

MIIPHER: A ROBUST SPEECH RESTORATION MODEL INTEGRATING SELF-SUPERVISED SPEECH AND TEXT REPRESENTATIONS

Yuma Koizumi¹, Heiga Zen¹, Shigeki Karita¹, Yifan Ding¹, Kohei Yatabe², Nobuyuki Morioka¹,
Yu Zhang³, Wei Han³, Ankur Bapna³, Michiel Bacchiani¹

¹ Google, Japan

² Tokyo University of Agriculture & Technology, Japan

³ Google, USA

ABSTRACT

Speech restoration (SR) is a task of converting degraded speech signals into high-quality ones. In this study, we propose a robust SR model called *Miipher*, and apply *Miipher* to a new SR application: increasing the amount of high-quality training data for speech generation by converting speech samples collected from the Web to studio-quality. To make our SR model robust against various degradation, we use (i) a speech representation extracted from w2v-BERT for the input feature, and (ii) a text representation extracted from transcripts via PnG-BERT as a linguistic conditioning feature. Experiments show that *Miipher* (i) is robust against various audio degradation and (ii) enable us to train a high-quality text-to-speech (TTS) model from restored speech samples collected from the Web. Audio samples are available at our demo page: google.github.io/df-conformer/miipher/.

Index Terms— Speech restoration, speech enhancement, text-to-speech, self-supervised learning

1. INTRODUCTION

Speech restoration (SR) is a task of converting degraded speech signals into high-quality speech signals [1–5]. It is a comprehensive task including enhancement [6], dereverberation [7–9], declipping [10], and super-resolution [11, 12]. SR has been addressed to improve speech intelligibility for helping human listening, and its quality has evolved over the past few years with the evolution of the speech generative models [13–15]. Parametric resynthesis [1]-based methods [2–4] and diffusion model-based methods [5] have been proposed. They can convert lecture and historical speech into high-quality speech as if these were recorded in a studio.

By converting speech samples in-the-wild to studio-quality using advanced SR methods, we are trying to increase the amount of high-quality training data for speech generation. The performance of deep learning depends on both size and quality of the training dataset [16–18]. As the cost for collecting studio-recorded samples is expensive, recent research has tried to alleviate the requirement for dataset by developing specific schemes, such as self-supervised learning (SSL) [19–21], and using (low-quality) speech samples from the Web [22, 23]. However, unlike automatic speech recognition (ASR), training of text-to-speech (TTS) models using samples in-the-wild is still a challenging problem [24, 25] because the quality of speech generation is directly affected by that of the training samples.

The main challenge for the application of SR in this study is the robustness aspect. If SR fails to restore speech samples and produces artifacts, they will negatively affect the subsequent training of speech generative models. Thus, we must reduce the number of restoration failed samples as much as possible.

In this study, we propose *Miipher*¹, a robust parametric re-synthesis SR model for restoring speech samples in-the-wild. Based on our preliminary investigation, we especially focus on the following two difficult degradations where SR frequently fails:

Phoneme masking: Speech signals are sometimes masked by noise and/or reverberation. It is difficult to discriminate speech and noise without having additional information.

Phoneme deletion: Important frequency components of some phonemes could be missing from the signal due to codecs and/or down-sampling. If a noisy sample lacks a phoneme, it introduces an unrecoverable error.

To solve these issues, we introduce the following two techniques:

SSL features domain cleaning: For the input feature, instead of a log-mel spectrogram used in conventional methods [1, 4], we use a speech representation extracted from w2v-BERT [20], an SSL model trained on degraded speech samples. As it improves ASR performance, we expect its effectiveness on making SR models robust against speech degradation.

Transcript conditioning: We consider the deleted phoneme reconstruction problem as a text-conditioned speech inpainting [26]. As a linguistic conditioning feature, we use a text representation extracted from PnG-BERT [27], a text SSL model for TTS use.

Through experiments, we show that (i) *Miipher* is robust against various audio degradations, and (ii) applying *Miipher* on a large scale speech dataset collected from the Web enables us to train a TTS model whose quality is on par with one that is trained on a dataset designed for TTS training. Audio samples of the restored and TTS generated samples are available at our demo page².

