



MSR-NV: Neural Vocoder Using Multiple Sampling Rates

Kentaro Mitsui, Kei Sawada

rinna Co., Ltd., Japan

kemits@rinna.co.jp, keisawada@rinna.co.jp

Abstract

The development of neural vocoders (NVs) has resulted in the high-quality and fast generation of waveforms. However, conventional NVs target a single sampling rate and require re-training when applied to different sampling rates. A suitable sampling rate varies from application to application due to the trade-off between speech quality and generation speed. In this study, we propose a method to handle multiple sampling rates in a single NV, called the MSR-NV. By generating waveforms step-by-step starting from a low sampling rate, MSR-NV can efficiently learn the characteristics of each frequency band and synthesize high-quality speech at multiple sampling rates. It can be regarded as an extension of the previously proposed NVs, and in this study, we extend the structure of Parallel WaveGAN (PWG). Experimental evaluation results demonstrate that the proposed method achieves remarkably higher subjective quality than the original PWG trained separately at 16, 24, and 48 kHz, without increasing the inference time. We also show that MSR-NV can leverage speech with lower sampling rates to further improve the quality of the synthetic speech.

Index Terms: neural vocoder, speech synthesis, sampling rate, generative adversarial networks, Parallel WaveGAN

1. Introduction

In text-to-speech synthesis, singing voice synthesis, and music synthesis, studies to improve the quality of vocoders, which generate waveforms from acoustic features, have been extensively conducted. Neural vocoders (NVs), which use neural networks to generate waveforms, have greatly improved the quality of synthetic audio. WaveNet [1, 2], a representative of autoregressive (AR) NVs, achieved a significantly higher quality than conventional signal processing-based vocoders [3, 4] by predicting waveform samples individually. Although AR NVs can generate high-quality audio because they can use previous prediction results, they suffer from a slow generation speed. To address this issue, fast and high-quality waveform generation using non-autoregressive (non-AR) NVs has been actively studied. Various generative models have been used in this area, including inverse autoregressive flow (IAF) [5] used in Parallel WaveNet [6] and ClariNet [7], generative flow (Glow) [8] used in WaveGlow [9], the generative adversarial network (GAN) [10] used in several NVs [11, 12, 13, 14], and the denoising diffusion probabilistic model (DDPM) [15] used in WaveGrad [16] and DiffWave [17]. Moreover, several methods incorporating signal processing insights have been proposed, including the neural source-filter model [18] and LPCNet [19].

The sampling rate of the waveform plays a key role in waveform generation. It represents the resolution in the time domain when waveforms are being handled on a computer. The higher the sampling rate, the more accurately the original waveform can be expressed, but the larger will be the amount of data. With

regard to speech, its content and individuality are concentrated in the low frequency range. Therefore, sampling rates of 16, 22.05, and 24 kHz are often used in research on NVs. In an environment where only synthetic speech is heard, it does not matter even if the human audible range (20–20,000 Hz) is not entirely covered. However, in situations where synthetic speech is used for conversing with a human, or where it is played simultaneously with the sound of musical instruments in music, the synthetic speech may sound muffled compared to other sounds. Some studies (Full-band LPCNet [20] and PeriodNet [21]) have tackled this problem by generating 48 kHz waveforms. However, in general, the higher the sampling rate, the more difficult it is to generate high-quality waveforms because of the need to model long-term dependencies. Additionally, because conventional NVs target a specific sampling rate, speech with sampling rates lower than that of the target cannot be used for training. This is because even if we upsample a waveform with such low sampling rates, high-frequency components cannot be recovered. Although large speech corpora such as the LJSpeech dataset¹ and LibriTTS [22] have been recently released, 44.1 or 48 kHz datasets are still limited.

