



Reverberation Modeling for Source-Filter-based Neural Vocoder

Yang Ai¹, Xin Wang², Junichi Yamagishi^{2,3}, Zhen-Hua Ling¹

¹NELSLIP, University of Science and Technology of China, Hefei, P.R.China

²National Institute of Informatics, Japan

³CSTR, University of Edinburgh, UK

ay8067@mail.ustc.edu.cn, wangxin@nii.ac.jp, jyamagis@nii.ac.jp, zhling@ustc.edu.cn

Abstract

This paper presents a reverberation module for source-filter-based neural vocoders that improves the performance of reverberant effect modeling. This module uses the output waveform of neural vocoders as an input and produces a reverberant waveform by convolving the input with a room impulse response (RIR). We propose two approaches to parameterizing and estimating the RIR. The first approach assumes a global time-invariant (GTI) RIR and directly learns the values of the RIR on a training dataset. The second approach assumes an utterance-level time-variant (UTV) RIR, which is invariant within one utterance but varies across utterances, and uses another neural network to predict the RIR values. We add the proposed reverberation module to the phase spectrum predictor (PSP) of a HiNet vocoder and jointly train the model. Experimental results demonstrate that the proposed module was helpful for modeling the reverberation effect and improving the perceived quality of generated reverberant speech. The UTV-RIR was shown to be more robust than the GTI-RIR to unknown reverberation conditions and achieved a perceptually better reverberation effect.

Index Terms: reverberation, room impulse response, source-filter-based model, neural vocoder

1. Introduction

Recently several neural autoregressive models such as WaveNet [1], SampleRNN [2], and WaveRNN [3], have been proposed for raw audio generation. Their variants, such as knowledge-distilling-based models (e.g., parallel WaveNet [4] and ClariNet [5]) and flow-based models (e.g., WaveGlow [6]), were then proposed to further improve the performance and efficiency. These models can be used as *neural vocoders* [7, 8, 9, 10, 11, 12] wherein speech waveforms can be reconstructed from acoustic features for various tasks [13, 14, 15]. It was confirmed that these neural vocoders outperform vocoders using classical signal processing techniques. However, some limitations still exist — either a low generation speed, tricky training process, or complicated model structure.

Motivated by the limitations, new types of neural vocoders, such as the glottal neural vocoder [16, 17] and LPCNet [18], have been further proposed by combining speech production mechanisms with neural networks, and their performance is impressive. However, all of the above models operate under the autoregressive assumption and are slow in either waveform generation or training. Previously, we proposed non-autoregressive neural source-filter (NSF) [19] and HiNet vocoders [20, 21]. The NSF vocoder uses dilated convolutions to transform a sine-based source signal into an output waveform, following the idea of the source-filter speech production model [22]. The HiNet vocoder is composed of an amplitude spectrum predictor (ASP) and a phase spectrum predictor (PSP), where the PSP is built by using the NSF vocoder for better phase recovery.

The outputs of the ASP and PSP are combined to recover speech waveforms via short-time Fourier synthesis (STFS). Experimental results show that the NSF and HiNet vocoders can generate waveforms with high quality and high efficiency for speech [20, 21] and musical instrument sounds [23] recorded in acoustically isolated studios.

However, unlike the ideal data for speech or music synthesis, audio signals captured for real-life applications typically contain room reverberation. The reverberation poses a challenge to non-autoregressive neural vocoders, and the quality of synthesized speech usually degrades. Recently, Engel *et al.* tried to introduce a reverberation module with a trainable room impulse response (RIR) into a sinusoidal vocoder [24]. Their model successfully learned room reverberation effects on a solo violin dataset under a signal reverberation condition. However, learning the reverberation effects in multiple acoustic environments and applying the model for unseen acoustic environments have not yet been investigated, and neither has the model been evaluated on a reverberant speech dataset.

