

# WAVE-U-NET DISCRIMINATOR: FAST AND LIGHTWEIGHT DISCRIMINATOR FOR GENERATIVE ADVERSARIAL NETWORK-BASED SPEECH SYNTHESIS

Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, Shogo Seki

NTT Communication Science Laboratories, NTT Corporation, Japan

## ABSTRACT

In speech synthesis, a generative adversarial network (GAN), training a generator (speech synthesizer) and a discriminator in a min-max game, is widely used to improve speech quality. An ensemble of discriminators is commonly used in recent neural vocoders (e.g., HiFi-GAN) and end-to-end text-to-speech (TTS) systems (e.g., VITS) to scrutinize waveforms from multiple perspectives. Such discriminators allow synthesized speech to adequately approach real speech; however, they require an increase in the model size and computation time according to the increase in the number of discriminators. Alternatively, this study proposes a Wave-U-Net discriminator, which is a single but expressive discriminator with Wave-U-Net architecture. This discriminator is unique; it can assess a waveform in a sample-wise manner with the same resolution as the input signal, while extracting multilevel features via an encoder and decoder with skip connections. This architecture provides a generator with sufficiently rich information for the synthesized speech to be closely matched to the real speech. During the experiments, the proposed ideas were applied to a representative neural vocoder (HiFi-GAN) and an end-to-end TTS system (VITS). The results demonstrate that the proposed models can achieve comparable speech quality with a 2.31 times faster and 14.5 times more lightweight discriminator when used in HiFi-GAN and a 1.90 times faster and 9.62 times more lightweight discriminator when used in VITS.<sup>1</sup>


**Index Terms**— Speech synthesis, neural vocoder, text-to-speech, generative adversarial network, Wave-U-Net

## 1. INTRODUCTION

Speech plays an important role in human-human and human-machine communications. To obtain the desired or necessary speech for enriching communications, speech synthesis (e.g., text-to-speech (TTS), a technology for producing speech from text) has been actively studied.

One of the most successful approaches for speech synthesis is the two-stage approach. The first model predicts the intermediate representation (e.g., mel spectrogram) from the input data (e.g., text), and the second model synthesizes speech from a predicted intermediate representation. The second model is called a neural vocoder, and various neural vocoders (e.g., autoregressive models [1, 2], flow models [3, 4], generative adversarial network (GAN) models [5, 6, 7, 8, 9, 10, 11, 12], and diffusion probabilistic models [13, 14]) have been proposed. The merit of this approach is the portability of each model, and a learned neural vocoder is commonly used in other tasks, such as voice conversion (e.g. [15, 16, 17]).

<sup>1</sup>Audio samples are available at <https://www.kecl.ntt.co.jp/people/kaneko.takuhiro/projects/waveunetd/>.



**Fig. 1.** Overview of GAN training with a Wave-U-Net discriminator. A Wave-U-Net discriminator is unique in that it assesses a waveform in a sample-wise manner in the same resolution as the input signal while extracting multilevel features via an encoder and decoder with skip connections.

Another successful approach is the end-to-end approach (e.g., [18, 19, 20, 21]) that directly converts input data (e.g., text) to speech using a unified model. This approach is advantageous for reducing the cascaded errors caused by connecting two separately trained models and eliminating the bias caused by using nonoptimal intermediate representation.

In both approaches, the common objective is to obtain high-quality speech. Thus, a GAN [22], a framework that trains a generator (speech synthesizer) and discriminator in a two-player min-max game, has gained attention and has been widely used both in two-stage [5, 6, 7, 8, 9, 10, 11, 12] and end-to-end [19, 20, 21] approaches. In particular, an ensemble of discriminators is useful for assessing a waveform from multiple perspectives and is commonly used in recent neural vocoders (e.g., HiFi-GAN [8]) and end-to-end TTS systems (e.g., VITS [21]). Such discriminators succeed in bringing the synthesized speech adequately close to real speech; however, they require an increase in the model size and computation time according to the increase in the number of discriminators.

