc. INTERSPEECH*, 2016, pp. 352–356.
- [7] D. Dai, L. Chen, Y. Wang, M. Wang, R. Xia, X. Song, Z. Wu, and Y. Wang, "Noise robust tts for low resource speakers using pre-trained model and speech enhancement," *arXiv preprint arXiv:2005.12531*, 2020.
- [8] 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 *Proc. ICASSP*, 2019, pp. 5901–5905.
- [9] D. Tan, G. Zhang, and T. Lee, "Environment aware text-to-speech synthesis," *arXiv preprint arXiv:2110.03887*, 2021.
- [10] Y. Wang, D. Stanton, Y. Zhang, R. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, "Style Tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," *arXiv*, vol. abs/1803.09017, 2018.
- [11] Y. Wang, R. J. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomyriannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards end-to-end speech synthesis," *arXiv*, vol. abs/1609.03499, 2017.
- [12] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *Proc. ICASSP*, Calgary, Canada, Apr. 2018, pp. 4779–4783.
- [13] 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 *Proc. NeurIPS*, 2019.
- [14] J. Kim, J. Kong, and J. Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *Proc. ICML*, 2021, pp. 5530–5540.
- [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," in *Proc. ICLR*, 2021.
- [16] C. Kim, A. Misra, K. K. Chin, T. Hughes, A. Narayanan, T. N. Sainath, and M. Bacchiani, "Generation of large-scale simulated utterances in virtual rooms to train deep-neural networks for far-field speech recognition in google home," in *Proc. INTERSPEECH*, 2017, pp. 379–383.
- [17] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI*, 2015.
- [18] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, pp. 1256–1266, 2019.
- [19] J. Yamagishi, C. Veaux, and K. MacDonald, "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)," 2019.
- [20] G. Hu and D. Wang, "A tandem algorithm for pitch estimation and voiced speech segregation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 8, pp. 2067–2079, 2010.
- [21] ITU-R, "Recommendation ITU-R BS.1770-4: Algorithms to measure audio programme loudness and true-peak audio level," 2015.
- [22] R. Scheibler, E. Bezzam, and I. Dokmanić, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *Proc. ICASSP*, 2018, pp. 351–355.
- [23] M. Morise, "Harvest: A high-performance fundamental frequency estimator from speech signals," in *Proc. INTERSPEECH*, 2017, pp. 2321–2325.
- [24] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldii," in *Proc. INTERSPEECH*, 2017, pp. 498–502.
- [25] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "Librimix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [26] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *Proc. ICASSP*, South Brisbane, Australia, Apr. 2015, pp. 5206–5210.
- [27] 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. INTERSPEECH*, 2019, pp. 1526–1530.
- [28] G. W. and J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. L. Roux, "WHAM!: Extending speech separation to noisy environments," in *Proc. INTERSPEECH*, 2019, pp. 1368–1372.
- [29] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML*, 2015, p. 448–456.
- [30] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, Las Vegas, U.S.A., June 2016, pp. 770–778.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.
- [32] T. Fukada, K. Tokuda, T. Kobayashi, and S. Imai, "An adaptive algorithm for mel-cepstral analysis of speech," in *Proc. ICASSP*, San Francisco, U. S. A., Mar 1992, pp. 137–140.
- [33] R. J. Weiss, R. J. Skerry-Ryan, E. Battenberg, S. Mariooryad, and D. P. Kingma, "Wave-Tacotron: Spectrogram-free end-to-end text-to-speech synthesis," *Proc. ICASSP*, pp. 5679–5683, 2021.
- [34] T. Hayashi, R. Yamamoto, T. Yoshimura, P. Wu, J. Shi, T. Saeki, Y. Ju, Y. Yasuda, S. Takamichi, and S. Watanabe, "ESPnet2-TTS: Extending the edge of TTS research," *arXiv preprint arXiv:2110.07840*, 2021.