As an initial step towards robust reverberation modeling for speech data, this paper proposes a trainable reverberation module for neural vocoders. This module uses the output waveform of neural vocoders as an input and outputs a reverberant waveform by convolving the input with a RIR. We design two types of neural RIR estimators. One estimates the global time-invariant (GTI) RIR, which is invariant among a whole dataset and is regarded as a trainable variable of a model. This is similar to [24]. The other infers an utterance-level time-variant (UTV) RIR, which is invariant inside one utterance but varies among different utterances. The UTV-RIRs are predicted by an additional trainable neural network that uses the same conditional features as neural vocoders. We add the proposed reverberation module to the PSP of the HiNet vocoder, and experiments are conducted on a multi-speaker reverberant speech database with various types of reverberation conditions, including unseen ones. Furthermore, a multi-task training strategy that uses both reverberant and corresponding dry waveforms is also investigated.

This paper is organized as follows: In Section 2, we briefly review the NSF and HiNet vocoders. In Section 3, we give details on our proposed reverberation module and neural RIR estimators. Section 4 reports our experimental results. Conclusions are given in Section 5.

2. Brief view of NSF and HiNet vocoders

The NSF models [19] generate speech waveforms from input acoustic features through time-domain non-linear transformations. They include three modules: a conditional module that upsamples input acoustic features such as F0 and mel-spectrogram, a source module that outputs a sine-based source signal given the F0, and a dilated-convolution-based filter module that transforms the source signal into an output waveform.

The NSF models are suitable for applications where the users want to precisely control the F0 of the output waveform.

HiNet [20] is a neural vocoder that produces speech waveforms from input acoustic features by predicting amplitude and phase spectra hierarchically. The HiNet vocoder consists of two predictors, an ASP and a PSP. The ASP uses acoustic features as input and predicts frame-level log amplitude spectra (LAS). Then, the F0 and LAS predicted by the ASP are sent into the PSP for phase spectra prediction. Finally, the predicted amplitude and phase spectra are combined to reconstruct speech waveforms by STFS.

In our latest work [21], the ASP consisted of multiple convolutional layers for converting acoustic features into the LAS. Generative adversarial networks (GANs) were also newly introduced into the ASP; the ASP was used as a generator, and two discriminators were adopted. Both discriminators consisted of multiple convolutional layers, which operate convolution along with either the frequency or time axis of the input LAS, respectively. The training of the ASP is based on a Wasserstein GAN [25] loss together with the mean square error (MSE) between the predicted LAS and natural ones.

The PSP conducts two steps: neural waveform generation and phase spectrum extraction. The neural waveform generator was based on a customized NSF vocoder [19] with three modifications for better phase recovery: 1) the use of LAS as input, 2) pre-calculation of the initial phase of the sine-based excitation signal for each voiced segment at the training stage, and 3) the use of a combined loss function including MSE on amplitude spectra, waveform loss, and correlation loss.

3. Proposed methods

When a speech waveform signal $\mathbf{d} = [d_1, \dots, d_T]^\top$ of length T is produced in a closed room, it propagates to an observation point through a direct path, reflects off walls and surrounding objects and becomes a reverberant signal. By assuming that the RIR of a room can be approximated by the finite impulse response sequence $\mathbf{h} = [h_1, \dots, h_L]^\top$ [26, 27], where $h_1 = 1$ denotes the direct path, the received reverberant signal $\mathbf{r} = [r_1, \dots, r_T]^\top$ can be written as

$$\mathbf{r} = \mathbf{d} * \mathbf{h}. \quad (1)$$

On the basis of this principle, we propose a reverberation module for the HiNet vocoder. This module accepts the output waveform of the PSP in the HiNet vocoder as input, which is illustrated in Fig. 1.¹ Although we can directly compute Eq. (1) through convolution in the time domain, in order to reduce the computational cost, we implement the convolution in the frequency domain as

$$\mathbf{r} = \mathcal{F}^{-1}[\mathcal{F}(\mathbf{d}) \odot \mathcal{F}(\mathbf{h})], \quad (2)$$

where \mathcal{F} , \mathcal{F}^{-1} , and \odot represent the FFT, inverse FFT, and element-wise product, respectively.

There are two ways to parameterize and estimate the value of the RIR \mathbf{h} ²:

- **Global time-invariant (GTI) RIR:** inspired by DDSF [24], the RIR \mathbf{h} is assumed to be time-invariant and shared for all speech data, and the values of its coefficients $\{h_1, \dots, h_L\}$ are learned from the training data. Note that, because $h_1 = 1$, only $L - 1$ elements in \mathbf{h} need to be learned.