This fact motivates the need to address the following question: “Can one replace an ensemble of discriminators with a single expressive discriminator?” As an answer, this study proposes a *Wave-U-Net discriminator*, a novel discriminator with Wave-U-Net architecture [23, 24], inspired by the recent success of a U-Net discriminator [25] (discriminator with a U-Net architecture [26]) in image synthesis. Figure 1 presents an overview of GAN training with a Wave-U-Net discriminator. As shown in Figure 1, the Wave-U-Net discriminator is unique in that it assesses a waveform in a sample-wise manner with the same resolution as the input signal while extracting multilevel features using an encoder and decoder with skip connections. This architecture provides a generator (speech synthesizer) with sufficiently rich information for synthesized speech to approach real speech. This enabled replacing a typical ensemble of discriminators with a single discriminator.

During the experiments, the general validity of the proposed

ideas was investigated by evaluating Wave-U-Net discriminators in various situations: (1) evaluation on neural vocoders with HiFi-GAN [8] in various datasets, including LJSpeech [27] (single English speaker), VCTK [28] (multiple English speakers), and JSUT [29] (single Japanese speaker), and (2) evaluation on end-to-end TTS with VITS [21]. The results demonstrate that the proposed model can achieve comparable speech quality with a 2.31 times faster and 14.5 times more lightweight discriminator when used in HiFi-GAN and a 1.90 times faster and 9.62 times more lightweight discriminator when used in VITS.

The remainder of this paper is organized as follows: Section 2 briefly reviews GAN-based speech synthesis. Section 3 describes the proposed Wave-U-Net discriminator. Section 4 presents the experimental results, and Section 5 concludes the paper with remarks on future research.

## 2. GAN-BASED SPEECH SYNTHESIS

### 2.1. Overview

A GAN [22] is a framework that trains a generator  $G$  and discriminator  $D$  in a two-player min-max game. In the context of speech synthesis, a generator is identical to a speech synthesizer that synthesizes a speech waveform  $x$  from the input data  $s$ , where  $s$  indicates the intermediate representation (e.g., mel spectrogram) in a two-stage approach and text in an end-to-end TTS system. In GAN-based speech synthesis, the generator and discriminator are trained with two GAN-related losses: adversarial and feature-matching losses.

### 2.2. Losses

**Adversarial losses.** Adversarial losses, particularly the variants of a least squares GAN [30] that are commonly used in speech synthesis, are defined as follows:

$$\mathcal{L}_{Adv}(D) = \mathbb{E}_{(x,s)}[(D(x) - 1)^2 + (D(G(s)))^2], \quad (1)$$

$$\mathcal{L}_{Adv}(G) = \mathbb{E}_s[(D(G(s)) - 1)^2], \quad (2)$$

where  $D$  attempts to distinguish between real and synthesized speech by minimizing  $\mathcal{L}_{Adv}(D)$  and  $G$  attempts to synthesize speech indistinguishable by  $D$  by minimizing  $\mathcal{L}_{Adv}(G)$ .

**Feature-matching loss.** To stabilize the adversarial training, feature-matching loss [31, 32] is simultaneously used:

$$\mathcal{L}_{FM}(G) = \mathbb{E}_{(x,s)} \left[ \sum_{i=1}^T \frac{1}{N_i} \|D^i(x) - D^i(G(s))\|_1 \right], \quad (3)$$


where  $T$  indicates the number of layers in  $D$ ,  $D^i$  and  $N_i$  denote the features and number of features in the  $i$ -th layer of  $D$ . Here,  $G$  attempts to synthesize speech close to the ground-truth speech by minimizing  $\mathcal{L}_{FM}(G)$ .

To further stabilize the training, losses other than GAN-related ones were also used in typical GAN-based speech synthesis. For example, reconstruction losses, such as spectrogram loss [6] and mel-spectrogram loss [8], are widely used. Owing to space limitations, this study omits explanations of them. Please refer to the corresponding studies for more details.

