



FastFit: Towards Real-Time Iterative Neural Vocoder by Replacing U-Net Encoder With Multiple STFTs

Won Jang, Dan Lim, Heayoung Park

Kakao Enterprise Corporation, Republic of Korea

{taylor.martin, satoshi.2020, abigail.p}@kakaocommerce.com

Abstract

This paper presents FastFit, a novel neural vocoder architecture that replaces the U-Net encoder with multiple short-time Fourier transforms (STFTs) to achieve faster generation rates without sacrificing sample quality. We replaced each encoder block with an STFT, with parameters equal to the temporal resolution of each decoder block, leading to the skip connection. FastFit reduces the number of parameters and the generation time of the model by almost half while maintaining high fidelity. Through objective and subjective evaluations, we demonstrated that the proposed model achieves nearly twice the generation speed of baseline iteration-based vocoders while maintaining high sound quality. We further showed that FastFit produces sound qualities similar to those of other baselines in text-to-speech evaluation scenarios, including multi-speaker and zero-shot text-to-speech.

Index Terms: Neural vocoder, text-to-speech, U-Net, short-time Fourier transform

1. Introduction

Neural vocoders generate speech that conforms to the given input conditions by modeling short- and long-term dependencies. Owing to these features, these architectures have been applied[1, 2], wholly or partially, to various applications[3, 4, 5] that output speech and audio as well as text-to-speech applications[6, 7, 8, 9]. Moreover, the use of generative adversarial networks (GANs)[10] for neural waveform generation has further improved neural vocoders[2, 11, 12, 13, 14, 15]. However, according to recent text-to-speech studies, some vocoders require additional training (i.e., fine-tuning) using pairs of ground-truth waveforms and model-predicted features to adapt to low-quality audio features generated by an acoustic model[4, 12, 13, 14].

Recent research has shown that image generation models utilizing denoising diffusion probabilistic models (DDPMs)[16] outperform traditional GAN-based models[17]. Several studies have successfully applied DDPM to neural vocoders, with some reporting superior performance over conventional models[18, 19, 20, 21]. However, the trade-off between generation speed and quality owing to the need for repeated denoising is considered a barrier to the commercialization of these models. Subsequent studies have attempted to overcome this by maintaining robust performance with fewer iterations[20, 22, 23].

The symmetric architecture of U-Net[24] has made it an attractive choice for iteration-style models. To use an existing GAN-based vocoder as a decoder, some studies have added an encoder connected with skip connections[18, 20]. However, this doubles the size of the model and slows the generation speed by approximately half.

To improve efficiency, we propose FastFit, a new architecture that replaces the encoders in U-Net with multiple short-time Fourier transforms (STFTs) to trade a small fidelity degradation for a high generation speed gain. Our work is inspired by the work of Kaneko *et al.* [14], who replaced some of the blocks with an inverse STFT. We extended the GAN-based vocoder proposed by Jang *et al.* [13] to U-Net and replaced each encoder block with an STFT with parameters corresponding to the shape of its skip connection. This modification of the model preserves the advantages of skip connection in U-Net, while expecting more efficient intermediate feature encoding of raw waveforms because the computational cost of STFT is less than that of the neural encoder.

We applied the iteration-style principle proposed by Koizumi *et al.* [23] to the proposed architecture. To compare the performance of FastFit, we used iteration-based vocoders, which reported fast generation speeds as baselines and conducted objective and subjective evaluations. The results showed that FastFit achieved twice the generation speed with statistically similar speech quality despite having half the baseline parameters. Further, we conducted experiments by applying each vocoder to multi-speaker and zero-shot text-to-speech without fine-tuning, and FastFit was found to be one of the best-performing models¹.

2. Related work

The proposed model is influenced by several improvements from previous studies on GAN-based vocoders. Parallel WaveGAN[11] applied multi-resolution STFT (MR-STFT) loss as an auxiliary loss to a vocoder to facilitate stable adversarial training. HiFi-GAN[12] and UnivNet[13] include a multi-period discriminator (MPD) and multi-resolution spectrogram discriminator (MRSD), respectively, which are discriminators that can observe real or generated waveforms with various patterns and resolutions. iSTFTNet[14] replaced the back part of the residual blocks of HiFi-GAN with an inverse short-time Fourier transform (iSTFT), trading off a small reduction in quality for higher synthesis speeds. Recently, BigVGAN[15] succeeded in the adversarial training of large-scale generators with more than 100M parameters, achieving overall state-of-the-art fidelity, including out-of-distribution robustness.