¹We also tried to add a reverberation module based on a causal convolution network for the ASP, but it was not effective.

²We also tried to parameterize the RIR as an exponentially decaying function with a trainable decay rate, but the learned RIR was intractable.

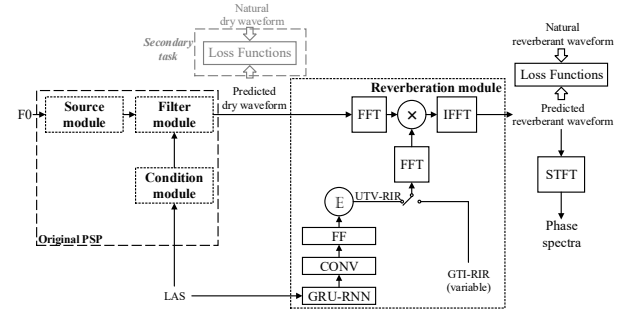


Figure 1: Flowchart of PSP with reverberation module. Here FF, CONV, and GRU-RNN represent feed-forward, convolutional, and unidirectional GRU-based recurrent layers, respectively, and \times and \mathbb{E} denote element-wise product and temporal averaging operation, respectively.

- **Utterance-level time-variant (UTV) RIR:** the RIR \mathbf{h} is assumed to be invariant for one utterance, but different utterances acquire different \mathbf{h} . The value of \mathbf{h} is predicted from the input LAS by a small conditional network that consists of trainable recurrent layers, convolutional layers, feed-forward layers, and a temporal average pooling layer as shown in Fig. 1. The temporal average pooling layer averages the hidden features of all of the frames and gives a single vector as the predicted \mathbf{h} .

The GTI-RIR is expected to be suitable for scenarios where we want to learn the RIR of one time-invariant acoustic environment, while the UTV-RIR is suitable for more general cases where the speech data is recorded in several different acoustic environments. During training, the reverberation module and the PSP are jointly optimized by a loss function consisting of multi-resolution spectral distortions [19] between the output of the reverberation module and the natural reverberant waveform.

For cases where the dry waveforms of the reverberant data are also available (e.g., when reverberation data are generated from clean data through simulation or replaying), we further investigate a multi-task training strategy that uses not only reverberant data but also dry waveforms. As the gray region in Fig. 1 shows, the loss function of the secondary task is a combination of MSE on LAS, waveform loss, and correlation loss [20] between a generated dry waveform and the natural dry waveform. The whole loss function is the sum of the loss functions of the main and secondary tasks.

4. Experiments

4.1. Data and feature configuration

A multi-speaker reverberant speech database³ [28] was used in our experiments. From the database, we used a reverberant subset of 28 speakers that contained 11,572 utterances and 18 reverberation types (9 rooms \times 2 microphone positions). We randomly divided this subset into a training set (11,012 utterances) and validation set (560 utterances). Regarding the test set, there were three scenarios below in our experiments:

- T1 Two unseen speakers' reverberant data with 6 unseen reverberation types (3 rooms \times 2 microphone positions), 824 utterances in total;
- T2 Two unseen speakers' reverberant data with the same 18 reverberation types as in the training set, 832 utterances in total;
- T3 Dry speech version of T1.

³<https://doi.org/10.7488/ds/1425>

The original 48-kHz waveforms were down sampled to 24-kHz for the experiments. The acoustic features included the 80-dimensional mel-spectrogram, F0 extracted using YAPPT [29], and a voiced/unvoiced flag. The LAS used by HiNet was computed using 2048 FFT points. All features were extracted with a frame shift and length of 12 and 50 ms, respectively.

4.2. Experimental models

We compared the following variants in the experiments⁴:

N-BL: The harmonic-plus-noise NSF model [30] was included as a reference model without the reverberation module. The number of model parameters was around 1.2×10^6 .

H-BL: Baseline HiNet vocoder without the proposed reverberation module. We used the baseline ASP configuration in our previous work [21] and added two convolution layers having 512 and 1024 channels, respectively, to the discriminator that operated along with the frequency axis. This was motivated by the increased number of FFT points of the LAS. We adopted the same PSP used as the baseline in our previous work [21], but GANs were not used here. The loss function was a combination of MSE on LAS, waveform loss, and correlation loss. Note that the NSF module in the PSP is slightly different from **N-BL** (see details in [20]). The number of model parameters was around 6.2×10^7 for the ASP and 7.2×10^6 for the PSP.