## 3. WAVE-U-NET DISCRIMINATOR

### 3.1. Overview

As discussed in Section 2, in a GAN, the generator is trained and guided by an adversarial discriminator. To guide the learning of the



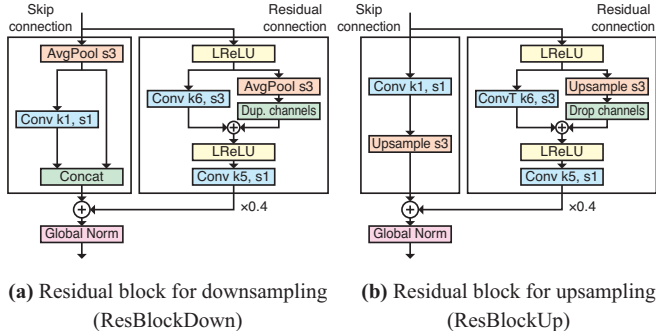
**Fig. 2.** Overview architectures of a Wave-U-Net discriminator.  $cx$ ,  $ky$ , and  $sz$  indicate the number of output channels of  $x$ , kernel size of  $y$ , and stride of  $z$ , respectively. The details of residual blocks for downsampling (ResBlockDown) and upsampling (ResBlockUp) are presented in Figure 3.

generator in the correct direction, the discriminator must capture a speech waveform sufficiently and propagate adequate information to the generator. In particular, speech waveforms have multilevel (e.g., multiscale and multiperiod) structures; therefore, the discriminator must adequately capture such structures.

To satisfy this requirement with a single discriminator, this study introduced a Wave-U-Net discriminator with Wave-U-Net architecture [23, 24]. Figure 2 presents an overview of the Wave-U-Net discriminator architecture. As shown in this figure, a Wave-U-Net discriminator differs from a typical discriminator in that a typical discriminator only constitutes an encoder and judges the reality of speech using abstracted features obtained via downsampling. By contrast, a Wave-U-Net discriminator has an encoder-decoder architecture and assesses a waveform in a sample-wise manner with the same resolution as the input waveform by applying upsampling after downsampling. In this process, multilevel features are efficiently extracted using an encoder and decoder with skip connections. This study suggests that the architectural design provides a rich information source for synthesized speech to approach real speech, similar to a typical ensemble of discriminators. We empirically demonstrate the validity of this statement in various situations through the experiments described in Section 4.

### 3.2. Techniques for stabilizing GAN training

It is known that GANs are sensitive to architectural design because they conduct an unstable min-max game during training. This also holds true in this case, and preliminary experiments indicated that the direct use of a conventional Wave-U-Net [23, 24] architecture saturated adversarial losses. This phenomenon could occur when a discriminator restricts itself to specific features that make it possible to distinguish real from fake data easily and experiences difficulties in propagating adequate information to a generator. Through extensive exploration, it was determined that careful normalization and the introduction of residual connections are necessary and sufficient to alleviate this deterioration.



**Fig. 3.** Detailed architectures of residual blocks in a Wave-U-Net discriminator.  $k_x$  and  $s_y$  indicate the kernel size of  $x$  and stride of  $y$ , respectively. The number of channels increases and decreases in the first convolution layer (Conv) and the first transposed convolution layer (ConvT) for downsampling and upsampling, respectively. A leaky rectified linear unit (LReLU) [33] with a negative slope of 0.1 is used as an activation function. “Dup. channels” is the abbreviation for “Duplicate channels”, “+” denotes the addition of features, and “ $\times 0.4$ ” indicates that the features are scaled by a constant of 0.4.

