



Unified Source-Filter GAN with Harmonic-plus-Noise Source Excitation Generation

Reo Yoneyama, Yi-Chiao Wu, Tomoki Toda

Nagoya University, Japan

yoneyama.reo@g.sp.m.is.nagoya-u.ac.jp, yichiao.wu@g.sp.m.is.nagoya-u.ac.jp,
tomoki@icts.nagoya-u.ac.jp

Abstract

This paper introduces a unified source-filter network with a harmonic-plus-noise source excitation generation mechanism. In our previous work, we proposed unified Source-Filter GAN (uSFGAN) for developing a high-fidelity neural vocoder with flexible voice controllability using a unified source-filter neural network architecture. However, the capability of uSFGAN to model the aperiodic source excitation signal is insufficient, and there is still a gap in sound quality between the natural and generated speech. To improve the source excitation modeling and generated sound quality, a new source excitation generation network separately generating periodic and aperiodic components is proposed. The advanced adversarial training procedure of HiFiGAN is also adopted to replace that of Parallel WaveGAN used in the original uSFGAN. Both objective and subjective evaluation results show that the modified uSFGAN significantly improves the sound quality of the basic uSFGAN while maintaining the voice controllability.

Index Terms: Speech synthesis, neural vocoder, source-filter model, generative adversarial networks

1. Introduction

Speech generation is a technology for generating speech waveforms on the basis of text or acoustic features. In particular, models that generate waveforms on the basis of acoustic features are called vocoders. In recent years, neural vocoders [1–13] have been extensively studied, and the sound quality of the synthesized speech is almost as good as that of human speech. However, these vocoders have several drawbacks due to their fully data-driven manner. One is the lack of robustness to unknown data. For example, when fundamental frequency (F_0) deviates from the F_0 range of the training data, these neural vocoders often fail to synthesize high-fidelity speech, or the F_0 of the synthesized speech is not consistent with the given F_0 . Another shortage is the disability to independently control acoustic parameters, which is essential to the expressive speech generation applications such as entertaining text-to-speech and voice conversion. Consequently, most neural vocoders suffer from lack of the flexible controllability of pitch, timbre, and aperiodicity, which are usually achieved with the conventional parametric vocoders [14, 15]. To address this issue, studies [6–9] have attempted to improve the F_0 controllability by introducing the source-filter theory [16] into their deep neural networks. However, the ad hoc constraints and signal processing can cause several problems such as the degradations of sound quality.

Our previous work, unified Source-Filter GAN (uSFGAN) [17], mimics the source-filter model using a generative adversarial networks (GAN)-based [18] unified network. The key concepts of uSFGAN are the regularization loss on the out-

put signal of the source-network which requires flat spectral envelopes, the pitch-dependent source excitation signal generation, and the GAN-based training. This approach jointly optimizes the source excitation generation (source-network) and the resonance filtering (filter-network) modules, which outperforms Neural Source-Filter (NSF) [7, 8] and Quasi-Periodic Parallel WaveGAN (QPPWG) [5] in both sound quality and F_0 controllability. However, there is still a gap in sound quality between natural speech and the generated speech of uSFGAN.

In this paper, to bridge the gap, we propose several improvements to uSFGAN. First, we replace the regularization loss with a new one that utilizes the residual spectrum as the target for stable training. Secondly, to tackle the undesired periodicity in the aperiodic components of the source excitation signals caused by the pitch-dependent source-network, we propose a parallel periodic and aperiodic source excitation generation inspired by the success of the harmonic-plus-noise decomposing neural vocoders [8–11]. Thirdly, we adopt the training procedure of HiFiGAN [12] instead of that of Parallel WaveGAN (PWG) [3] to take the F_0 estimation errors into account. According to the objective and subjective evaluation results, the proposed harmonic plus noise uSFGAN (HN-uSFGAN) significantly improves the sound quality while keeping the F_0 controllability.