H-GTI: HiNet with the GTI-RIR-based reverberation module integrated into PSP. The RIR length was 6,000. The model configuration and size were the same as **H-BL** except for the increased 5,999 trainable parameters for GIT-RIR.

H-UTV: HiNet with the UTV-RIR-based reverberation module integrated into the PSP. The RIR length was 6,000. The trainable neural network that converts LAS to RIR consisted of a unidirectional GRU layer with 1,024 nodes, a convolution layer with 1,024 nodes and a kernel-size of 11, and a feed-forward linear layer with 5,999 output nodes. Other settings were the same as those of **H-BL**. The number of model parameters for the PSP was increased by 2.4×10^7 compared with **H-BL**. Since the UTV-RIR is non-autoregressive, the increased model size did not cause an obvious degradation of generation efficiency.

H-UTV-MT: same as **H-UTV** but with the secondary task using dry waveforms during training.

4.3. Main experiments

Our main experiments focused on the reverberation effect and speech quality. We compared **N-BL**, **H-BL**, **H-GTI**, and **H-UTV** under testing scenarios **T1** and **T2** using both objective and subjective evaluations.

4.3.1. Objective evaluation – T60 comparisons –

T60 estimation errors [31] were used as the objective metric to evaluate the reverberation effects. T60 is also called the reverberation time, and it is defined as the time it takes for sound to decay by 60 dB after the source has been switched off. We used an open source toolkit [32] to blindly estimate T60 from the reverberant speech. The T60 estimation errors were calculated as the difference between the estimated T60 and the ground-truth T60 (T60n) reported in the database paper [28].

We calculated the T60 estimation errors for the output waveforms from all of the experimental models. For reference,

⁴Examples of generated speech can be found at <http://home.ustc.edu.cn/~ay8067/reverb/demo.html>. Scripts and toolkits for the NSF model can be found at <https://github.com/nii-yamagishilab/project-CURRENNT-scripts>




Figure 2: Box plots of T60 estimation errors for utterances with $T60n = 0.362s$ under test scenario **T1**.

we also calculated the errors for the natural reverberant waveform and the output waveform from the PSP in the HiNet models (denoted by **P-***). Figure 2 shows box plots of T60 estimation errors for utterances with $T60n = 0.362s$ under test scenario **T1**. Note that the T60 estimated errors for natural reverberant speech were non-zero because blind estimation of T60 is not error-free.

Figure 2 demonstrates that both **P-GTI** and **P-UTV** had smaller errors than **P-BL**, which indicates the usefulness of the proposed reverberation module in the PSP component of HiNet. Furthermore, **P-UTV** had a smaller error than **P-GTI**, suggesting that UTV-RIR is more effective than GTI-RIR in modeling unseen reverberation types. By comparing **P-*** with **H-***, we see that **H-*** had smaller errors than **P-***. This suggests that the ASP is able to produce the reverberation effect by a moderate amount even though the ASP has no explicit reverberation module. The performance differences among **H-*** vocoders are small. Additionally we can observe that **N-BL** had smaller errors than **P-BL** while **H-BL** had marginally smaller errors than **N-BL**.

4.3.2. Subjective evaluation

We conducted two types of listening tests on the crowdsourcing platform Amazon Mechanical Turk⁵ with anti-cheating considerations [33] to evaluate the reverberation effect and speech quality, respectively. In each test, 20 test-set utterances were generated for each test scenario by each experimental model, and these utterances were evaluated by about 40 English native listeners.

Reverberation effect: The first test was a similarity test on the reverberation effect. Listeners were asked to first listen to the natural dry and reverberant audio tracks. They were then asked to listen to a few test audio tracks and assign a score from 1 to 9 to each, where a higher score denoted a reverberation effect more similar to that in the natural reverberant audio tracks. The audio tracks generated from the PSP in the HiNet models were directly used for the listening test, and they are denoted as **P-***. Furthermore, to investigate the impact of the proposed reverberation module, the listening test used the audio tracks generated from the PSP after the trained reverberation module was removed, and they are denoted as **P-*(dry)**.