**Normalization.** A previous study [5] showed that in speech synthesis, a discriminator is sensitive to normalization, and normalization commonly used in image synthesis, such as spectral normalization [34] and instance normalization [35], does not work well. Alternatively, a previous study [5] applied weight normalization [36], which performed well with the discriminator. However, this study determined that weight normalization is not sufficient to stabilize training with the Wave-U-Net discriminator, possibly because the Wave-U-Net discriminator, which consists of an encoder and a decoder, is deeper than typical discriminators, which include only an encoder, and it is difficult to avoid restricting itself to specific features. To alleviate this problem, this study proposes global normalization (Global Norm), which is defined as

$$b = a \sqrt{\frac{1}{N} \sum_{i=1}^N (a^i)^2 + \epsilon}, \quad (4)$$

where  $\epsilon = 10^{-8}$ ;  $N$  denotes the total number of features;  $a$  and  $b$  represent the original and normalized feature vectors, respectively; and  $a^i$  indicates the  $i$ -th feature in  $a$ . This normalization prevents the discriminator from restricting itself to specific features by regularizing the norms of the features within a constant range.

One possible alternative is layer normalization [37]. Unlike global normalization, layer normalization uses trainable bias and gain parameters and conducts a sample-wise zero-mean shift. Preliminary experiments indicated that the first property invoked training instability and the second deteriorated cFW2VD [11], which is a metric for assessing speech quality and is explained in detail in Section 4.1. Based on these preliminary experiments, this study used global normalization, which only scales features without any trainable parameters.

**Residual connections.** This study empirically determined global normalization to be effective even for a non-residual Wave-U-Net. However, adversarial losses still approached saturation at the end of the training, possibly because the Wave-U-Net discriminator was deeper, as previously discussed. Hence, this study introduces residual connections [38] to prevent the vanishing gradient problem that

typically occurs in deep networks. Inspired by the success of a previous GAN [39], the output of the residual connection was scaled by a constant of 0.4. This technique contributed slightly to stabilizing the training. Figure 3 shows the detailed architecture of the residual blocks equipped with the aforementioned modules.

## 4. EXPERIMENTS

Experiments were conducted to examine the general validity of the proposed ideas, that is, whether a Wave-U-Net discriminator can be used as an alternative to a typical ensemble of discriminators in various situations. This study first examined the dataset independence by applying the presented ideas to neural vocoders (particularly HiFi-GAN [8]) on various datasets. Subsequently, task independence was investigated by applying the presented ideas to an end-to-end TTS system (particularly VITS [21]). The details are discussed in Sections 4.1 and 4.2. Audio samples are available from the link on the first page.<sup>1</sup>

### 4.1. Evaluation on neural vocoders

**Implementation.** We first examined the effectiveness of a Wave-U-Net discriminator (hereafter denoted by *Wave-U-Net D* for simplicity) when used to train a neural vocoder. Specifically, *HiFi-GAN* [8] was used as the baseline and the difference in performance was examined when a Wave-U-Net discriminator was used as an alternative to the original ensemble of discriminators. The proposed model was implemented based on the official implementation of HiFi-GAN,<sup>2</sup> and the default training settings were used. The networks were trained using the AdamW optimizer [40] with  $\beta_1 = 0.8$ ,  $\beta_2 = 0.99$ , and weight decay  $\lambda = 0.01$ . The learning rate decayed exponentially by 0.999 at fixed intervals with an initial learning rate of  $2 \times 10^{-4}$ . The batch size was set to 16 with a segment length of 8192.

**Datasets.** To examine the dataset independence, the models were evaluated based on three different datasets: *LJSpeech* [27], which consists of 13,100 audio clips of a single English speaker, and 12,500, 100, and 500 audio clips were used for training, validation, and evaluation, respectively. Specifically, data splitting was conducted based on an open-source code.<sup>5</sup> *VCTK* [28], which includes 44,081 audio clips from 108 different English speakers, and 41,921, 1,080, and 1,080 audio clips were used for training, validation, and evaluation, respectively. The audio clips were divided based on an open-source code.<sup>3</sup> *JSUT* [29], which is composed of 7,696 audio clips from a single Japanese speaker, and 7,196, 250, and 250 audio clips were used for training, validation, and evaluation, respectively. Specifically, the audio clips were split according to an open-source code.<sup>3</sup> Following the study on HiFi-GAN [8], audio clips were sampled at 22.05 kHz, and 80-dimensional log-mel spectrograms with an FFT size of 1024, hop length of 256, and window length of 1024 were extracted from the audio clips. Neural vocoders were individually trained for each dataset.