2. SPEECH RESTORATION MODEL

2.1. Model overview

Let the T -sample time-domain signal $\mathbf{x} \in \mathbb{R}^T$ be a degraded signal of an original signal $\mathbf{s} \in \mathbb{R}^T$. The goal of SR is to recover \mathbf{s} from \mathbf{x} as $\mathbf{y} \in \mathbb{R}^T$. In this study, the sampling rate of \mathbf{s} and \mathbf{y} is 24 kHz, because 24 kHz sampling signals are often used in speech generation tasks such as TTS [14].

To achieve this goal, we use a parametric re-synthesis framework [1] as shown in Fig. 1 (a): a noisy feature is cleaned by a feature cleaner network, then the restored waveform is synthesized using a neural vocoder. To make our SR model robust against various speech degradation, we used a speech representation extracted

¹Miipher: multiple features integrated speech restoration

²google.github.io/df-conformer/miipher/

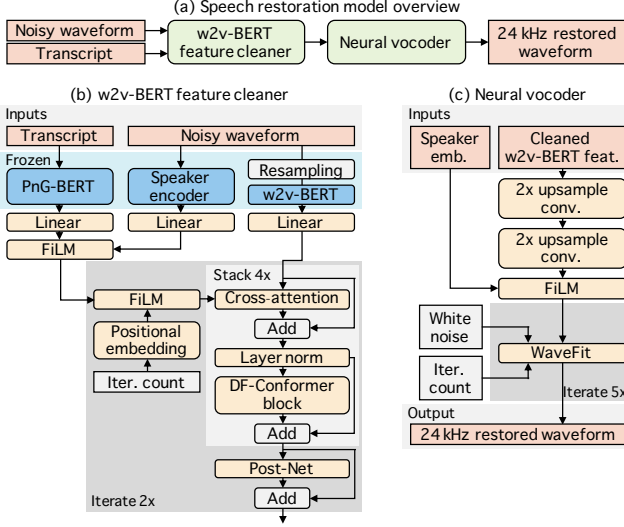


Figure 1: (a) Overview of Miipher. (b) w2v-BERT and PnG-BERT are used as feature extractors, and a DF-Conformer-based feature cleaner predicts clean w2v-BERT features. Then, (c) a WaveFit synthesizes a restored waveform from cleaned w2v-BERT features.

from w2v-BERT [20] for the input feature of SR, and linguistic features conditioning extracted by PnG-BERT [27] from the transcript.

2.2. Feature cleaner

The overview of the feature cleaner is shown in Fig. 1 (b). Our feature cleaner predicts w2v-BERT features of the clean speech $\mathbf{S} \in \mathbb{R}^{K \times D}$ from that of degraded speech $\mathbf{X} \in \mathbb{R}^{K \times D}$ as $\hat{\mathbf{S}} \in \mathbb{R}^{K \times D}$. Here, K is the number of time-frames and D is the dimension of w2v-BERT feature. As a conditioning of SR, we use linguistic features extracted from a transcript via PnG-BERT $\mathbf{E} \in \mathbb{R}^{M \times W}$ where M is the number of tokens and W is the dimension of PnG-BERT features. In addition, since SSL features often lose speaker information [28], we use a speaker embedding $\mathbf{d} \in \mathbb{R}^Q$ extracted from \mathbf{x} as speaker conditioning. Note that since w2v-BERT is trained from 16 kHz audio, we resample \mathbf{x} at 16 kHz for extracting \mathbf{X} .

First, the feature dimension of \mathbf{X} , \mathbf{d} and \mathbf{E} are aligned to D_b by linear layers, where D_b is the input dimension of the DF-Conformer block. Then, \mathbf{d} is combined to \mathbf{E} using a CNN-based simple feature-wise linear modulation (FiLM) layer [29]. Here a FiLM layer is defined as $\text{FiLM}(\mathbf{A}, \mathbf{b}) = \text{CNN}_2(\text{LeakyReLU}(\text{CNN}_1(\mathbf{A})) + \mathbf{b})$ where $\mathbf{A} \in \mathbb{R}^{K \times D}$ and $\mathbf{b} \in \mathbb{R}^Q$. CNN_1 changes \mathbf{A} 's 2nd dimension D to Q . Then, \mathbf{b} is added to the CNN_1 output while broadcasting the 1st dimension. Here, the slope parameter of LeakyReLU is 0.1. Finally, CNN_2 returns the 2nd dimension to D . Both CNNs are 1D-CNN layer with a kernel-size 3 and stride 1.