2. Baseline uSFGAN

This section presents the baseline uSFGAN, which mimics a source-filter model using a single unified neural network. Specifically, as shown in Fig. 1, the overall architecture of the baseline uSFGAN is a general GAN including a neural generator (G) and a neural discriminator (D). The uSFGAN generator includes source and filter networks, and a regularization loss is applied to the output of the source network for the excitation signal generation. In addition, not only the adversarial losses of the generator and discriminator but also an auxiliary spectral loss is adopted for improving the training stability [3].

2.1. Unified Source-Filter Networks

The generator of uSFGAN is factorized into the source excitation generation and resonance filtering networks. The regularization loss on the output signal of the source-network is designed based on the conventional source-filter models [14, 15] which assume that the power of any source excitation signal is constant and its spectral envelopes are flat. Therefore, the regularization loss is formulated as

$$L_{\text{reg}}(G) = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, I)} \left[\frac{1}{N} \|\log \hat{E}\|_2 \right], \quad (1)$$

where $\|\cdot\|_2$, \hat{E} and N denote the L2 norm, the magnitude of the source spectral envelopes, and the number of elements in

the magnitude, respectively. \hat{E} is extracted using the simplified CheapTrick [19] algorithm described in [17]. The baseline uSFGAN jointly optimizes the cascaded source excitation and resonance filtering network to generate reasonable source excitation signals and high-fidelity speech with the source regularization loss and the GAN-based structure. The network architecture is interpretable and tractable with the flexibility of controlling voice components such as F_0 .

2.2. Pitch-dependent Source Excitation Generation

To capture the long term periodicity of source excitation signals and retain the robustness to unseen F_0 in the inference stage, pitch-dependent dilated convolution neural networks (PDCNNs) [4,5], which dynamically change the sizes of the receptive fields by multiplying the convolution dilation factors according to the input F_0 , are adopted in the source-network. On the other hand, the usual dilated convolution networks (DCNN) with the same structure as PWG [3] are adopted in the filter-network. Furthermore, inspired by NSF [7], the inputs of the generator are white noise and the sinusoidal signals corresponding to the input F_0 for efficiently modeling the periodic components of the source excitation signal.

2.3. PWG-based Adversarial Training

The training procedure of uSFGAN is based on that of PWG. The PWG discriminator is trained to identify natural samples as *real* and generated samples as *fake* by minimizing the following optimization criterion:

$$L_D(G, D) = \mathbb{E}_{\mathbf{x} \sim p_{data}} [(1 - D(\mathbf{x}))^2] + \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, I)} [D(G(\mathbf{z}))^2], \quad (2)$$

where \mathbf{x} denotes the natural samples, p_{data} denotes the data distribution of the natural samples and \mathbf{z} is random noise distributed from Gaussian distribution. On the other hand, the generator is trained to deceive the discriminator by minimizing the following adversarial loss:

$$L_{adv}(G, D) = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, I)} [(1 - D(G(\mathbf{z})))^2]. \quad (3)$$

GAN-based vocoders often use an auxiliary loss in the spectral-domain to stabilize the training process and encourage the generator to generate speech according to the input auxiliary features. Following PWG, uSFGAN uses the multi-resolution short-time Fourier transform (STFT) loss as the auxiliary spectral loss (L_{spc}). Therefore, the final loss function of the generator can be written as the summation of the regularization loss, the auxiliary spectral loss, and the adversarial loss:

$$L_G(G, D) = L_{reg}(G) + \lambda_{spc} L_{spc}(G) + \lambda_{adv} L_{adv}(G, D), \quad (4)$$

where λ_{spc} and λ_{adv} are loss balancing hyperparameters.

3. Proposed Harmonic + Noise uSFGAN

For improving source excitation modeling and generated sound quality, a new residual spectra regularization loss for the generated source excitation signals and a harmonic plus noise source excitation generation are proposed. We adopt the HiFiGAN-based adversarial training for further speech quality improvement instead of the PWG-based one.