**Evaluation metrics.** We employed four metrics to evaluate the performance: *Mean opinion score (MOS) tests* were conducted to evaluate the perceptual quality. Twenty utterances were randomly selected from the evaluation set, and the mel spectrograms extracted from the utterances were used as the vocoder input. The *ground-truth* speech

<sup>2</sup><https://github.com/jik876/hifi-gan>

<sup>3</sup><https://github.com/kan-bayashi/ParallelWaveGAN>

**Table 1.** Comparison of MOS with 95% confidence intervals, cFW2VD, processing time, and model size among neural vocoders on LJSpeech.

Model	MOS $\uparrow$	cFW2VD $\downarrow$	Time $\downarrow$ (s/batch)	# Param $\downarrow$ (M)
Ground truth	4.72 $\pm$ 0.07	0.000	–	–
HiFi-GAN	4.57 $\pm$ 0.09	0.014	0.028	70.7
w/ Wave-U-Net $D$	4.54 $\pm$ 0.09	0.014	0.012	4.9
MelGAN	3.26 $\pm$ 0.14	0.193	–	–

**Table 2.** Comparison of MOS with 95% confidence intervals, cFW2VD, processing time, and model size among neural vocoders on VCTK.

Model	MOS $\uparrow$	cFW2VD $\downarrow$	Time $\downarrow$ (s/batch)	# Param $\downarrow$ (M)
Ground truth	4.62 $\pm$ 0.07	0.000	–	–
HiFi-GAN	4.37 $\pm$ 0.10	0.040	0.028	70.7
w/ Wave-U-Net $D$	4.44 $\pm$ 0.09	0.041	0.012	4.9
MelGAN	3.60 $\pm$ 0.13	0.301	–	–

**Table 3.** Comparison of MOS with 95% confidence intervals, cFW2VD, processing time, and model size on among neural vocoders on JSUT.

Model	MOS $\uparrow$	cFW2VD $\downarrow$	Time $\downarrow$ (s/batch)	# Param $\downarrow$ (M)
Ground truth	4.75 $\pm$ 0.07	0.000	–	–
HiFi-GAN	4.67 $\pm$ 0.07	0.039	0.028	70.7
w/ Wave-U-Net $D$	4.69 $\pm$ 0.07	0.038	0.012	4.9
MelGAN	3.68 $\pm$ 0.13	0.222	–	–

and speech synthesized by *MelGAN* [5] were used as anchor samples. These tests were conducted online with 11 listeners participating. The listeners were asked to judge the speech quality using a five-point scale: 1 = bad, 2 = poor, 3 = fair, 4 = good, and 5 = excellent. The *conditional Fréchet wav2vec distance (cFW2VD)* [11] was used as an objective metric for assessing the speech quality. This metric measures the distance between the real and generative distributions in a wav2vec 2.0 [41] feature space conditioned on the text. The smaller the value, the higher is the similarity between the real and synthesized speech. *Number of parameters (# Param)* was used to investigate the model size. The smaller the value, the more lightweight the model. For the training speed, the *time* required for a discriminator to forward-propagate real and synthesized speech in a batch was measured on a single NVIDIA A100 GPU. Here, the smaller the value, the faster the speed.

**Results.** Tables 1, 2, and 3 present the results for LJSpeech, VCTK, and JSUT, respectively. It was determined that, across all the datasets, the proposed model (*w/ Wave-U-Net D*) achieved comparable speech quality with the baseline (*HiFi-GAN*) in terms of MOS<sup>4</sup> and cFW2VD with a 2.31 times faster and 14.5 times more lightweight discriminator. These results indicate that the Wave-U-Net discriminator can be used as an alternative to the HiFi-GAN discriminator regardless of the dataset.