The results for test scenarios **T1** (unseen reverberant type) and **T2** (seen reverberant type) are plotted in Figure 3. As expected, the similarity scores of **P-GTI** and **P-UTV** had higher means and medians than those of **P-BL** in both **T1** and **T2**. This means that the proposed module generated reverberation that was perceptually more similar to the natural reverberant speech. Furthermore, **P-UTV** outperformed **P-GTI** in **T1**. These results were consistent with the results for the T60 estimation errors in Section 4.3.1. For **T2**, however, **P-UTV** did not

⁵<https://www.mturk.com>.

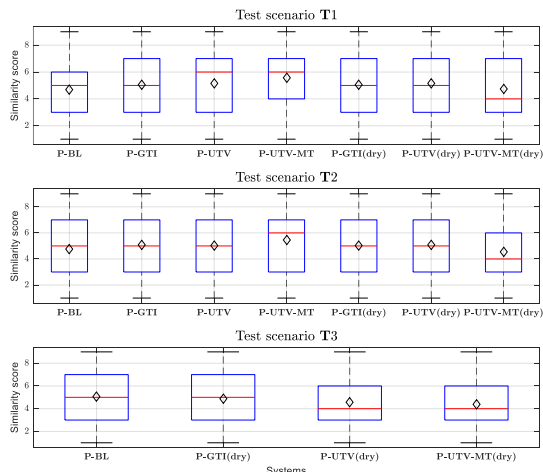


Figure 3: Box plot of reverberant effect similarity scores for all test scenarios. Here, black diamonds and red lines represent mean and median.

outperform **P-GTI**, which indicates that **P-UTV** may be more suitable for unknown reverberation conditions. Unfortunately, the similarity scores of **P-*(dry)** remained similar to **P-BL**. It seems that the evaluated models did not have de-reverberation ability, i.e., they could not generate perfect dry waveforms giving reverberant acoustic features. One possible reason may be that the reverberation module was jointly trained with the rest of the network. This point is further investigated together with multi-task learning in the next section.

Quality: The second test was a MUSHRA (Multiple Stimuli with Hidden Reference and Anchor) test [34] done to compare the quality of the generated waveforms. The average MUSHRA scores and their 95% confidence intervals are shown in Figure 4. The reference audio tracks in the MUSHRA test for **T1** and **T2** were the natural reverberant waveforms.

As Figure 4 shows, **H-GTI** and **H-UTV** had higher MUSHRA scores than **H-BL** for both **T1** and **T2**, suggesting that the reverberation module in the PSP was helpful for improving the quality of synthetic speech for HiNet. The MUSHRA scores for **T2** were higher than those for **T1**, which is reasonable since unseen reverberation conditions are more challenging to model. We can also see that the difference between **H-GTI** and **H-UTV** was not significant. Utterance-dependent RIR estimation seems to be important for modeling multiple reverberation types as the T60 comparison and the similarity test suggest, but it does not improve the perceived quality of generated waveforms. Finally, **H-BL** outperformed **N-BL**, and this indicates that the reverberant speech generated from the HiNet vocoder sounded better than that from the NSF vocoder with the current configurations.

4.4. Additional analysis

Finally, we analyzed two additional configurations.

Multi-task training using dry waveforms: Models using the multi-task training are denoted as ***-UTV-MT**. By comparing **P-UTV-MT** with **P-UTV** in Figure 2, we see that using the multi-task training did not reduce the T60 estimation error. However, as the similarity test results in Figure 3 show, **P-UTV-MT** had a higher mean and median than **P-UTV**. Furthermore, the median of **P-UTV-MT(dry)** was 4.0, while that of **P-UTV-MT** was 6.0 for **T1** and **T2**, and the differences were larger than those between **P-UTV** and **P-UTV(dry)**. These results suggest that multi-task training using dry waveforms as




Figure 4: Average MUSHRA scores with 95% confidence interval for all test scenarios.

additional supervision makes the functional role of the proposed reverberation module more explicit. Regarding the quality, there was no obvious difference between between **H-UTV-MT** and **H-UTV** as shown in Figure 4.