Our work is based on the iteration-style vocoding principle proposed by WaveFit[23]. According to the fixed-point iteration theorem, if a mapping \mathcal{T} has a fixed point $\mathbf{x} = \mathcal{T}(\mathbf{x})$ and is firmly quasi-nonexpansive (as described in Section 17.2.2 in Yamada *et al.* [26]), then the mapping point $\mathcal{T}(\mathbf{y})$ of an arbitrary point \mathbf{y} always has a smaller Euclidean distance from \mathbf{x}

¹Audio demo samples can be found at the following URL:
<https://kallavinka8045.github.io/is2023/>

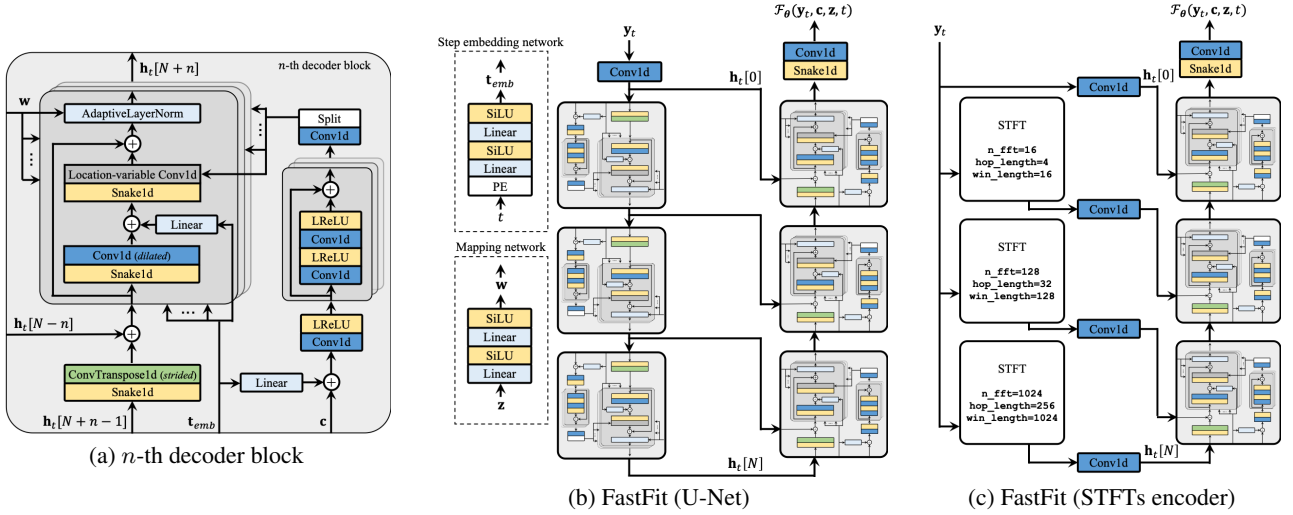


Figure 1: *FastFit* architecture. (a) The n -th decoder block. For example, the 1-st decoder block computes $\mathbf{h}_t[N + 1]$ from the output $\mathbf{h}_t[N]$ of the last N -th encoder block and the skip-connection $\mathbf{h}_t[N - 1]$, the output of the $N - 1$ -th encoder. (b) *FastFit* based on U-Net. \mathbf{c} , \mathbf{w} , and \mathbf{t}_{emb} are used as inputs to each block, but we omitted them in this figure for brevity. PE denotes a positional encoding operation[25]. (c) *FastFit* with multiple STFTs encoder, based on the proposed U-Net version. Each channel size of STFT is converted to fit the channel of the decoder block through each convolution layer.

than \mathbf{y} . \mathcal{T} can be extended to the form of iterative denoising $\mathbf{y}_{t-1} = \mathcal{T}(\mathbf{y}_t)$. If an arbitrary initial point \mathbf{y}_T is iteratively refined at each t from T to 1, then \mathbf{y}_{t-1} always moves closer than \mathbf{y}_t to the clean signal \mathbf{x} , which is the fixed point of \mathcal{T} . WaveFit proposed a denoising mapping and loss function for a vocoder that satisfies this property.