Next, \mathbf{S} is predicted from \mathbf{X} by using a DF-Conformer [30]-based network. We build our feature cleaner by stacking N blocks, each consisting of a cross attention, a layer norm and a DF-Conformer-block. In each block, the combined conditioning feature is combined to the input w2v-BERT features using a cross attention layer. Then, layer norm is applied to the output, and finally, a DF-Conformer block [30] with $\text{mod}(n, 2)$ dilation factor is applied where $n \in \{0, 1, \dots, N-1\}$ is the block number. After refining the input w2v-BERT features, we applied the 5-layer convolutional

Post-Net proposed in Tacotron2 [14] which predicts a residual to add to the prediction to improve the overall reconstruction.

We adopt the iterative feature refinement into our feature cleaner since a fixed-point iteration-like processing achieved high-fidelity waveform generation [18]. Specifically, we iterate twice the entire feature cleaning process consisting of the feature cleaner and the Post-Net, where the parameters of the layers are shared. As with WaveGrad [31] and WaveFit [18], the number of iterations is embedded using a positional embedding layer, and the embedded iteration count is mixed into text and speaker embedding using a FiLM layer. Note that it is not strictly fixed-point iteration since the function being applied is slightly different at each iteration.

We used a combined loss function of the mean-absolute-error, the mean-squared-error, and a spectral convergence loss [32]-like loss as $\mathcal{L} = \|\mathbf{S} - \hat{\mathbf{S}}\|_1 + \|\mathbf{S} - \hat{\mathbf{S}}\|_2^2 + \|\mathbf{S} - \hat{\mathbf{S}}\|_2^2 / \|\mathbf{S}\|_2^2$, where $\|\cdot\|_p$ is the vector p -norm given by $\|\mathbf{S}\|_p = (\sum_k \sum_d |S_{k,d}|^p)^{1/p}$. This loss value is calculated before and after the Post-Net like in Tacotron2 [14], and calculated for all iterations like in WaveFit [18].

2.3. Neural vocoder

As shown in Fig. 1 (c), we use a WaveFit neural vocoder [18] to synthesize \mathbf{y} from $\hat{\mathbf{S}} \in \mathbb{R}^{K \times D}$. The original WaveFit synthesizes waveforms from log-Mel spectrograms; here we adopt WaveFit to synthesize waveforms from w2v-BERT features.

First, the frame rate of w2v-BERT feature and log-mel spectrogram does not match, hence $\hat{\mathbf{S}}$ is 4x upsampled along the time axis. Since the network architecture of WaveFit consists of a CNN-based U-Net [31], a deep upsampling network is required to synthesize a waveform from low frame-rate features. To keep the core network size close to the original one, we applied two upsampling CNN layers: each layer consists of a transposed convolution with ReLU, channel size 1 where the kernel- and stride-sizes are 4 and 2, respectively. Then, the conditioning speaker embedding \mathbf{d} is combined with the upsampled w2v-BERT feature $\hat{\mathbf{S}}' \in \mathbb{R}^{4K \times D}$ using a FiLM layer as $\text{FiLM}(\hat{\mathbf{S}}', \mathbf{d})$.

Finally, the speech waveform is synthesized by WaveFit. Note that we used white noise as the initial noise waveform instead of the SpecGrad noise [33] used in the original WaveFit [18] because we cannot know the spectral envelope of \mathbf{s} from w2v-BERT features. In addition, since the maximum amplitude of \mathbf{s} cannot be known from w2v-BERT features, the gain normalization of WaveFit is replaced to $\mathcal{G}(\mathbf{y}) = \lambda \cdot \mathbf{y} / \max(\text{abs}(\mathbf{y}))$ where $\lambda = 0.9$ and $\text{abs}(\cdot)$ returns element-wise absolute values of the input vector.