<sup>4</sup>The  $p$ -values for the results in Tables 1, 2, and 3 were 0.622, 0.533, and 0.708, respectively, in the Mann-Whitney U tests. The results indicate that the two models do not differ significantly regarding the  $p$ -value  $< 0.05$ .

**Table 4.** Comparison of MOS with 95% confidence intervals, cFW2VD, processing time, and model size among TTS systems.

Model	MOS $\uparrow$	cFW2VD $\downarrow$	Time $\downarrow$ (s/batch)	# Param $\downarrow$ (M)
Ground truth	4.65 $\pm$ 0.07	0.000	–	–
VITS	4.51 $\pm$ 0.09	0.109	0.030	46.7
w/ Wave-U-Net $D$	4.54 $\pm$ 0.09	0.101	0.016	4.9
Tacotron 2 + HiFi-GAN	3.78 $\pm$ 0.14	0.261	–	–

## 4.2. Evaluation on end-to-end TTS

**Implementation.** To investigate task independence, the proposed ideas were applied to end-to-end TTS. Specifically, *VITS* [21] was used as the baseline, and the performance difference was examined when the original ensemble of discriminators was replaced with a Wave-U-Net discriminator. The proposed model was implemented based on the official implementation of *VITS*<sup>5</sup> and the default training settings were used. The networks were trained using the AdamW optimizer [40] with  $\beta_1 = 0.8$ ,  $\beta_2 = 0.99$ , and weight decay  $\lambda = 0.01$ . The learning rate decayed exponentially by  $0.999^{\frac{1}{8}}$  in fixed intervals with an initial learning rate of  $2 \times 10^{-4}$ . The batch size was set to 64 with a segment length of 8192.

**Dataset.** LJSpeech [27] was used for the experiment. The data splitting setting is the same as that used in the experiment described in Section 4.1.

**Evaluation metrics.** The same evaluation metrics are used to evaluate the neural vocoders (Section 4.1). In the MOS test, the anchor samples were altered because of the task differences. In this experiment, *ground-truth* speech and speech synthesized using a combination of Tacotron 2 [42] and HiFi-GAN [8] (*Tacotron 2 + HiFi-GAN*) were used as anchor samples.

**Results.** Table 4 summarizes the results. The results show that the proposed model (*w/ Wave-U-Net D*) achieved speech quality comparable to that of the baseline model (*VITS*) in terms of the MOS<sup>6</sup> and cFW2VD with a 1.90 times faster and 9.62 times more lightweight discriminator. These results verified the task independence of the Wave-U-Net discriminator.

## 5. CONCLUSION

This study proposes a Wave-U-Net discriminator, which is a single but expressive discriminator that assesses a waveform in a sample-wise manner with the same resolution as the input signal while extracting multilevel features via an encoder and decoder with skip connections. The experimental results demonstrate that a Wave-U-Net discriminator can be used as an alternative to a typical ensemble of discriminators while maintaining speech quality, reducing the model size, and accelerating the training speed. Although the general utility of a Wave-U-Net discriminator has been demonstrated to some extent, there are several other tasks that are beyond the scope of this study, such as singing speech synthesis, emotional speech synthesis, and music synthesis. Utilization of the proposed ideas is a valuable future research topic.

**Acknowledgements.** This work was supported by JST CREST Grant Number JPMJCR19A3, Japan.

<sup>5</sup><https://github.com/jaywalnut310/vits>

<sup>6</sup>The  $p$ -value was 0.345 in the Mann-Whitney U test. This score indicates that the two models do not differ significantly regarding the  $p$ -value  $< 0.05$ .