3. Description of the proposed model

The proposed model, *FastFit*, begins with an initial point \mathbf{y}_T . At each iteration step, $t = T, T - 1, \dots, 1$, denoising mapping is applied to \mathbf{y}_t to obtain the denoised signal \mathbf{y}_{t-1} . A model \mathcal{F} parameterized by θ was trained to predict the noise components of \mathbf{y}_t . \mathcal{F}_θ was conditioned on the log-mel-spectrogram \mathbf{c} , latent noise $\mathbf{z} \sim \mathcal{N}(0, \mathbf{I})$, and current step t as $\mathcal{F}_\theta(\mathbf{y}_t, \mathbf{c}, \mathbf{z}, t)$. The objective of the vocoder is to make \mathbf{y}_t at each iteration, including the final output \mathbf{y}_0 , close to the target waveform, \mathbf{x} .

3.1. Improving the architecture of the residual block

Our U-Net model has N encoder and decoder blocks with mapping and step embedding networks for the intermediate latent \mathbf{w} and step embedding \mathbf{t}_{emb} , respectively, as shown in Figure 1(b). Each n -th decoder block computes $\mathbf{h}_t[N + n]$ with $\mathbf{h}_t[N + n - 1]$, \mathbf{c} , \mathbf{w} , \mathbf{t}_{emb} , and $\mathbf{h}_t[N - n]$ as inputs, as shown in Figure 1(a). \mathbf{t}_{emb} is conditioned to be broadcast and added to the following features after \mathbf{c} and after each dilated convolution.

Each decoder block is based on the UnivNet[13] generator with three main changes. First, we added an adaptive layer normalization (AdaLN) after each residual connection to inject noise \mathbf{z} into the vocoder for improving the performance and training stability[27]. Second, we applied the snake activation function[15] to the model. This trainable activation function controls the output of each layer in the form of a periodic frequency and contributes to out-of-distribution robustness[15]. Finally, the gated activation units were removed to improve the generation speed. Although this layer contributes to a slight improvement in quality according to UnivNet, it doubles the

number of channels in the previous layer. The encoder block is like the decoder block, with a few differences: there is no skip connection, and the upsampling layer is replaced with downsampling using strided convolution layer.

3.2. Replacing U-Net encoder with multiple STFTs

To approach a real-time iterative vocoder application, we propose an intuitive methodology: replacing the U-Net encoder with multiple STFTs. As shown in Figure 1(c), we used a frame shift interval equal to the temporal resolution of each decoder block, leading to the skip connection as a parameter for each STFT. Inspired by iSTFTNet[14], the number of points in the Fourier transform and Hann window length was set to four times the respective frame shift interval. To match the channel size, a convolution layer was placed between each encoder and decoder blocks. Because the computation speed of STFT is high compared to that of the neural encoder block, our proposed model could reduce the number of parameters by almost half, approximately doubling the generation speed. We expect limited degradation of speech quality using the methodology because skip connections are still used, which is the basis for the high performance of the U-Net architecture.

We conducted an ablation study to determine the optimal representations of STFT. Consequently, we chose the Cartesian form (concatenation of real and imaginary channels) as the STFT representation.

3.3. Denoising mapping and training losses

As mentioned in the previous section, our model is based on denoising mapping and the loss function proposed using WaveFit. The denoised signal \mathbf{y}_{t-1} is then computed as follows:

$$\tilde{\mathbf{y}}_t = \mathbf{y}_t - \mathcal{F}_\theta(\mathbf{y}_t, \mathbf{c}, \mathbf{z}, t) \quad (1)$$

$$\mathbf{y}_{t-1} = (P_c / (P_{\tilde{\mathbf{y}}_t} + s)) \tilde{\mathbf{y}}_t \quad (2)$$

where $s = 10^{-8}$ is a constant used to avoid numerical errors.

The denoising mapping is defined by subtracting the noise component predicted by \mathcal{F}_θ from \mathbf{y}_t to obtain $\tilde{\mathbf{y}}_t$ and adjusting the power of $\tilde{\mathbf{y}}_t$ to P_c . $P_{\tilde{\mathbf{y}}_t}$ and P_c can be obtained by computing the power spectrograms of $\tilde{\mathbf{y}}_t$ and \mathbf{c} , respectively, and then taking the element-wise mean. Specifically, the power spectrogram of \mathbf{c} can be obtained by multiplying \mathbf{c} with the pseudoinverse of the mel-compression matrix and then squaring it. By scaling the power of the signal to a constant power of \mathbf{c} at each step, the power of \mathbf{y}_{t-1} can be kept constant until denoising is repeated for all t and the final output \mathbf{y}_0 is obtained.