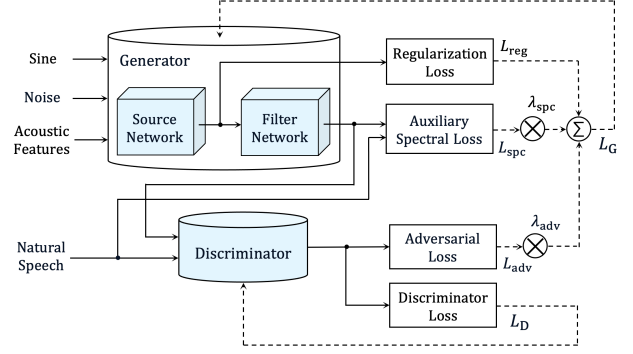


Figure 1: Overall architecture of unified Source-Filter GAN.




Figure 2: Architecture of proposed harmonic-plus-noise source excitation generation network. \oplus denotes element-wise addition.

3.1. Residual Spectra Regularization Loss

According to the source-filter theory [16], the process of human speech production can be modeled by the source excitation generation (controlling the F_0 and the power) and the resonance filtering (controlling the timbre). While the characteristics of the resonance filtering are represented by the spectral envelopes in the spectral-domain, the characteristics of source excitation generation can be obtained as the residual spectra between the natural speech and the spectral envelopes in the log spectra. Therefore, the residual spectra have reasonable information about source excitation signal.

We propose a new regularization loss utilizing residual spectra as the target of the output source excitation signal. The residual spectra are calculated from the natural speech using the spectral envelopes estimated by CheapTrick [19]. The residual spectra regularization loss can be written as

$$L_{reg} = \mathbb{E}_{\mathbf{x} \sim p_{data}, \mathbf{z} \sim \mathcal{N}(0, I)} \left[\frac{1}{N} \|\log f(S) - \log f(\hat{S})\|_1 \right], \quad (5)$$

where \hat{S} is the amplitude spectra of the output source excitation signal, S is the amplitude residual spectra which have the same frame-wise average power as that of the target speech, and N is the number of elements in the amplitude spectra. To flexibly model the source excitation while tolerating perceptually acceptable distortion, mel-filter bank f is applied to both amplitude spectra. Unlike the regularization loss of the baseline uSFGAN, the new loss makes the power estimation of the source-network like a real human speech production process.

3.2. Harmonic-plus-Noise Source Excitation Generation

Although PDCNNs are very effective to improve F_0 controllability, they are not suitable for generating unvoiced segments of the source excitation signal. Specifically, the periodic operation sometimes causes undesired periodic components in the unvoiced segments. To improve the modeling of both periodic

and aperiodic components of the source excitation signals, we introduce a parallel harmonic-plus-noise source excitation generation architecture as shown in Fig. 2.

The harmonic-plus-noise source-network consists of three modules: harmonic-network, noise-network, and periodicity-estimator. The harmonic-network is composed of several PDCNN-based blocks, and the latent feature $l^{(h)}$ corresponding to the periodic components of the source excitation signal is generated on the basis of the input sinusoidal signal and auxiliary features. On the other hand, the noise-network is composed of several CNN-based blocks, and the latent feature $l^{(n)}$ corresponding to the aperiodic components of the source excitation signal is generated on the basis of the input random noise signal and auxiliary features.

The two latent features are element-wise summed using the channel-wise and sample-wise weights \mathbf{a} . These weights are corresponding to the speech periodicity and estimated by the periodicity-estimator according to the input auxiliary features. The source excitation latent feature l is formulated as

$$l_{t,i} = a_{t,i} \cdot l_{t,i}^{(h)} + (1 - a_{t,i}) \cdot l_{t,i}^{(n)}, \quad (6)$$

where the subscripts indicate the i^{th} channel of the t^{th} sample of each latent feature or the weights. After the harmonic-plus-noise excitation generation, l is passed to the filter-network to generate the synthesized speech and also to the 1×1 convolution to generate the source excitation signal by reducing the number of channels to one for the regularization loss calculation. We set the the number of channels in l to 64 as the same number of the residual channel.