## 6. REFERENCES

- [1] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, “WaveNet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [2] Nal Kalchbrenner, Erich Elsen, Karen Simonyan, Seb Noury, Norman Casagrande, Edward Lockhart, Florian Stimberg, Aäron van den Oord, Sander Dieleman, and Koray Kavukcuoglu, “Efficient neural audio synthesis,” in *Proc. ICML*, 2018, pp. 2410–2419.
- [3] Aäron van den Oord, Yazhe Li, Igor Babuschkin, Karen Simonyan, Oriol Vinyals, Koray Kavukcuoglu, George van den Driessche, Edward Lockhart, Luis Cobo, Florian Stimberg, Norman Casagrande, Dominik Grewe, Seb Noury, Sander Dieleman, Erich Elsen, Nal Kalchbrenner, Heiga Zen, Alex Graves, Helen King, Tom Walters, Dan Belov, and Demis Hassabis, “Parallel WaveNet: Fast high-fidelity speech synthesis,” in *Proc. ICML*, 2018, pp. 3918–3926.
- [4] Ryan Prenger, Rafael Valle, and Bryan Catanzaro, “WaveGlow: A flow-based generative network for speech synthesis,” in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [5] Kundan Kumar, Rithesh Kumar, Thibault de Boissiere, Lucas Gestin, Wei Zhen Teoh, Jose Sotelo, Alexandre de Brébisson, Yoshua Bengio, and Aaron Courville, “MelGAN: Generative adversarial networks for conditional waveform synthesis,” in *Proc. NeurIPS*, 2019, pp. 14910–14921.
- [6] Ryuichi Yamamoto, Eunwoo Song, and Jae-Min Kim, “Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram,” in *Proc. ICASSP*, 2020, pp. 6199–6203.
- [7] Jinhyeok Yang, Junmo Lee, Youngik Kim, Hoonyoung Cho, and Injung Kim, “VocGAN: A high-fidelity real-time vocoder with a hierarchically-nested adversarial network,” in *Proc. Interspeech*, 2020, pp. 200–204.
- [8] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” in *Proc. NeurIPS*, 2020, pp. 17022–17033.
- [9] Geng Yang, Shan Yang, Kai Liu, Peng Fang, Wei Chen, and Lei Xie, “Multi-band MelGAN: Faster waveform generation for high-quality text-to-speech,” in *Proc. SLT*, 2021, pp. 492–498.
- [10] Ahmed Mustafa, Nicola Pia, and Guillaume Fuchs, “StyleMelGAN: An efficient high-fidelity adversarial vocoder with temporal adaptive normalization,” in *Proc. ICASSP*, 2021, pp. 6034–6038.
- [11] Takuhiro Kaneko, Kou Tanaka, Hirokazu Kameoka, and Shogo Seki, “iSTFTNet: Fast and lightweight mel-spectrogram vocoder incorporating inverse short-time Fourier transform,” in *Proc. ICASSP*, 2022, pp. 6207–6211.
- [12] Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, and Shogo Seki, “MISRNet: Lightweight neural vocoder using multi-input single shared residual blocks,” in *Proc. Interspeech*, 2022, pp. 1631–1635.
- [13] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J. Weiss, Mohammad Norouzi, and William Chan, “WaveGrad: Estimating gradients for waveform generation,” in *Proc. ICLR*, 2021.
- [14] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro, “DiffWave: A versatile diffusion model for audio synthesis,” in *Proc. ICLR*, 2021.
- [15] Kaizhi Qian, Yang Zhang, Shiyu Chang, Xuesong Yang, and Mark Hasegawa-Johnson, “Auto-VC: Zero-shot voice style transfer with only autoencoder loss,” in *Proc. ICML*, 2019, pp. 5210–5219.
- [16] Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, and Nobukatsu Hojo, “CycleGAN-VC3: Examining and improving CycleGAN-VCs for mel-spectrogram conversion,” in *Proc. Interspeech*, 2020, pp. 2017–2021.
- [17] Takuhiro Kaneko, Hirokazu Kameoka, Kou Tanaka, and Nobukatsu Hojo, “MaskCycleGAN-VC: Learning non-parallel voice conversion with filling in frames,” in *Proc. ICASSP*, 2021, pp. 5919–5923.
- [18] Wei Ping, Kainan Peng, and Jitong Chen, “ClariNet: Parallel wave generation in end-to-end text-to-speech,” in *Proc. ICLR*, 2019.