FastFit is adversarially trained with the least squares GAN (LSGAN)[28] as the GAN loss and discriminators D , which are a combination of MPD (as described in Appendix B.2 in Kim *et al.* [9]) and MRSD[13]. The overall losses $\mathcal{L}_{\text{disc}}$ and \mathcal{L}_{gen} are defined as follows.

$$\mathcal{L}_{\text{disc}} = \frac{1}{TK} \sum_{t=0}^{T-1} \sum_{k=0}^{K-1} \left[\mathbb{E}_{\mathbf{x}} [(D_k(\mathbf{x}) - 1)^2] + \mathbb{E}_{\mathbf{y}_t} [D_k(\mathbf{y}_t)^2] \right] \quad (3)$$

$$\begin{aligned} \mathcal{L}_{\text{gen}} = & \frac{1}{T} \sum_{t=0}^{T-1} \left[\lambda_{\text{aux}} \mathcal{L}_{\text{aux}}(\mathbf{y}_t, \mathbf{x}) \right. \\ & \left. + \frac{1}{K} \sum_{k=0}^{K-1} \left[\mathbb{E}_{\mathbf{y}_t} [(D_k(\mathbf{y}_t) - 1)^2] + \lambda_{\text{fm}} \mathcal{L}_{\text{fm}}(D_k; \mathbf{y}_t, \mathbf{x}) \right] \right] \quad (4) \end{aligned}$$

where K denotes the number of sub-discriminators. We used MR-STFT[11] as the auxiliary loss \mathcal{L}_{aux} and set λ_{aux} to 2.5. Additionally, we applied the scaled feature matching loss \mathcal{L}_{fm} proposed by Yang *et al.* [29]; $\lambda_{\text{fm}} = \lambda_{\text{aux}} \mathcal{L}_{\text{aux}} / \mathcal{L}_{\text{fm}}$.

According to WaveFit, the initial point \mathbf{y}_T is sampled using the noise generation algorithm of SpecGrad[21], which is defined as follows:

$$\mathbf{y}_T = \mathbf{G}^+ \mathbf{M} \mathbf{G} \boldsymbol{\epsilon} \quad (5)$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$, \mathbf{G} and \mathbf{G}^+ denote STFT and iSTFT, respectively, and \mathbf{M} denotes a filter computed from \mathbf{c} for prior adaptation. SpecGrad estimated a cepstrum-based spectral envelope from the spectrogram obtained by multiplying \mathbf{c} with the pseudoinverse of the mel-compression matrix (like the previous paragraph) and used it as \mathbf{M} . However, SpecGrad reported that using the spectrogram from \mathbf{c} directly did not provide satisfactory results. In contrast, we experimentally verified that, unlike the case of the DDPM algorithm on which SpecGrad is based, using the spectrogram as \mathbf{M} yields higher sound quality under our denoising mapping and architecture. The experimental design and results are described in detail in the following sections.

4. Experiments

4.1. Data configurations and evaluation metrics

We adopted LibriTTS[30], a multi-speaker English dataset with 24 kHz sampling rate waveforms, for the training and evaluation of the vocoder models. We used the “train-clean-360” dataset to train the models, with 5% and 2% of the dataset for validation and testing, respectively, with all speakers included in each of the three splits. For ground-truth mel-spectrogram evaluation (GT mel evaluation) including the ablation study, a “test-clean” dataset was prepared. The STFT parameters used to extract the 100-band, 0-12 kHz log-mel-spectrograms are 1024-point Fourier transform, 256 sample frame shift, and 1024 sample Hann window length.

Two objective evaluation metrics, PESQ and MR-STFT, were used to evaluate the performance of each model. An open-source library² was used to calculate the wideband PESQ, and the parameters required to calculate the MR-STFT metric were set to the same values as those in Yamamoto *et al.* [11].

To clarify the comparison between the proposed model and baselines, a 5-point mean opinion score (MOS) was used for the TTS evaluation and a 7-point comparative MOS (CMOS) evaluation was used for GT mel evaluation. To collect 400 ratings for each evaluation item, we randomly sampled 20 speech samples and collected 20 ratings for each sample from 20 listeners located in the United States using Amazon Mechanical Turk. The loudness of all speech samples used was normalized to -23 LUFS. Other details of the subjective evaluations were based on Loizou’s work[31].