For these reasons, a method that: 1) can properly model speech with high sampling rates, and 2) also use speech with a sampling rate lower than that of the target for training, is required. In this study, we propose multiple sampling rate (MSR) -NV as a method with the aforementioned characteristics. MSR-NV can generate speech waveforms step-by-step, starting from a low sampling rate. More specifically, after generating a waveform with a certain sampling rate, sinc interpolation is performed to upsample the waveform, and a neural network is used to predict the residual high-frequency components to generate a waveform with a higher sampling rate. This would allow different networks to capture the features contained in each frequency band and efficiently model waveforms with high sampling rates. Additionally, MSR-NV enables us to train a part of the model using speech with a low sampling rate. As a result, it is expected that speech with sampling rates that could not be used in the past can be used together for training, which leads to the realization of a more general-purpose NV.

We conducted experimental evaluations using a model structure based on Parallel WaveGAN (PWG) [12] to demonstrate the effectiveness of the proposed method. First, we compared the subjective quality of the speech generated at 16, 24, and 48 kHz using the proposed method with that of the speech generated using the baseline model that was trained separately at these sampling rates. To investigate the data efficiency of the proposed method, we evaluated the quality of the synthetic speech when the amount of training data was varied from 1 min to 8 h. Finally, we confirmed that the quality of synthetic speech can be improved using speech with lower sampling rates in conjunction with the original training data.

¹<https://keithito.com/LJ-Speech-Dataset/>

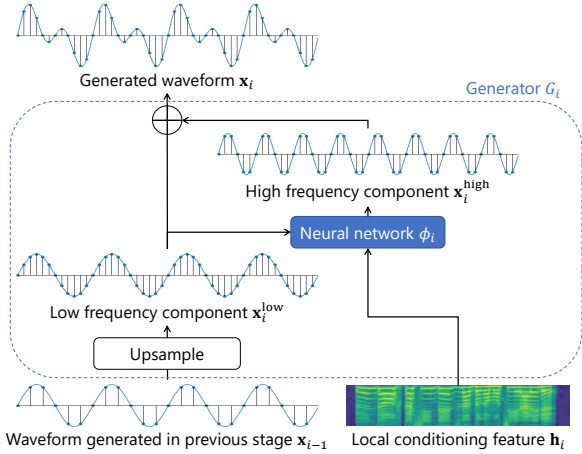


Figure 1: Algorithm of generating waveforms of multiple sampling rates step-by-step.

2. Related works

Obtaining a high-resolution output is a common challenge not only for speech waveform generation, but also for image generation. Progressive growing GAN [23] enables the generation of images with unprecedented 1024×1024 pixels by the gradual addition of layers corresponding to higher resolutions. The later published StyleGAN [24] enables the generation of high-resolution images without changing the topology of the network by preparing layers for multiple resolutions in advance. Each layer of the network can represent features corresponding to different resolutions using these methods, which can handle multiple resolutions in a stepwise manner.

Several methods have been proposed for speech waveform generation to handle multiple sampling rates. MelGAN [11] and HiFi-GAN [14] are methods that can predict waveform by repeatedly upsampling and transforming features. HiFi-GAN achieves a quality comparable to natural speech at 22.05 kHz. VocGAN [13] has a generator similar to that of MelGAN, but generates and evaluates waveforms with $\times 1/n$ ($n = 2, 4, 8, 16$) sampling rates. However, while the proposed method directly upsamples the waveform and predicts the residual high-frequency components, VocGAN upsamples the features; thus, when a waveform is viewed at multiple sampling rates, the residual structure is not explicitly used. Additionally, when a high sampling rate (e.g., 44.1, 48 kHz) that covers the entire human audible range is targeted, it is difficult to use speech data with lower sampling rates (e.g., 16, 22.05 kHz) for training. Even if we upsample them, high-frequency components cannot be recovered.

3. MSR-NV

3.1. Sequential waveform generation of multiple sampling rates

In this section, we describe the procedure for generating speech waveforms at multiple sampling rates using the proposed method. Let $f_1 < \dots < f_i < \dots < f_I$ be the sequence of the sampling rates to be generated.