This architecture is inspired by multi-band harmonic-plus-noise PWG (HN-PWG) [11] in which periodic and aperiodic output signals through band-pass-filters are summed up with estimated energy for each subband. However, our approach applies harmonic-plus-noise decomposition to the source excitation generation while HN-PWG applies it to the final speech generation. Moreover, we don't adopt band-pass-filtering but estimate channel-wise energy in latent space, promoting the network decomposing periodic and aperiodic components in more flexible way.

3.3. HiFiGAN-based Adversarial Training

Inspired by the work of You et al. [13] showing that HiFiGAN's multi-period and multi-scale discriminators [12] are very effective to improve the sound quality of neural vocoders, the effectiveness of these discriminators for our proposed method is explored. Moreover, the auxiliary mel-spectral loss from HiFiGAN is also adopted. Specifically, since our method depends on an external F_0 extractor, the error propagations from the feature extraction errors are inevitable, and these errors usually cause significant degradation of the training losses, such as the multi-resolution STFT loss in PWG and the feature matching loss in HiFiGAN. Consequently, even if the generator can generate sufficiently natural speech according to the input F_0 , the inevitable mismatches of F_0 s and phases still lead to large multi-resolution STFT and feature matching losses. On the other hand, the mel-spectral loss has a relatively large tolerance for those mismatches because the resolution of each frequency band is reduced according to human perception. Since the generator can still learn the phase information via the adversarial training mechanism, the multi-resolution STFT loss is not necessary for the proposed model. Although the feature matching loss is adopted in HiFiGAN, the proposed method excludes it to avoid the mismatch issues.

4. Experimental Evaluations

4.1. Experimental Setting

4.1.1. Data Preparation

In the following experiments, we used the VCTK corpus [20] composed of 110 speakers, and the sampling frequency was set to 24000 Hz. We limited F_0 range of training data from 70 Hz to 340 Hz and excluded two speakers (p271, p300) from training data to use for evaluation. Consequently, we used 99 speakers as seen speakers and the rest as unseen speakers. The testing data consisted of 100 randomly selected utterances from the seen speakers and 100 randomly selected utterances from the unseen speakers. The total number of the training utterances was 35671, which was around 33 hours.

4.1.2. Model Details

We compared proposed HN-uSFGAN with two baselines and four ablation models as follows:

- **WORLD**: Conventional source-filter vocoder [15].
- **uSFGAN**: Baseline uSFGAN vocoder [17]. λ_{spc} and λ_{adv} were set to 1.0 and 4.0, respectively.
- **HN-uSFGAN**: Proposed model with the harmonic-network (20 PDCNN layers with 4 dilation cycles), the noise-network (5 CNN layers), and the filter-network (30 DCNN layers with 3 dilation cycles). λ_{spc} and λ_{adv} were set to 15.0 and 1.0, respectively.
- **-Reg-loss**: Proposed model trained without the residual spectra loss.
- **-HN-SN**: Proposed model without the harmonic-plus-noise source-network but with the source-network of baseline uSFGAN.
- **-HiFi-D**: Proposed model without the HiFiGAN discriminators [12] but with the PWG one.
- **-Mel-loss**: Proposed model trained without the mel-spectra loss but with the multi-resolution STFT loss [3].

The baseline uSFGAN was trained with only the auxiliary losses for the first 100k iterations and trained with the discriminator in the rest 500k steps using the RAdam optimizer [21]. On the other hand, HN-uSFGAN followed the HiFiGAN training procedure to simultaneously train the generator and the discriminators from scratch for 600k iterations using the Adam optimizer [22].

The acoustic features extracted by WORLD with a 5 ms shift length were adopted for all neural vocoders. Specifically, we used one-dimensional continuous F_0 , 41 dimensional mel-cepstral coefficients, and three-dimensional coded aperiodicity.

4.2. Experimental Results

We conducted both objective and subjective evaluations. For the objective evaluation, three measurements were used: root mean square error of $\log F_0$: RMSE [Hz], voiced or unvoiced decision error: V/UV [%], and mel-cepstral distortion: MCD [dB]. For the subjective evaluation, we conducted opinion tests on sound quality with 10 subjects, and each subject evaluated 20 utterances per method and condition.¹

¹We provide audio samples at <https://chomeyama.github.io/HN-UnifiedSourceFilterGAN-Demo/>.