4.2. Model settings

The proposed model uses the hyperparameters of each block that follow UnivNet-c32[13], with the number of dilated convolutions reduced to three, each with a dilation of {1,3,9}, to improve speed. The channel size of each convolution in the MRSD was set to 16. The dimensions of the latent noise \mathbf{z} and the number of iterations T were set to 100 and 3, respectively. The step embedding network used the structure proposed by Kong *et al.* [19], and the mapping network used the same structure, but the channel size of each layer was set to 256. The minimum phase response based on the homomorphic filter method was used to calculate the filter \mathbf{M} . The same optimizer and learning rate as UnivNet were used to train FastFit and all the models for the ablation study up to 1M steps, with a batch size of 64. All other architectural details followed the settings of the studies on which they were based.

The performance of the proposed model was compared with three baselines: UnivNet, FastDiff, and WaveFit. These models are based on three main methodologies of vocoder research: GAN, DDPM, and fixed-point iteration. For UnivNet, we used the “c32” version of our implementation. FastDiff was implemented using the official repository³, with $T = 4$ as suggested by Huang *et al.* [20]. We implemented WaveFit with $T = 3$ following Koizumi *et al.* [23], using an unofficial implementation⁴ of WaveGrad with 15.8M parameters as the base model. The upsampling ratios were set to {4,4,4,2,2} to fit our experimental setting. To improve the training stability, we replaced the GAN loss with an LSGAN, which resulted in more stable training and a lower auxiliary loss. All the models were trained up to 1M steps using four NVIDIA V100 GPUs, and no additional fine-tuning was applied.

For the multi-speaker TTS evaluation, we trained the JDI-T[7] acoustic model with the adapter-based multi-speaker methodology of Hsieh *et al.* [32] using the LibriTTS train-clean-360 subset with 100 speakers. For the zero-shot TTS evaluation, we used an open-source zero-shot TTS program named TorToiSe⁵. The recordings of the LibriTTS “test-clean” subset with 10 speakers were input into the program with an “ultrafast” offset to synthesize mel-spectrograms for evaluation. Sentences for all TTS evaluations were extracted from the utterances of speakers who were not included in the training. All speaker selection and data processing details for TTS evaluations followed Hsieh *et al.*’s approach.

²<https://github.com/vBaiCai/python-pesq>

³<https://github.com/Rongjiehuang/FastDiff>

⁴<https://github.com/ivanvovk/WaveGrad>

⁵<https://github.com/neonbjb/tortoise-ts>

Table 1: *The ablation study results.*

Model	PESQ \uparrow	MR-STFT \downarrow	CMOS \uparrow
Recordings	-	-	0.251
FastFit	3.712	0.866	-
FastFit (U-Net)	3.754	0.868	0.062
Without AdaLN	3.411	0.974	-0.168
Without skip-connections	3.449	0.936	-0.199
$y_T \sim$ Spectral envelope	3.422	0.969	-0.175
$y_T \sim$ Griffin-Lim	3.685	0.872	0.069
Magnitude STFTs encoder	3.677	0.875	-0.031
Polar STFTs encoder		Failed to train	
Polar+Cartesian STFTs encoder		Failed to train	

5. Results

5.1. Ablation studies

We conducted ablation studies to evaluate the proposed improvements under the GT mel conditions. According to the results presented in Table 1, FastFit achieved the non-significant objective and subjective metrics compared with U-Net with a neural encoder. We observed significant destabilization of the training when the model was trained without AdaLN, resulting in worse metrics. To demonstrate that the proposed methodology maintains quality by maintaining skip connections, we removed all skip connections and connected only one STFT to the first decoder block, which resulted in relatively poor metrics.

Moreover, we have made several attempts to define an effective initial point y_T . Our experiments computing y_T from the spectral envelope performed worse than the proposed method of computing y_T directly from the spectrogram from c. We also attempted to convert the spectrogram to a waveform with 32 Griffin-Lim iterations and use it as y_T . This did not produce significant differences in metrics but resulted in a slight reduction in generation speed, so we did not adopt it.

We were intrigued by Webber *et al.* [33]’s work with different representations of STFT output and tested to find the optimal one. The magnitude spectrogram, used as a representation, showed no significant metric difference from the Cartesian form of the proposed model. However, internal tests showed a decrease in quality in some samples; therefore, we did not adopt it. For the polar form, training collapsed early on, so using phase as a representation was not appropriate for the proposed model.