First, we consider $i \geq 2$, where we predict the waveform \mathbf{x}_i with a sampling rate f_i from waveform \mathbf{x}_{i-1} with a sampling rate f_{i-1} . From the linearity of the Fourier transform, \mathbf{x}_i can be expressed as the sum of a waveform $\mathbf{x}_i^{\text{low}}$ with a component at frequency $[0, f_{i-1}/2]$ and a waveform $\mathbf{x}_i^{\text{high}}$ with a component at frequency $[f_{i-1}/2, f_i/2]$. Therefore, by repeating the

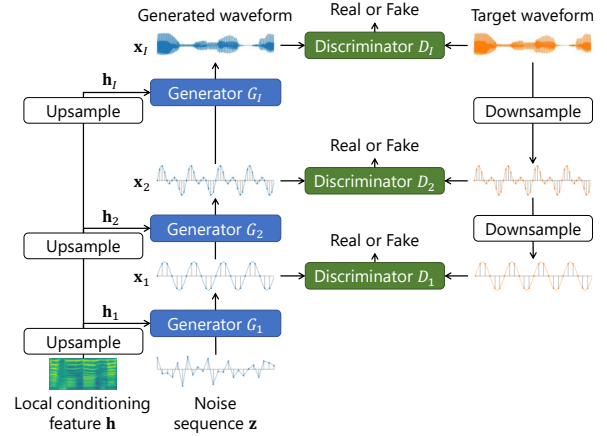


Figure 2: Parallel WaveGAN-based model structure extended using the proposed method.

following procedure, we can sequentially generate a waveform with a higher sampling rate starting from a low sampling rate:

1. Upsample \mathbf{x}_{i-1} to sampling rate f_i to approximate $\mathbf{x}_i^{\text{low}}$
2. Based on the upsampled waveform $\mathbf{x}_i^{\text{low}}$ and the conditioning feature \mathbf{h}_i (e.g. upsampled melspectrogram), predict $\mathbf{x}_i^{\text{high}} = \phi_i(\mathbf{x}_i^{\text{low}}, \mathbf{h}_i)$ using a neural network ϕ_i
3. Calculate the sum of 1 and 2: $\mathbf{x}_i = \mathbf{x}_i^{\text{low}} + \mathbf{x}_i^{\text{high}}$

The module that performs the abovementioned procedure is referred to as Generator G_i . Fig. 1 illustrates these steps. Regarding step 1, transposed convolution is often used for upsampling in NVs [11, 13, 14]. However, this requires that f_i is divisible by f_{i-1} and may include high-frequency components in the range $[f_{i-1}/2, f_i/2]$ that were not included before upsampling. For these reasons, the proposed method uses upsampling based on sinc interpolation.

Then, we consider the case where $i = 1$. When generating the waveform \mathbf{x}_1 of the lowest sampling rate f_1 , the previous waveform \mathbf{x}_0 does not exist. Thus, the function ϕ_1 predicts $\mathbf{x}_1 = \phi_1(\mathbf{h}_1)$ with only the conditioning feature \mathbf{h}_1 as the input. In the proposed method, f_1 can be set arbitrarily. According to Oura et al. [25] and Hono et al. [21], a sinusoidal input corresponding to the fundamental frequency of speech is effective in predicting the periodic component of speech. Therefore, we set f_1 such that the fundamental frequency is included in $[0, f_1/2]$ to make \mathbf{x}_1 similar to the sine wave corresponding to the fundamental frequency. We expect that \mathbf{x}_1 will facilitate waveform generation in the subsequent stages.

3.2. Model details

The proposed method described in section 3.1 can be applied to NVs of various structures and enables training at high sampling rates, which is difficult with conventional NVs. In the case of $I = 1$, the model structure matches that of a conventional NV, and in the case of $I > 1$, it can be regarded as an extension using the proposed method. We used a model structure based on PWG [12]. Although PWG has undergone several improvements since it was first introduced [26, 27, 28, 29, 30], we used the original PWG to verify the effectiveness of the proposed method with the simplest model structure. A conceptual model is depicted in Fig. 2. As the function ϕ_i described in section 3.1, we use the same structure as the generator of PWG, that is, a