For the loss function, in addition to the original adversarial loss function proposed in WaveFit [18], we used the multi-period discriminator (MPD) [34]. We used the same MPD hyper-parameters used in [34] except for additionally using three more prime-number periods (13, 17, 19) because the sampling rate of \mathbf{y} is 24 kHz. In addition, the parameters of the STFT loss was changed to the same as those of Parallel WaveGAN settings [32].

2.4. Model training

Frozen feature extraction models: We used the “w2v-BERT XL” model, where $D = 1024$, trained on 60k hours of English speech samples from the LibriVox repository. We used the 8th layer Conformer outputs without quantization as w2v-BERT features. For the PnG-BERT model, we used a $W = 512$ PnG-BERT model [27] pre-trained on a plain text corpus mined from Wikipedia, containing 131M English sentences. To extract speaker embedding, we

Table 1: Random codec parameters. Codec was randomly selected based on the probability listed in the “Prob.” column. Codec bitrate was randomly selected from those listed in the “Bit rate” column.

Codec	Prob.	Bit rate
MP3	0.5	16k, 32k, 64k, 128k
Vorbis	0.075	32k, 48k, 64k
A-law	0.025	64k
AMR-WB	0.025	6.6k, 8.85k, 12.65k, 14.25k, 15.85k
OPUS	0.375	18.25k, 19.85k, 23.05k, 23.85k
		8k, 16k, 32k, 64k, 128k

used a streaming Conformer-based speaker encoding model consists of 12 Conformer [35] layers each of $Q = 256$, followed by an attentive temporal pooling module [36]. The model was trained on the dataset described in [37] while minimizing the generalized end-to-end extended-set softmax (GE2E-XS) loss [38].

Trainable model parameters: We stacked $N = 4$ DF-Conformer blocks [30] where the input and hidden dimensions of the multi-head attention are $D_b = 128$ and 512, respectively. Other hyperparameters were the same as those of the original paper [30]. For the cross-attention layer before the DF-Conformer block, the input and hidden dimensions were 128 and 512, respectively. The hyperparameters of the Post-Net are the same as those of Tacotron2 [14]. For the neural vocoder, we used WaveFit-5 with the same hyperparameters described in [18] except for the upsampling factors of the five upsampling blocks (UBlock) are (5, 4, 2, 2, 2) because the frame rate of the upsampled w2v-BERT feature is 100 [frame/sec]. To accelerate training, we randomly picked up 15 frames (0.6 seconds) from \mathbf{X} as input.

Training dataset: We trained the proposed model with a proprietary dataset that contains 2,680 hours of noisy and studio-quality speech pairs. The target speech dataset contains 670 hours of studio-recorded Australia, United Kingdom, India, Nigeria, and United States English at 24 kHz sampling. For the noise dataset, we used the TAU Urban Audio-Visual Scenes 2021 dataset [39], internally collected noise snippets that simulate conditions like cafe, kitchen, and cars, and noise and music sources. The noisy utterances were generated by mixing randomly selected speech and noise samples from these datasets with signal-to-noise ratio (SNR) from 5 dB to 30 dB. In addition, we augmented the noisy dataset with 4 patterns depending on the presence or absence of reverberation and codec artifacts. A room impulse response (RIR) for each sample was generated by a stochastic RIR generator using the image method [40]. Its parameters were drawn from the following uniform distributions \mathcal{U} : the reverberation times (RT60) ranging from $\mathcal{U}(0.2, 0.5)$ [sec], the length of one side for x- and y-axis (width) room was $\mathcal{U}(2.0, 10.0)$ [m], and that for z-axis (height) $\mathcal{U}(2.0, 5.0)$ [m]. For simulating codec artifacts, we randomly applied one of MP3, Vorbis, A-law, Adaptive Multi-Rate Wideband (AMR-WB), and OPUS with a random bit-rate. The codec parameters were shown in Table 1 which were decided with reference to [5].

3. EXPERIMENTS

3.1. Experimental setup

We evaluated the effectiveness of w2v-BERT and PnG-BERT while comparing the sound quality between the studio-recorded original speech and restored samples from artificially contaminated samples.