Table 1: Results of objective and subjective evaluations. The best results and those that are not significantly different from them are highlighted in bold.

method	RMSE ↓	V/UV ↓	MCD ↓	MOS ↑
Reconstruction (original F_0)				
Natural	–	–	–	3.90 ± 0.04
WORLD	0.05	12	3.51	3.64 ± 0.06
uSFGAN	0.05	12	3.09	3.66 ± 0.05
HN-uSFGAN	0.05	9	2.82	3.79 ± 0.05
- Reg-loss	0.05	10	2.95	3.66 ± 0.05
- HN-SN	0.06	11	3.22	3.80 ± 0.05
- HiFi-D	0.05	8	2.65	3.56 ± 0.05
- Mel-loss	0.06	11	3.04	3.76 ± 0.05
F_0 transformation ($2^{-0.75} \times F_0$)				
WORLD	0.08	16	3.62	3.60 ± 0.07
uSFGAN	0.11	14	3.33	3.47 ± 0.07
HN-uSFGAN	0.11	14	3.08	3.76 ± 0.06
- Reg-loss	0.13	18	3.20	3.70 ± 0.06
- HN-SN	0.13	17	3.34	3.80 ± 0.06
- HiFi-D	0.09	11	2.90	3.62 ± 0.06
- Mel-loss	0.15	17	3.43	3.67 ± 0.06
F_0 transformation ($2^{0.75} \times F_0$)				
WORLD	0.09	12	3.92	3.39 ± 0.06
uSFGAN	0.08	17	3.62	3.30 ± 0.06
HN-uSFGAN	0.10	12	3.37	3.48 ± 0.06
- Reg-loss	0.11	16	3.45	3.36 ± 0.07
- HN-SN	0.10	12	3.80	3.45 ± 0.06
- HiFi-D	0.10	13	3.32	3.28 ± 0.07
- Mel-loss	0.10	13	3.39	3.42 ± 0.06

4.2.1. Speech Reconstruction

First, we evaluated the performance of speech reconstruction (i.e., original F_0) and the evaluation results are shown in the top rows of Table 1. Since there was not significant difference between seen and unseen speakers, we have combined all the results together. The results show that HN-uSFGAN outperforms the baselines in all objective metrics and perceptual sound quality. According to the ablation results, we can find that the proposed method achieves the top performances in both the objective and subjective evaluations, which demonstrates the effectiveness of each proposed module especially the residual spectral loss and the HiFiGAN discriminators for perceptual quality improvements.

4.2.2. Speech Generation with Transformed F_0

Next, we evaluated the performances of F_0 transformation with factors of $2^{-0.75} \simeq 0.59$ and $2^{0.75} \simeq 1.68$. The ground truth F_0 s were determined by multiplying the F_0 s extracted from natural speech with the scale factors, and they were also adopted as the auxiliary F_0 for all models. The evaluation results are shown in the middle and bottom rows of Table 1. According to the results, we can find that the proposed HN-uSFGAN generates speech with better V/UV error rates, MCDs, and MOS scores for both F_0 factors than WORLD while maintaining almost the same F_0 accuracy as WORLD. Again, these results show the effectiveness of the proposed model to generate high-fidelity speech while attaining the high F_0 controllability. Since flexibly manipulating F_0 is an essential feature of conventional vocoders, the proposed HN-uSFGAN vocoder meets the definition of a vocoder better than the other neural vocoders. In addition, the baseline uSFGAN significantly inferior to other models in sound quality. We attribute the degradation to its capability to deal with multiple speakers and unseen speakers.




Figure 3: Waveforms of output periodic and aperiodic source excitation signals and the natural speech from top to bottom.




Figure 4: Waveforms of output source excitation signals with different F_0 scaling factors (0.5, 1.0, and 2.0 from top to bottom). Original F_0 values over this segment are around 180 Hz.