non-causal WaveNet. To predict $\mathbf{x}_i^{\text{high}} = \phi_i(\mathbf{x}_i^{\text{low}}, \mathbf{h}_i)$ with a WaveNet-based structure, we must match the temporal resolution of $\mathbf{x}_i^{\text{low}}$ and \mathbf{h}_i . Thus, we obtain \mathbf{h}_i by properly upsampling the original conditioning feature \mathbf{h} with sinc interpolation. In the case of $i = 1$, we follow PWG and predict $\mathbf{x}_1 = \phi_1(\mathbf{z}, \mathbf{h}_1)$ using white noise \mathbf{z} as input, in addition to the conditioning feature \mathbf{h}_1 .

The proposed method generates I waveforms from \mathbf{x}_1 to \mathbf{x}_J . Therefore, we use I discriminators $\{D_i\}_{i=1}^I$ to identify whether the corresponding waveform is real or fake. All the discriminators have identical structure to that of PWG. A multi-scale discriminator (MSD) [31] has also been proposed as a structure to identify downsampled waveforms. However, MSD loses high-frequency components owing to low-pass processing associated with average pooling. Thus, we downsampled the natural speech in advance using an anti-aliasing filter and sinc interpolation, and used it as a target at each sampling rate.

We simply extend the loss function used in PWG [12] to multiple sampling rates. For each sampling rate f_i , we add multi-resolution short time Fourier transform (MR-STFT) loss $\mathcal{L}_{\text{aux}}(G_i)$ to adversarial loss $\mathcal{L}_{\text{adv}}(G_i, D_i)$ weighted by λ_{adv} , and use the sum of these as the loss function $\mathcal{L}_G(G, D)$ for the entire generator. For the discriminator, we use the sum of the losses $\mathcal{L}_{\text{dis}}(G_i, D_i)$ for the discrimination results of D_i as the loss function $\mathcal{L}_D(G, D)$ for the entire discriminator. The above can be expressed as follows:

$$\mathcal{L}_G(G, D) = \sum_{i=1}^I \{\mathcal{L}_{\text{aux}}(G_i) + \lambda_{\text{adv}}\mathcal{L}_{\text{adv}}(G_i, D_i)\} \quad (1)$$

$$\mathcal{L}_D(G, D) = \sum_{i=1}^I \mathcal{L}_{\text{dis}}(G_i, D_i) \quad (2)$$

When we use speech with a sampling rate f_J ($J < I$) lower than f_I for training, we can train a part of the network by changing the range of summation in Eqs. (1) and (2) from $i = 1, \dots, I$ to $i = 1, \dots, J$.

4. Experiments

4.1. Experimental conditions

Experiments were conducted in two settings: single-speaker and multi-speaker. In the single-speaker experiment, a total of 14,375 utterances (approximately 8 h) uttered in a normal speaking style by a female Japanese speaker were used. For the development and evaluation sets, respectively, 300 utterances were randomly selected, and the remaining 13,775 utterances were used as the training set. In the multi-speaker experiment, a total of 31,936 utterances (approximately 25 h) by 26 Japanese speakers (18 female and 8 male), including three speaking styles (normal, happy, and sad) were used. A total of 520 utterances, 20 from each speaker, were used for the development and evaluation sets, respectively, and the remaining 30,896 utterances were used for the training set. All the experiments were conducted using 48 kHz/16 bit speech signals. They were trimmed so that the silent interval before and after the speech was approximately 200 ms. Eighty-dimensional log mel spectrograms with bands in the range 80–7,600 Hz were used as the conditioning feature \mathbf{h} described in section 3.2. They were extracted with a frame and window length of 2,048 points (approximately 42.7 ms) and a frame shift of 240 points (5 ms) and were normalized to zero mean and unit variance.

The sampling rates handled using the proposed method were set to $I = 7$ and $\{f_i\}_{i=1}^7 = \{1, 2, 4, 8, 16, 24, 48\}$ kHz. f_1 was set to 1 kHz so that the fundamental frequency would be

Table 1: Results of MOS evaluation on closeness to recorded speech quality with 95% confidence intervals.