Table 2: Results on the synthesized test dataset with their 95% confidence intervals. “L”, “W” and “S” means log-mel, w2v-BERT, transcript and speaker embedding conditioning, respectively.

Feat.	Method	MOS (\uparrow)	WER (\downarrow)	SPK (\uparrow)
-	Clean	4.69 ± 0.06	13.7	N/A
-	Noisy	3.28 ± 0.13	15.1	0.703
L	Full	3.08 ± 0.12	19.4	0.767
	w/o T	2.16 ± 0.13	39.6	0.698
	w/o S	3.17 ± 0.08	19.5	0.711
	w/o T & S	1.40 ± 0.09	61.4	0.496
W	Full (ours)	4.54 ± 0.08	13.5	0.727
	w/o T	4.48 ± 0.08	16.4	0.735
	w/o S	4.39 ± 0.08	14.0	0.646
	w/o T & S	4.26 ± 0.09	17.2	0.637

Test dataset: Test dataset was generated by contaminating the clean data with the same way for generating the training dataset described in Sec. 2.4. The clean speech samples were studio-recorded US-English containing 534 utterances spoken by 1 female and 1 male. There was no overlap of this data with the training set in terms of the utterances themselves or the speakers. The additive noise dataset was the DEMAND noise dataset [41]. We applied reverberation to all clean speech samples using the same stochastic RIR generator with the same parameters as the training dataset. After that, the noisy samples were generated by mixing randomly selected reverberant speech and noise samples from these datasets with SNR from 5 dB to 30 dB. Finally, we simulated codec artifacts to the noisy samples with the same parameters with the training dataset.

Comparison methods: The main novelties of the proposed method are to use (i) w2v-BERT features instead of log-mel spectrogram, (ii) PnG-BERT for text conditioning, and (iii) a speaker embedding for speaker conditioning. Therefore, we compared eight models with and without these three patterns: i.e. $\{\text{w2v-BERT, log-mel}\} \times \{\text{full, w/o PnG-BERT, w/o speaker embedding, w/o PnG-BERT \& speaker embedding}\}$. For log-mel models, we used 128-dimensional log-mel spectrograms (50 ms Hann window, 12.5 ms frame shift, 2048-point FFT, and 20 Hz and 12 kHz lower and upper frequency cutoffs, respectively). As the neural vocoder for the log-mel models, we used the original WaveFit-5 [18] except for using the MPD [34] with the same parameters as w2v-BERT models. Other parameters were the same among the all comparison models. We trained the feature-cleaner 400k steps. Individually, we pre-trained the WaveFit 200k steps to reconstruct waveform from clean features, after that, we fine-tuned the WaveFit 100k steps to synthesize the clean waveform from the predicted features by the trained feature cleaner.

3.2. Results

To evaluate subjective quality, we rated speech quality through mean-opinion-score (MOS). The scale of MOS was a 5-point scale (1: Bad, 2: Poor, 3: Fair, 4: Good, 5: Excellent) with rating increments of 0.5. Test stimuli were randomly chosen and each stimulus was evaluated by one subject. Each subject was allowed to evaluate up to six stimuli, that is, 89 subjects participated in this experiment to evaluate 534 samples in each condition. The subjects were paid native English speakers in the United States. They were requested to use headphones in a quiet room. In addition, in order to confirm

Table 3: MOS scores with their 95% confidence intervals of TTS outputs trained on each restored dataset.

Training dataset	Speaker ID	MOS (\uparrow)
Common Voice	-	3.63 \pm 0.21
LibriVox	4788	4.38 \pm 0.15
	10801	4.15 \pm 0.17
	5968	4.11 \pm 0.19
	8138	4.05 \pm 0.18

whether the text-content and speaker identity in the restored speech samples are retained, we evaluated the word error rate (WER) of an ASR experiment and computed a speaker similarity measure. The WER was computed between the reference transcripts and the output of the “Pre-trained Conformer XXL” ASR model [42]. Note that the ASR model was trained on noisy speech and transcripts normalized by a different text-normalizer. Therefore, we only check trends in WER differences since absolute value of WER for our evaluation dataset will be high. To evaluate speaker similarity, we computed a cosine similarity (SPK) of the speaker embedding between the clean and restored samples [43, 44].