Use of dry acoustic features (T3): If the proposed framework is well generalized, it should be able to generate dry speech with high quality when we input dry acoustic features. This was the purpose of **T3**. The results in Figure 3 show that the medians or the mean similarity scores of **P-GTI(dry)**, **P-UTV(dry)**, and **P-UTV-MT(dry)** were lower than those of the corresponding models in **T1** and **T2**. In other words, these models generated waveforms that were perceptually closer to the natural dry waveforms when using dry input acoustic features. These results are encouraging. However, the generated waveforms were not sufficiently close to the natural dry waveforms, so there is still room for improvement. The results of the quality comparisons are shown in Figure 4. The reference tracks used for the MUSHRA test were natural waveforms without reverberation. From the MUSHRA listening test, we can see that the quality scores for **T3** were similar to those of **T1** for unseen conditions.

5. Conclusion

In this paper, we proposed a neural reverberation module and integrated it into non-autoregressive source-filter-based neural vocoders. The reverberation module uses RIRs for convolving waveforms generated by the vocoders as the standard signal processing method does, but the RIRs are estimated jointly with other parameters of the neural vocoder or predicted by another trainable network. The former approach, called **GIT-RIR**, uses a globally invariant vector and is directly trained on a reverberant dataset. The latter approach, called **UTV-RIR**, uses another network to estimate a different RIR for each utterance. We conducted experiments by adding the proposed reverberation module to the PSP of the HiNet vocoder. Objective and subjective evaluation results indicated that the proposed reverberation module is helpful for modeling the reverberation effect and improving the quality of reverberant speech generated by the HiNet vocoder. We also confirmed that the **UTV-RIR** was better than the **GTI-RIR** when modeling multiple unseen reverberation types. For future work, we plan to apply the reverberation module to other neural vocoders.

6. Acknowledgements

This work was partially supported by a JST CREST Grant (JPMJCR18A6, VoicePersonae project), Japan, MEXT KAKENHI Grants (19K24371, 16H06302, 17H04687, 18H04120, 18H04112, 18KT0051), Japan, and the National Natural Science Foundation of China under Grant 61871358 and the China Scholarship Council (CSC). The experiments were partially conducted on TSUBAME 3.0 of Tokyo Institute of Technology.