Dataset	Ref	PWG	MSR-PWG
single/16k	2.48±0.16	1.50±0.11	2.26±0.14
single/24k	3.97±0.15	2.21±0.15	3.47±0.15
single/48k	4.43±0.14	1.95±0.15	3.80±0.16
multi/16k	3.11±0.14	2.41±0.13	2.35±0.12
multi/24k	4.25±0.13	2.50±0.15	3.06±0.15
multi/48k	4.55±0.10	1.84±0.12	3.29±0.15

included in $[0, f_1/2]$ in most of the frames. Hereafter, we denote the experimental conditions for generating 48 kHz speech using the proposed method as **MSR-PWG-48k**. All $\{\phi_i\}_{i=1}^7$ have identical structures of a ten-layer, one-stack non-causal WaveNet with dilation set to 1, 2, 4, ..., 512. Following the experimental conditions of PWG, the number of channels for the residual block and skip connection was set to 64, and the size of the convolutional filter was set to three. Upsampling of waveforms and conditioning features was conducted in a differentiable manner using torchaudio [32]. All discriminators were constructed with the same structure as that of the PWG, that is, ten-layer non-causal dilated convolutions with a leaky ReLU activation function ($\alpha = 0.2$).

The parameters of the MR-STFT loss were set to different values for each sampling rate. For $f_7 = 48$ kHz, we set the frame length to $\{2048, 4096, 1024\}$ points, the window length to $\{1200, 2400, 480\}$ points, and the frame shift to $\{240, 480, 100\}$ points. For $\{f_i\}_{i=1}^6$, these parameters were multiplied by f_i/f_7 . The weight of adversarial loss λ_{adv} was set to 1.0. A mini-batch was constructed by randomly clipping speech (0.5 s) from eight utterances for each training step. Training was conducted for 400,000 steps using RAdam optimizer ($\epsilon = 10^{-6}$). Only the generators were optimized in the first 200,000 steps, and the generators and discriminators were jointly optimized in the following 200,000 steps. The initial learning rates were set to 10^{-3} for both the generator and discriminator, and they were reduced by a factor of 0.5, after 300,000 steps of training.

As baseline models, we trained the original PWG with 16, 24, and 48 kHz speech waveforms, respectively (hereinafter referred to as **PWG- $\{16, 24, 48\}$ k**)². To extract log mel spectrograms, the frame and window lengths were set to 512, 1,024, and 2,048 points for **PWG-16k**, **PWG-24k**, and **PWG-48k**, respectively. The frame shift was set to 5 ms for all settings. For **PWG-16k** and **PWG-24k**, the parameters of the MR-STFT loss are the same as in the previous study [12], and they were doubled for **PWG-48k**.

4.2. Evaluation

4.2.1. Comparison to baseline in terms of quality and speed

A mean opinion score (MOS) test was conducted to evaluate the subjective quality of synthetic speech. In both of the single-speaker and multi-speaker experiments, each rater evaluated ten sets of 9 speech samples, for a total of 90 samples, where each set consisted of a combination of three sampling rates (16 kHz, 24 kHz, and 48 kHz) and three methods (reference: **Ref**, conventional method: **PWG**, and proposed method: **MSR-PWG**)³

²<https://github.com/kan-bayashi/ParallelWaveGAN> was used.

³Speech samples are available at the following URL: <https://rinnakk.github.io/research/publications/MSR-NV>.

Table 2: Real-time factor of inference averaged over 300 utterances.

Method	16 kHz	24 kHz	48 kHz
PWG	0.027	0.039	0.074
MSR-PWG	0.027	0.042	0.068

Speech samples were evaluated in terms of closeness to the quality of the 48 kHz reference samples using a five-point scale from 1 (very far) to 5 (very close). The raters should be able to distinguish between 16, 24, and 48 kHz speech samples; however, depending on the playback equipment and the subject’s hearing, it may not be possible to distinguish between these different sampling rates. Thus, we conducted an internal preliminary test and asked twenty raters who were able to distinguish between them to participate in the evaluation.