Table 2 shows the results. The proposed method (w2v-BERT, full) achieved the best MOS among all SR comparison models. In addition, the MOS of the proposed model was only 0.15 points behind that of the studio-recorded clean signals. By comparing MOS of w2v-BERT models and log-mel models, the scores of w2v-BERT models were significantly higher. This result indicates that the use of w2v-BERT features has a significant effect on improving SR performance. The WER of “W w/o T” was 2.9 point behind of the proposed method, and the WER of the proposed method was comparable to that of the clean signal. This shows that the use of PnG-BERT features extracted from a transcript is effective in preserving text content of the restored speech. The SPK score of “W w/o S” model was significantly lower than that of the proposed method. This result indicates the use of a speaker embedding is effective for keeping speaker characteristics. Audio samples are available at our demo page².

4. TTS MODEL TRAINING WITH NON-TTS DATASET

Thanks to crowd-sourcing projects for creating license-free speech corpus, a huge amount of English text-speech pair data is publicly available. However, since these samples are not studio-recorded quality, their use-cases are mainly limited to ASR. To increase the value of such public speech samples, we restored the samples in Common Voice [22] and LibriVox datasets using Miipher. Then, we trained TTS models using the restored datasets. The details of this experiment is as follows.

TTS model: The TTS model consisted of an unsupervised Non-Attentive Tacotron (UNAT) with a fine-grained variational auto-encoder [45] and a WaveRNN neural vocoder [46]. For UNAT, we used the same hyper-parameters and training parameters listed in the original paper [45]. The WaveRNN [46] consisted of a single long short-term memory layer with 512 hidden units, 5 convolutional layers with 512 channels as the conditioning stack to process the mel-spectrogram features, and a 10-component mixture of logistic distributions as its output layer. We trained UNAT and WaveRNN for 100k and 400k steps and concatenated them without fine-tuning.

Common Voice: We restored samples in the version 5.1 (June 22, 2020) snapshot with approximately 1,500 hours. We removed speech samples of less than 2 seconds or more than 15 seconds. The training dataset consisted of 685 hours speech samples. Note that since Common Voice is prohibited from identifying the speaker, the TTS model was trained without using speaker IDs and trained as a single-speaker TTS model.

LibriVox: We collected over 25,000 hours of speech samples (Jan. 7, 2020). There was no overlap of this data with LibriTTS [47]. We aligned the long-form audio recordings with the corresponding texts, and split them into sentence-level segments as the same way of building LibriTTS [47]. We removed speech samples of less than 2 seconds or more than 15 seconds. The training dataset consisted of 13,270 hours speech samples spoken by 4,000 speakers. We trained a multi-speaker TTS model.

We synthesized waveforms with the sentences from the LJSpeech [48] test split. For the LibriVox multi-speaker model, we used 4 speakers (two female and two male) where the female and male reader IDs were (14788, 10801) and (5968, 8168), respectively.

We evaluated the subjective sound quality using MOS. To give baseline MOS on LJSpeech test split sentences, we also trained the same TTS model using Miipher applied LJSpeech [48]. The MOS of human spoken and TTS generated speech samples of LJSpeech were 4.33 and 4.36 respectively.

Table 3 shows the MOS of each model. Even though the LibriVox are not datasets designed for TTS like LJSpeech, and the training of multi-speaker TTS is difficult compared with single-speaker TTS, TTS models trained using restored LibriVox achieved MOS that were almost equal to those using restored LJSpeech. Furthermore, in Common Voice, despite we trained a single-speaker TTS model using multiple speakers samples, the MOS of the TTS generated samples was 3.63. We complement that we could not train the same WaveRNN neural vocoder using the original Common Voice dataset: the training failed to converge due to the noise in the target speech. These results indicate that applying Miipher on a large scale speech dataset collected from the Web enables us to train high-quality TTS models. Audio samples of the restored and TTS generated samples are available at our demo page².