5.2. Comparison with baselines

To measure the speed of the evaluation models, we generated 6 second segments 20 times using an NVIDIA V100 GPU and measured the average time. As shown in Table 2, Fast-

Fit achieved approximately twice the synthesis speed of the iteration-based vocoders despite having approximately half the number of parameters, and none of the metrics scored significantly worse in terms of speech quality. Although UnivNet’s synthesis speed was superior to other baselines, it performed poorly in the CMOS evaluation owing to the occasional blurring of the harmonic component of the GT mel-spectrogram. FastDiff performed best on PESQ but worst on MR-STFT, which calculates the numerical distances. This is because the model confused segments with noise-like spectral shapes, such as consonants, breaths, and high-frequency spectral bands, with noise and denoised them. WaveFit performed well overall but had the slowest generation speed.

5.3. Application to text-to-speech synthesis

We characterized the mel-spectrograms generated by the models for two TTS evaluation tasks and found that the mel-spectrograms for multi-speaker TTS had relatively blurry shapes because they were not generated by high-quality models, such as GAN or DDPM. In contrast, the mel-spectrograms for zero-shot TTS had relatively more realistic shapes because they used a DDPM-based acoustic model that produced high-quality output regardless of the synthesis rate. However, each vocoder was not fine-tuned using the predicted mel-spectrograms.

For multi-speaker TTS, FastFit produced better MOS scores than UnivNet and FastDiff, with a slight difference from WaveFit. For the zero-shot TTS, all models except FastDiff produced similar MOS. UnivNet was observed to produce noise artifacts in some segments of the multi-speaker TTS experiment, which may be responsible for its worse MOS. FastDiff recorded the worst MOS owing to the incorrect denoising of noise-like components and produced blurred harmonic components. The remaining models produced statistically similar MOS, with FastFit (U-Net) producing the best results, but the confidence intervals of the MOS overlapped.

6. Conclusion

By improving the architecture of an iteration-based neural vocoder, we could double the generation rate while maintaining a high fidelity. As the U-Net architecture is widely used in speech processing applications, we expect our simple yet effective idea of replacing the encoder with STFTs to be applied in a variety of speech-based research and applications in the future.

7. Acknowledgements

The authors would like to thank James Bekter for developing TorToiSe, an outstanding open-sourced TTS model. We also would like to thank Jaesam Yoon, Sunghye Jung, Gyeonghwan O and Bongwan Kim for providing insightful feedback.

Table 2: *Results of comparison with baseline models. “Speed” indicates each model’s generation speed relative to real time.*

Model	Model complexity		GT mel evaluation			Multi-speaker TTS	Zero-shot TTS
	Params \downarrow	Speed \uparrow	PESQ \uparrow	MR-STFT \downarrow	CMOS \uparrow	MOS \uparrow	MOS \uparrow
UnivNet	14.86M	$\times 314.49$	3.705	0.853	-0.295	3.46 \pm 0.06	3.84 \pm 0.09
FastDiff	15.36M	$\times 52.35$	3.786	1.385	0.078	3.32 \pm 0.07	3.68 \pm 0.10
WaveFit	15.85M	$\times 43.10$	3.639	0.921	0.072	3.70 \pm 0.08	3.83 \pm 0.09
FastFit	6.81M	$\times 101.40$	3.712	0.866	-	3.67 \pm 0.08	3.86 \pm 0.08
FastFit (U-Net)	12.94M	$\times 59.88$	3.754	0.868	0.062	3.75\pm0.07	3.90\pm0.09
Recordings	-	-	-	-	0.251	-	-