The results are presented in Table 1. Note that the scores of **PWG** are lower than those of the original paper [12] owing to the differences in the experimental condition and evaluation criterion. In the single-speaker experiment, **MSR-PWG** achieved higher scores than **PWG** at any sampling rate. The score of **PWG-48k** was lower than that of **PWG-24k**, even though it was closer to the reference speech in terms of the sampling rate. This is because **PWG** had difficulty generating waveforms with a high sampling rate, resulting in quality degradation. However, **MSR-PWG** achieved a high score even at 48 kHz, which verifies the effectiveness of stepwise waveform generation starting from a low sampling rate. In the multi-speaker setting, the scores of **MSR-PWG** were comparable to **PWG** at 16 kHz and significantly higher at 24 kHz and 48 kHz. However, there was a gap between the scores of **Ref** and those of **MSR-PWG**. In particular, the speech quality was degraded for male speakers, whose amount of data was relatively small in this experiment. Improving the performance of multi-speaker multi-style modeling is a future challenge for the proposed method.

In terms of model size, the numbers of parameters were 1.44M for **PWG** and 3.06M for **MSR-PWG**. When we expanded **PWG** to 70 layers to align the number of its parameters with **MSR-PWG**, the training and inference time increased significantly without any quality improvement. Thus, we conclude that not the increased number of parameters but the proposed structure improved the speech quality.

We also measured the training and inference speed of the baseline and proposed methods using NVIDIA Tesla P40. The training required 103 h for **PWG-16k**, 153 h for **PWG-24k**, 252 h for **PWG-48k**, and 132 h for **MSR-PWG**. Although the proposed method handled up to 48 kHz waveform generation, it achieved high-quality synthesis with a training time shorter than that of **PWG-24k**. For inference, the real-time factor (RTF), which is the time required to generate a waveform in one second, was measured, and the results obtained are presented in Table 2. Although the generator of the proposed method consists of 70 layers, compared to 30 in the baseline, the inference time did not show an increase. This is because the input length is proportional to f_i , which is especially short in the early stages of the model, and the computational cost is reduced.

4.2.2. Training data amount and synthesis quality

Six models were trained with single-speaker data of 1, 3, 5, 10, 30 min, and 8 h (full data). The quality of the 48 kHz synthetic speech was evaluated using the MOS test in the same manner as in section 4.2.1. The results are presented in Table 3. Although the score was significantly low when the amount of training data

Table 3: Comparison of MOS for various training data amount with 95% confidence intervals.

Training data amount	MOS
1 min	1.49±0.13
3 min	3.07±0.15
5 min	3.43±0.16
10 min	3.50±0.15
30 min	3.55±0.16
8 h (full data)	3.75±0.16

Table 4: Comparison of MOS for different training sets with 95% confidence intervals.

Training set	MOS
1min	1.76±0.14
3×1min	2.67±0.15
3min	3.32±0.17

was only 1 min, it greatly improved when 5 min of training data was used. This result indicates that the proposed method can generate speech of adequate quality even with 5 min of training data.

4.2.3. Use of speech data with low sampling rates

We investigated whether the quality could be improved using 16 and 24 kHz speech for training when there is little 48 kHz speech available. The following three conditions were evaluated in the single-speaker setting using the MOS test similarly to section 4.2.1: **1min**: 1 min of 48 kHz speech was used for training, **3×1min**: 1 min of different 16, 24, and 48 kHz speech, for 3 min, was used for training, and **3min**: 3 min of 48 kHz speech was used for training. The results are presented in Table 4. **3×1min** showed a significantly higher score than **1min**, confirming that the quality of synthetic speech can be improved using speech with a lower sampling rate for training. Although we used only one minute each of 16 and 24 kHz speech here, the synthesis quality of 48 kHz speech could be further improved by increasing the amount of these lower sampling rate speech data. However, because 16 and 24 kHz speech do not contain any components above 8 and 12 kHz, respectively, only a part of the network can be trained with these data. This is why the score of **3×1 min** was lower than that of **3 min**, which was trained using the same amount of 48 kHz speech.