5. CONCLUSIONS

In this study, we proposed Miipher, a robust SR model based on the parametric re-synthesis framework for restoring speech samples in-the-wild. Based on our preliminary investigation, we especially focused on two difficult degradation where SR frequently fails: phoneme masking and deletion. To solve these issues, for the input features, we used speech and text representations extracted from w2v-BERT [20] and PnG-BERT [27], respectively. Through experiments, we show that Miipher was robust against various audio degradation, could restore public speech samples from the Web, and enabled us to train high-quality TTS models from the restored samples. From these results, we conclude that Miipher can increase the value of speech samples in-the-wild by improving the speech quality as the training data for speech generation tasks.

In the future, we will investigate the use of transcripts predicted using ASR to extract linguistic features and restore speech samples for which transcriptions do not exist. Also, we will develop multilingual Miipher in order to train high-quality TTS models for low-resource languages. Training of SSL models to obtain speech representation from 24 kHz sampling signals is also a future work.

6. REFERENCES

- [1] S. Maiti and M. I. Mandel, "Parametric resynthesis with neural vocoders," in *WASPAA*, 2019.
- [2] T. Saeki, S. Takamichi, *et al.*, "SelfRemaster: Self-supervised speech restoration with analysis-by-synthesis approach using channel modeling," in *Interspeech*, 2022.
- [3] J. Su, Z. Jin, and A. Finkelstein, "HiFi-GAN-2: Studio-quality speech enhancement via generative adversarial networks conditioned on acoustic features," in *WASPAA*, 2021.
- [4] H. Liu, Q. Kong, *et al.*, "VoiceFixer: Toward general speech restoration with neural vocoder," *arXiv:2109.13731*, 2021.
- [5] J. Serrà, S. Pascual, *et al.*, "Universal speech enhancement with score-based diffusion," *arXiv:2206.03065*, 2022.
- [6] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM TASLP*, vol. 26, no. 10, p. 1702–1726, 2018.
- [7] T. Nakatani, T. Yoshioka, *et al.*, "Speech dereverberation based on variance-normalized delayed linear prediction," *IEEE TASLP*, 2010.
- [8] K. Han, Y. Wang, and D. Wang, "Learning spectral mapping for speech dereverberation," in *ICASSP*, 2014.
- [9] K. Kinoshita, M. Delcroix, *et al.*, "Neural network-based spectrum estimation for online WPE dereverberation," in *Interspeech*, 2017.
- [10] P. Závitska, P. Rajmic, *et al.*, "A survey and an extensive evaluation of popular audio declipping methods," *IEEE J. Sel. Top. Signal Process.*, 2021.
- [11] K. Li and C.-H. Lee, "A deep neural network approach to speech bandwidth expansion," in *ICASSP*, 2015.
- [12] V. Kuleshov, S. Z. Enam, and S. Ermon, "Audio super resolution using neural networks," in *ICLR*, 2017.
- [13] A. van den Oord, S. Dieleman, *et al.*, "WaveNet: A generative model for raw audio," *arXiv:1609.03499*, 2016.
- [14] J. Shen, R. Pang, *et al.*, "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *ICASSP*, 2018.
- [15] F. Biadsy, R. J. Weiss, *et al.*, "Parrottron: An end-to-end speech-to-speech conversion model and its applications to hearing-impaired speech and speech separation," in *Interspeech*, 2019.
- [16] W. Chan, D. Park, *et al.*, "Speechstew: Simply mix all available speech recognition data to train one large neural network," *arXiv:2104.02133*, 2021.
- [17] S.-g. Lee, W. Ping, *et al.*, "BigVGAN: A universal neural vocoder with large-scale training," in *ICLR*, 2023.
- [18] Y. Koizumi, K. Yatabe, *et al.*, "WaveFit: An iterative and non-autoregressive neural vocoder based on fixed-point iteration," in *SLT*, 2023.
- [19] A. Baevski, H. Zhou, *et al.*, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, 2020.