8. References

- [1] A. van den 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," in *ISCA Speech Synthesis Workshop (SSW)*, 2016.
- [2] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. Courville, "MelGAN: Generative Adversarial Networks for Conditional Waveform Synthesis," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [3] D. Stoller, S. Ewert, and S. Dixon, "Wave-U-Net: A Multi-Scale Neural Network for End-to-End Audio Source Separation," in *International Society for Music Information Retrieval Conference (ISMIR)*, 2018.
- [4] A. Lee, P.-J. Chen, C. Wang, J. Gu, S. Popuri, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang *et al.*, "Direct Speech-to-Speech Translation With Discrete Units," in *Proc. Annual Meeting of the Association for Computational Linguistics (ACL)*, 2022.
- [5] Y. Zhang, J. Cong, H. Xue, L. Xie, P. Zhu, and M. Bi, "VISinger: Variational Inference with Adversarial Learning for End-to-End Singing Voice Synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [6] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan *et al.*, "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [7] D. Lim, W. Jang, G. O. H. Park, B. Kim, and J. Yoon, "JDI-T: Jointly trained Duration Informed Transformer for Text-To-Speech without Explicit Alignment," in *Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2020.
- [8] Y. Ren, C. Hu, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fast-Speech 2: Fast and High-Quality End-to-End Text-to-Speech," in *International Conference on Learning Representations (ICLR)*, 2021.
- [9] J. Kim, J. Kong, and J. Son, "Conditional Variational Autoencoder with Adversarial Learning for End-to-End Text-to-Speech," in *International Conference on Machine Learning (ICML)*, 2021.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative Adversarial Nets," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [11] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [12] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [13] W. Jang, D. Lim, J. Yoon, B. Kim, and J. Kim, "UnivNet: A Neural Vocoder with Multi-Resolution Spectrogram Discriminators for High-Fidelity Waveform Generation," in *Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2021.
- [14] T. Kaneko, K. Tanaka, H. Kameoka, and S. Seki, "iSTFTNet: Fast and Lightweight Mel-Spectrogram Vocoder Incorporating Inverse Short-Time Fourier Transform," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [15] S.-g. Lee, W. Ping, B. Ginsburg, B. Catanzaro, and S. Yoon, "BigVGAN: A Universal Neural Vocoder with Large-Scale Training," in *International Conference on Learning Representations (ICLR)*, 2023.
- [16] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [17] P. Dhariwal and A. Nichol, "Diffusion Models Beat GANs on Image Synthesis," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [18] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "WaveGrad: Estimating Gradients for Waveform Generation," in *International Conference on Learning Representations (ICLR)*, 2021.
- [19] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A Versatile Diffusion Model for Audio Synthesis," in *International Conference on Learning Representations (ICLR)*, 2021.
- [20] R. Huang, M. W. Y. Lam, J. Wang, D. Su, D. Yu, Y. Ren, and Z. Zhao, "FastDiff: A Fast Conditional Diffusion Model for High-Quality Speech Synthesis," in *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, 2022.
- [21] Y. Koizumi, H. Zen, K. Yatabe, N. Chen, and M. Bacchiani, "SpecGrad: Diffusion Probabilistic Model based Neural Vocoder with Adaptive Noise Spectral Shaping," in *Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2020.
- [22] Z. Chen, X. Tan, K. Wang, S. Pan, D. Mandic, L. He, and S. Zhao, "InferGrad: Improving Diffusion Models for Vocoder by Considering Inference in Training," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [23] Y. Koizumi, K. Yatabe, H. Zen, and M. A. U. Bacchiani, "WaveFit: An Iterative and Non-autoregressive Neural Vocoder based on Fixed-Point Iteration," in *IEEE Spoken Language Technology Workshop (SLT)*, 2022.
- [24] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [26] I. Yamada, M. Yukawa, and M. Yamagishi, "Minimizing the Moreau Envelope of Nonsmooth Convex Functions over the Fixed Point Set of Certain Quasi-Nonexpansive Mappings," *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pp. 345–390, 2011.
- [27] T. Karras, S. Laine, and T. Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks," in *Proc. IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 2019.
- [28] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, "Least Squares Generative Adversarial Networks," in *Proc. IEEE international conference on computer vision (ICCV)*, 2017.
- [29] J. Yang, J.-S. Bae, T. Bak, Y. Kim, and H.-Y. Cho, "GANSpeech: Adversarial Training for High-Fidelity Multi-Speaker Speech Synthesis," in *Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2021.
- [30] 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. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2019.
- [31] P. C. Loizou, "Speech Quality Assessment," *Multimedia analysis, processing and communications*, pp. 623–654, 2011.
- [32] C.-P. Hsieh, S. Ghosh, and B. Ginsburg, "Adapter-Based Extension of Multi-Speaker Text-to-Speech Model for New Speakers," *arXiv preprint arXiv:2211.00585*, 2022.
- [33] J. J. Webber, C. Valentini-Botinhao, E. Williams, G. E. Henter, and S. King, "Autovocoder: Fast Waveform Generation from a Learned Speech Representation using Differentiable Digital Signal Processing," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.