5. Conclusions

In this study, we proposed MSR-NV, a method to handle multiple sampling rates using a single neural vocoder. Experimental evaluations using a PWG-based structure demonstrated that the proposed method could generate high-quality waveforms at multiple sampling rates, including 48 kHz, while maintaining a fast generation speed. We also demonstrated the data efficiency of the proposed method by varying the training data amount and evaluating the speech quality. In addition, we could improve speech quality using 16 and 24 kHz speech when the amount of 48 kHz speech was limited. The proposed MSR-NV eliminates the need to re-train the model for different sampling rates, which broadens the range of applications.

Future work includes verifying the effectiveness of the proposed method when model structures other than PWG are used. It is also necessary to improve the speech quality for multiple speakers and speaking styles to realize a more versatile vocoder.

6. References

- [1] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “WaveNet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, Sep. 2016.
- [2] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, “Speaker-dependent WaveNet vocoder,” in *Proc. INTERSPEECH*, Stockholm, Sweden, Aug. 2017, pp. 1118–1122.
- [3] H. Kawahara, M. Morise, T. Takahashi, R. Nisimura, T. Irino, and H. Banno, “Tandem-STRAIGHT: A temporally stable power spectral representation for periodic signals and applications to interference-free spectrum, f0, and aperiodicity estimation,” in *Proc. ICASSP*, Las Vegas, U.S.A., Apr. 2008, pp. 3933–3936.
- [4] M. Morise, F. Yokomori, and K. Ozawa, “WORLD: a vocoder-based high-quality speech synthesis system for real-time applications,” *IEICE Transactions on Information and Systems*, vol. 99, no. 7, pp. 1877–1884, Jul. 2016.
- [5] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, “Improved variational inference with inverse autoregressive flow,” in *Proc. NIPS*, vol. 29, Barcelona, Spain, Dec. 2016, pp. 4743–4751.
- [6] A. Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu *et al.*, “Parallel WaveNet: Fast high-fidelity speech synthesis,” in *Proc. ICML*, Stockholm, Sweden, Jul. 2018, pp. 3918–3926.
- [7] W. Ping, K. Peng, and J. Chen, “Clarinet: Parallel wave generation in end-to-end text-to-speech,” in *Proc. ICLR*, New Orleans, U.S.A., May 2019.
- [8] D. P. Kingma and P. Dhariwal, “Glow: Generative flow with invertible 1x1 convolutions,” in *Proc. NeurIPS*, vol. 31, Barcelona, Spain, Dec. 2018, pp. 10 236–10 245.
- [9] R. Prenger, R. Valle, and B. Catanzaro, “WaveGlow: A flow-based generative network for speech synthesis,” in *Proc. ICASSP*, Brighton, U.K., May 2019, pp. 3617–3621.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Proc. NIPS*, vol. 27, Montreal, Canada, Dec. 2014, pp. 2672–2680.
- [11] K. Kumar, R. Kumar, T. de Boissiere, L. Gestein, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. Courville, “MelGAN: Generative adversarial networks for conditional waveform synthesis,” in *Proc. NeurIPS*, vol. 32, Vancouver, Canada, Dec. 2019, pp. 14 881–14 892.
- [12] 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 *Proc. ICASSP*, online, May 2020, pp. 6199–6203.
- [13] J. Yang, J. Lee, Y. Kim, H. Cho, and I. Kim, “VocGAN: A high-fidelity real-time vocoder with a hierarchically-nested adversarial network,” in *Proc. INTERSPEECH*, online, Oct. 2020, pp. 200–204.
- [14] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” in *Proc. NeurIPS*, vol. 33, online, Dec. 2020, pp. 17 022–17 033.