- [20] Y.-A. Chung, Y. Zhang, *et al.*, "w2v-BERT: Combining contrastive learning and masked language modeling for self-supervised speech pre-training," in *ASRU*, 2021.
- [21] C.-C. Chiu, J. Qin, *et al.*, "Self-supervised learning with random-projection quantizer for speech recognition," in *ICML*, 2022.
- [22] R. Ardila, M. Branson, *et al.*, "Common Voice: A massively-multilingual speech corpus," in *LREC*, 2020.
- [23] C. Wang, M. Riviere, *et al.*, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *ACL*, 2021.
- [24] W.-N. Hsu, Y. Zhang, *et al.*, "Hierarchical generative modeling for controllable speech synthesis," in *ICLR*, 2019.
- [25] —, "Disentangling correlated speaker and noise for speech synthesis via data augmentation and adversarial factorization," in *ICASSP*, 2019.
- [26] Z. Borsos, M. Sharifi, and M. Tagliasacchi, "SpeechPainter: Text-conditioned speech inpainting," in *Interspeech*, 2022.
- [27] Y. Jia, H. Zen, *et al.*, "PnG BERT: Augmented BERT on phonemes and graphemes for neural TTS," in *Interspeech*, 2021.
- [28] A. Polyak, Y. Adi, *et al.*, "Speech resynthesis from discrete disentangled self-supervised representations," in *Interspeech*, 2021.
- [29] E. Perez, F. Strub, *et al.*, "FiLM: Visual reasoning with a general conditioning layer," in *AAAI*, 2018.
- [30] Y. Koizumi, S. Karita, *et al.*, "DF-Conformer: Integrated architecture of Conv-TasNet and Conformer using linear complexity self-attention for speech enhancement," in *WASPAA*, 2021.
- [31] N. Chen, Y. Zhang, *et al.*, "WaveGrad: Estimating gradients for waveform generation," in *ICLR*, 2021.
- [32] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *ICASSP*, 2020.
- [33] Y. Koizumi, H. Zen, *et al.*, "SpecGrad: Diffusion probabilistic model based neural vocoder with adaptive noise spectral shaping," in *Interspeech*, 2022.
- [34] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS*, 2020.
- [35] A. Gulati, C.-C. Chiu, *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," in *Interspeech*, 2020.
- [36] Q. Wang, Y. Yu, *et al.*, "Attentive temporal pooling for Conformer-based streaming language identification in long-form speech," in *Odyssey*, 2022.
- [37] J. Pelecanos, Q. Wang, *et al.*, "Parameter-free attentive scoring for speaker verification," in *Odyssey*, 2022.
- [38] J. Pelecanos, Q. Wang, and I. L. Moreno, "Dr-Vectors: Decision residual networks and an improved loss for speaker recognition," in *Interspeech*, 2021.
- [39] S. Wang, A. Mesaros, *et al.*, "A curated dataset of urban scenes for audio-visual scene analysis," in *ICASSP*, 2021.
- [40] J. B. Allen and D. A. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Am.*, 1979.
- [41] J. Thiemann, N. Ito, and E. Vincent, "The diverse environments multi-channel acoustic noise database (DEMAND): A database of multi-channel environmental noise recordings," *Meet. Acoust.*, 2013.
- [42] Y. Zhang, J. Qin, *et al.*, "Pushing the limits of semi-supervised learning for automatic speech recognition," in *NeurIPS SAS 2020 Workshop*, 2020.
- [43] Y. Jia, Y. Zhang, *et al.*, "Transfer learning from speaker verification to multispeaker text-to-speech synthesis," in *NeurIPS*, 2018.
- [44] Y. Chen, Y. Assael, *et al.*, "Sample efficient adaptive text-to-speech," in *ICLR*, 2019.
- [45] J. Shen, Y. Jia, *et al.*, "Non-attentive Tacotron: Robust and controllable neural TTS synthesis including unsupervised duration modeling," *arXiv:2010.04301*, 2020.
- [46] N. Kalchbrenner, W. Elsen, *et al.*, "Efficient neural audio synthesis," in *ICML*, 2018.
- [47] H. Zen, R. Clark, *et al.*, "LibriTTS: A corpus derived from LibriSpeech for text-to-speech," in *Interspeech*, 2019.
- [48] K. Ito and L. Johnson, "The LJ speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.