- [19] Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, “FastSpeech 2: Fast and high-quality end-to-end text to speech,” in *Proc. ICLR*, 2021.
- [20] Jeff Donahue, Sander Dieleman, Mikołaj Bińkowski, Erich Elsen, and Karen Simonyan, “End-to-end adversarial text-to-speech,” in *Proc. ICLR*, 2021.
- [21] Jaehyeon Kim, Jungil Kong, and Juhee Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *Proc. ICML*, 2021, pp. 5530–5540.
- [22] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Proc. NIPS*, 2014, pp. 2672–2680.
- [23] Santiago Pascual, Antonio Bonafonte, and Joan Serra, “SEGAN: Speech enhancement generative adversarial network,” in *Proc. Interspeech*, 2017, pp. 3642–3646.
- [24] Daniel Stoller, Sebastian Ewert, and Simon Dixon, “Wave-U-Net: A multi-scale neural network for end-to-end audio source separation,” in *Proc. ISMIR*, 2018, pp. 334–340.
- [25] Edgar Schonfeld, Bernt Schiele, and Anna Khoreva, “A U-Net based discriminator for generative adversarial networks,” in *Proc. CVPR*, 2020, pp. 8207–8216.
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *Proc. MICCAI*, 2015, pp. 234–241.
- [27] Keith Ito and Linda Johnson, “The LJ speech dataset,” <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [28] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald, “CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit,” *The Centre for Speech Technology Research*, 2016.
- [29] Ryosuke Sonobe, Shinnosuke Takamichi, and Hiroshi Saruwatari, “JSUT corpus: free large-scale Japanese speech corpus for end-to-end speech synthesis,” *arXiv preprint arXiv:1711.00354*, 2017.
- [30] Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley, “Least squares generative adversarial networks,” in *Proc. ICCV*, 2017, pp. 2794–2802.
- [31] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther, “Autoencoding beyond pixels using a learned similarity metric,” in *Proc. ICML*, 2016, pp. 1558–1566.
- [32] Takuhiro Kaneko, Hirokazu Kameoka, Kaoru Hiramatsu, and Kunio Kashino, “Sequence-to-sequence voice conversion with similarity metric learned using generative adversarial networks,” in *Proc. Interspeech*, 2017, pp. 1283–1287.
- [33] Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng, “Rectifier nonlinearities improve neural network acoustic models,” in *Proc. ICML*, 2013.
- [34] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida, “Spectral normalization for generative adversarial networks,” in *Proc. ICLR*, 2018.
- [35] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky, “Instance normalization: The missing ingredient for fast stylization,” *arXiv preprint arXiv:1607.08022*, 2016.
- [36] Tim Salimans and Diederik P. Kingma, “Weight normalization: A simple reparameterization to accelerate training of deep neural networks,” in *Proc. NIPS*, 2016, pp. 901–909.
- [37] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton, “Layer normalization,” *arXiv preprint arXiv:1607.06450*, 2016.
- [38] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016, pp. 770–778.
- [39] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin, “Which training methods for GANs do actually converge?,” in *Proc. ICML*, 2018, pp. 3481–3490.
- [40] Ilya Loshchilov and Frank Hutter, “Decoupled weight decay regularization,” in *Proc. ICLR*, 2019.
- [41] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, 2020, pp. 12449–12460.
- [42] Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, RJ Skerrv-Ryan, Rif A. Saurous, Yannis Agiomyrgiannakis, and Yonghui Wu, “Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions,” in *Proc. ICASSP*, 2018, pp. 4779–4783.