- [15] J. Ho, A. Jain, and P. Abbeel, “Denosing diffusion probabilistic models,” in *Proc. NeurIPS*, vol. 33, online, Dec. 2020, pp. 6840–6851.
- [16] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, “WaveGrad: Estimating gradients for waveform generation,” in *Proc. ICLR*, online, May 2021.
- [17] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, “DiffWave: A versatile diffusion model for audio synthesis,” in *Proc. ICLR*, online, May 2021.
- [18] X. Wang, S. Takaki, and J. Yamagishi, “Neural source-filter waveform models for statistical parametric speech synthesis,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 402–415, Nov. 2019.
- [19] J.-M. Valin and J. Skoglund, “LPCNet: Improving neural speech synthesis through linear prediction,” in *Proc. ICASSP*, Brighton, U.K., May 2019, pp. 5891–5895.
- [20] K. Matsubara, T. Okamoto, R. Takashima, T. Takiguchi, T. Toda, Y. Shiga, and H. Kawai, “Full-Band LPCNet: A real-time neural vocoder for 48 kHz audio with a CPU,” *IEEE Access*, vol. 9, pp. 94 923–94 933, Jun. 2021.
- [21] Y. Hono, S. Takaki, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, “PeriodNet: A non-autoregressive waveform generation model with a structure separating periodic and aperiodic components,” in *Proc. ICASSP*, online, Jun. 2021, pp. 6049–6053.
- [22] 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*, Graz, Austria, Sep. 2019, pp. 1526–1530.
- [23] T. Karras, T. Aila, S. Laine, and J. Lehtinen, “Progressive growing of GANs for improved quality, stability, and variation,” in *Proc. ICLR*, Vancouver, Canada, May 2018.
- [24] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *Proc. CVPR*, California, U.S.A., Jun. 2019, pp. 4401–4410.
- [25] K. Oura, K. Nakamura, K. Hashimoto, Y. Nankaku, and K. Tokuda, “Deep neural network based real-time speech vocoder with periodic and aperiodic inputs,” in *Proc. SSW10*, vol. 32, Vienna, Austria, Sep. 2019, pp. 13–18.
- [26] Y.-C. Wu, T. Hayashi, T. Okamoto, H. Kawai, and T. Toda, “Quasi-periodic Parallel WaveGAN: A non-autoregressive raw waveform generative model with pitch-dependent dilated convolutional neural network,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 792–806, Jan. 2021.
- [27] E. Song, R. Yamamoto, M.-J. Hwang, J.-S. Kim, O. Kwon, and J.-M. Kim, “Improved Parallel WaveGAN vocoder with perceptually weighted spectrogram loss,” in *Proc. SLT*, online, Jan. 2021, pp. 470–476.
- [28] R. Yamamoto, E. Song, M.-J. Hwang, and J.-M. Kim, “Parallel waveform synthesis based on generative adversarial networks with voicing-aware conditional discriminators,” in *Proc. ICASSP*, online, Jun. 2021, pp. 6039–6043.
- [29] K. Mizuta, T. Koriyama, and H. Saruwatari, “Harmonic WaveGAN: GAN-based speech waveform generation model with harmonic structure discriminator,” in *Proc. INTERSPEECH*, online, Sep. 2021, pp. 2192–2196.
- [30] M.-J. Hwang, R. Yamamoto, E. Song, and J.-M. Kim, “High-fidelity Parallel WaveGAN with multi-band harmonic-plus-noise model,” in *Proc. INTERSPEECH*, online, Sep. 2021, pp. 2227–2231.
- [31] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, “High-resolution image synthesis and semantic manipulation with conditional GANs,” in *Proc. CVPR*, Salt Lake City, U.S.A., Jun. 2018, pp. 8798–8807.
- [32] Y.-Y. Yang, M. Hira, Z. Ni, A. Chourdia, A. Astafurov, C. Chen, C.-F. Yeh, C. Puhusch, D. Pollack, D. Genzel *et al.*, “TorchAudio: Building blocks for audio and speech processing,” *arXiv preprint arXiv:2110.15018*, Oct. 2021.