4.3. Visualization of Source Excitation Signals

To investigate the behaviors of proposed HN-uSFGAN, we visualized its output source excitation signals. As shown in Fig. 3, the generated periodic and aperiodic source excitation signals demonstrate the high capability of HN-uSFGAN to separately model the periodic and aperiodic components. Moreover, Fig. 4 shows that HN-uSFGAN accurately reconstructs the source excitation signals according to the input F_0 and achieves a high capability to handle both the periodic and aperiodic components.

5. Conclusions

The proposed HN-uSFGAN has improved both the speech quality and F_0 controllability of the baseline uSFGAN by harmonic-plus-noise source excitation generation mechanism and the HiFiGAN-based adversarial training procedure. The proposed method has also shown comparative F_0 controllability as the WORLD vocoder which is a conventional vocoder with strong F_0 controllability while significantly outperforming WORLD in sound quality. For future work, we intend to further improve the sound quality of HN-uSFGAN to tackle the occasional buzzy noises in the unvoiced segments.

Acknowledgement: This work was supported in part by JST, CREST Grant Number JPMJCR19A3 and JSPS KAKENHI Grant Number 21H05054.

6. 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 *Proc. SSW9*, Sept. 2016, p. 125.
- [2] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," in *Proc. ICML*, July 2018, pp. 2415–2424.
- [3] 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*, May 2020, pp. 6199–6203.
- [4] Y.-C. Wu, T. Hayashi, P. L. Tobing, K. Kobayashi, and T. Toda, "Quasi-periodic wavenet: An autoregressive raw waveform generative model with pitch-dependent dilated convolution neural network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1134–1148, 2021.
- [5] 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 convolution neural network," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 792–806, 2021.
- [6] J.-M. Valin and J. Skoglund, "LPCNet: Improving neural speech synthesis through linear prediction," in *Proc. ICASSP*, May 2019, pp. 5891–5895.
- [7] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter-based waveform model for statistical parametric speech synthesis," in *Proc. ICASSP*, May 2019, pp. 5916–5920.
- [8] X. Wang and J. Yamagishi, "Neural Harmonic-plus-Noise Waveform Model with Trainable Maximum Voice Frequency for Text-to-Speech Synthesis," in *Proc. SSW*. ISCA: ISCA, 2019, pp. 1–6.
- [9] Z. Liu, K. Chen, and K. Yu, "Neural homomorphic vocoder," in *Proc. INTERSPEECH*, Oct. 2020, pp. 240–244.
- [10] 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 *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 6049–6053.
- [11] 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 2021*, 2021, pp. 2227–2231.
- [12] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," 2020.
- [13] J. You, D. Kim, G. Nam, G. Hwang, and G. Chae, "GAN Vocoder: Multi-Resolution Discriminator Is All You Need," in *Proc. Interspeech 2021*, 2021, pp. 2177–2181.
- [14] H. Kawahara, I. Masuda-Katsuse, and A. De Cheveigne, "Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based F0 extraction: Possible role of a repetitive structure in sounds," *Speech Communication*, vol. 27, no. 3-4, pp. 187–207, 1999.
- [15] 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, 2016.
- [16] R. McAulay and T. Quatieri, "Speech analysis/synthesis based on a sinusoidal representation," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 34, no. 4, pp. 744–754, 1986.
- [17] R. Yoneyama, Y.-C. Wu, and T. Toda, "Unified Source-Filter GAN: Unified Source-Filter Network Based On Factorization of Quasi-Periodic Parallel WaveGAN," in *Proc. Interspeech 2021*, 2021, pp. 2187–2191.
- [18] 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*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014.
- [19] M. Morise, "Cheaptrick, a spectral envelope estimator for high-quality speech synthesis," *Speech Communication*, vol. 67, pp. 1–7, 2015.
- [20] J. Yamagishi, C. Veaux, and K. MacDonald, "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2019. [Online]. Available: <https://doi.org/10.7488/ds/2645>
- [21] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," in *Proc. ICLR*, Apr. 2020.
- [22] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.