7. 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," *arXiv preprint arXiv:1609.03499*, 2016.
- [2] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. Courville, and Y. Bengio, "SampleRNN: An unconditional end-to-end neural audio generation model," in *arXiv preprint arXiv:1612.07837*, 2017.
- [3] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," in *Proc. ICML*, 2018, pp. 2410–2419.
- [4] A. v. d. Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. v. d. Driessche, E. Lockhart, L. C. Cobo, F. Stimberg *et al.*, "Parallel WaveNet: Fast high-fidelity speech synthesis," in *Proc. ICML*, 2018, pp. 3918–3926.
- [5] W. Ping, K. Peng, and J. Chen, "ClariNet: Parallel wave generation in end-to-end text-to-speech," in *arXiv preprint arXiv:1807.07281*, 2019.
- [6] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A Flow-based Generative Network for Speech Synthesis," in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [7] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent WaveNet vocoder," in *Proc. Interspeech*, 2017, pp. 1118–1122.
- [8] T. Hayashi, A. Tamamori, K. Kobayashi, K. Takeda, and T. Toda, "An investigation of multi-speaker training for WaveNet vocoder," in *Proc. ASRU*, 2017, pp. 712–718.
- [9] N. Adiga, V. Tsiaras, and Y. Stylianou, "On the use of WaveNet as a statistical vocoder," in *Proc. ICASSP*, 2018, pp. 5674–5678.
- [10] Y. Ai, H.-C. Wu, and Z.-H. Ling, "SampleRNN-based neural vocoder for statistical parametric speech synthesis," in *Proc. ICASSP*. IEEE, 2018, pp. 5659–5663.
- [11] Y. Ai, J.-X. Zhang, L. Chen, and Z.-H. Ling, "DNN-based spectral enhancement for neural waveform generators with low-bit quantization," in *Proc. ICASSP*, 2019, pp. 7025–7029.
- [12] J. Lorenzo-Trueba, T. Drugman, J. Latorre, T. Merritt, B. Putrycz, and R. Barra-Chicote, "Robust universal neural vocoding," *arXiv preprint arXiv:1811.06292*, 2018.
- [13] L.-J. Liu, Z.-H. Ling, Y. Jiang, M. Zhou, and L.-R. Dai, "WaveNet vocoder with limited training data for voice conversion," in *Proc. Interspeech*, 2018, pp. 1983–1987.
- [14] K. Kobayashi, T. Hayashi, A. Tamamori, and T. Toda, "Statistical voice conversion with WaveNet-based waveform generation," in *Proc. Interspeech*, 2017, pp. 1138–1142.
- [15] Z.-H. Ling, Y. Ai, Y. Gu, and L.-R. Dai, "Waveform Modeling and Generation Using Hierarchical Recurrent Neural Networks for Speech Bandwidth Extension," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 5, pp. 883–894, 2018.
- [16] Y. Cui, X. Wang, L. He, and F. K. Soong, "A new glottal neural vocoder for speech synthesis," in *Proc. Interspeech*, 2018, pp. 2017–2021.
- [17] L. Juvela, V. Tsiaras, B. Bollepalli, M. Airaksinen, J. Yamagishi, and P. Alku, "Speaker-independent raw waveform model for glottal excitation," in *Proc. Interspeech 2018*, 2018, pp. 2012–2016.
- [18] J.-M. Valin and J. Skoglund, "LPCNet: Improving neural speech synthesis through linear prediction," in *Proc. ICASSP*, 2019, pp. 5891–5895.
- [19] 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, 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/8915761/>
- [20] Y. Ai and Z.-H. Ling, "A neural vocoder with hierarchical generation of amplitude and phase spectra for statistical parametric speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 839–851, 2020.
- [21] Y. Ai and Z.-H. Ling, "Knowledge-and-data-driven amplitude spectrum prediction for hierarchical neural vocoders," *arXiv preprint arXiv:2004.07832*, 2020.
- [22] F. Gunnar, *The acoustic theory of speech production*. The Hague, The Netherlands: Mouton, 1960.
- [23] Y. Zhao, X. Wang, L. Juvela, and J. Yamagishi, "Transferring neural speech waveform synthesizers to musical instrument sounds generation," in *Proc. ICASSP*, 2020, pp. 6269–6273.
- [24] J. Engel, L. Hantrakul, C. Gu, and A. Roberts, "DDSP: Differentiable Digital Signal Processing," *Proc. ICLR*, 2020.
- [25] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of Wasserstein GANs," in *Advances in neural information processing systems*, 2017, pp. 5767–5777.
- [26] P. A. Naylor and N. D. Gaubitch, *Speech dereverberation*. Springer Science & Business Media, 2010.
- [27] J. Benesty, M. M. Sondhi, and Y. Huang, *Springer handbook of speech processing*. Springer, 2007.
- [28] C. Valentini-Botinhao and J. Yamagishi, "Speech enhancement of noisy and reverberant speech for text-to-speech," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 8, pp. 1420–1433, 2018.
- [29] K. Kasi and S. A. Zahorian, "Yet another algorithm for pitch tracking," in *Proc. ICASSP*, vol. 1, 2002, pp. 361–364.
- [30] X. Wang and J. Yamagishi, "Using cyclic noise as the source signal for neural source-filter-based speech waveform model," *arXiv preprint arXiv:2004.02191*, 2020.
- [31] N. D. Gaubitch, H. W. Loellmann, M. Jeub, T. H. Falk, P. A. Naylor, P. Vary, and M. Brookes, "Performance comparison of algorithms for blind reverberation time estimation from speech," in *Proc. IWAENC*, 2012, pp. 1–4.
- [32] M. Jeub, "Blind reverberation time estimation," 2015. [Online]. Available: <https://www.mathworks.com/matlabcentral/fileexchange/35740-blind-reverberation-time-estimation>
- [33] S. Buchholz and J. Latorre, "Crowdsourcing preference tests, and how to detect cheating," in *Proc. Interspeech*, 2011, pp. 3053–3056.
- [34] I. Recommendation, "Method for the subjective assessment of intermediate sound quality (MUSHRA)," *ITU, BS*, pp. 1543–1, 2001.