



# Using Cyclic Noise as the Source Signal for Neural Source-Filter-based Speech Waveform Model

Xin Wang<sup>1</sup>, Junichi Yamagishi<sup>1,2</sup>

<sup>1</sup>National Institute of Informatics, Japan

<sup>2</sup>CSTR, University of Edinburgh, UK

wangxin@nii.ac.jp, jyamagis@nii.ac.jp

## Abstract

Neural source-filter (NSF) waveform models generate speech waveforms by morphing sine-based source signals through dilated convolution in the time domain. Although the sine-based source signals help the NSF models to produce voiced sounds with specified pitch, the sine shape may constrain the generated waveform when the target voiced sounds are less periodic. In this paper, we propose a more flexible source signal called cyclic noise, a quasi-periodic noise sequence given by the convolution of a pulse train and a static random noise with a trainable decaying rate that controls the signal shape. We further propose a masked spectral loss to guide the NSF models to produce periodic voiced sounds from the cyclic noise-based source signal. Results from a large-scale listening test demonstrated the effectiveness of the cyclic noise and the masked spectral loss on speaker-independent NSF models in copy-synthesis experiments on the CMU ARCTIC database.

**Index Terms:** speech synthesis, source-filter model, harmonic-plus-noise waveform model, neural network

## 1. Introduction

Neural waveform models are neural network (NN)-based generative models that produce waveforms given acoustic features, linguistic features, other types of input data. In addition to the pioneering work using feedforward networks for speech enhancement [1, 2], many new types of neural waveform models have been proposed recently for text-to-speech (TTS) synthesis [3, 4, 5], band extension [6], music signal generation [7, 8], and other tasks. While new models such as WaveNet [4] and WaveGlow [5] are based on the latest deep learning methods, many other models combine NNs with classical signal processing techniques such as adaptive cepstral analysis [3], linear prediction coding (LPC) [9], and glottal waveform modeling [10]. These models can efficiently produce high-quality waveforms by leveraging signal processing algorithms that cannot be easily learned by neural networks [11].

For TTS, we have proposed a family of neural source-filter (NSF) waveform models [12, 13, 14] that combines the dilated convolution (CONV) NN with the idea of the source-filter speech production model. The NSF models share three common components: a condition module that pre-processes the input spectral features, a source module that produces sine-based source signals from input F0, and a dilated-CONV-based filter module that converts the source signals into output waveforms. Compared with other neural waveform models, NSF models are easy to train and work fast for waveform generation, and our previous study showed their good performance in a Japanese TTS system with a female voice [13].

The source signal is essential to NSF models as these NSF models produce output waveforms by gradually morphing the source signal in the time domain. In our previous works on TTS,

we used sine-based source signals because their periodicity can be accurately maintained in the generated voiced sounds [13]. In addition to the speech waveforms, the sine-based waveform also helps the NSF models to produce high-quality music signals for woodwind, string, and brass instruments [15].

However, for speech waveform generation, the sine-based source signal may cause artifacts when the target sounds are less periodic, e.g., when they are creaky or breathy [16]. To relax the constraint from using a specific waveform shape while maintaining the periodicity, we propose using a cyclic noise as the source signal. We compute this cyclic noise by convolving a pulse train of the specified F0 and a random noise with a tunable decaying rate, which keeps the quasi periodicity of the source signal and allows cyclostationary randomness across periods. We also propose a masked spectral loss that encourages the NSF models to produce periodic voiced sounds from the quasi-periodic cyclic noise. Experimental results demonstrated that the proposed cyclic noise outperformed the sine-based source signal on male speaker data while performing equally well on female speaker data.

After we briefly review the latest NSF model in Section 2, we describe the proposed source signal and loss in Section 3 and explain the experiment results in Section 4. We conclude in Section 5 with a brief summary. Although we focus on NSF models in this paper, we hope that our findings can be useful to related works that use random noises [5, 17, 18], LPC residuals [9], or other types of source signals [19, 20, 21].

## 2. Review of hn-sinc-NSF model

Our focus in this paper is the latest harmonic-plus-noise NSF model with sinc filters (hn-sinc-NSF) [14]. In the hn-sinc-NSF model, which is shown in Figure 1, the input frame-level F0s and Mel-spectrogram are transformed and up-sampled into feature sequences that have the same length  $T$  as the target waveform  $\mathbf{o}_{1:T} = \{o_1, \dots, o_T\}$ . The up-sampled F0  $\mathbf{f}_{1:T} = \{f_1, \dots, f_T\}$  is fed to the source module to produce a source signal  $\mathbf{e}_{1:T}$ , and the processed acoustic feature  $\tilde{\mathbf{c}}_{1:T}$  is fed to neural filter blocks that consist of dilated CONV layers and skip connections. The neural filter blocks form two branches: a harmonic branch that transforms  $\mathbf{e}_{1:T}$  into  $\hat{\mathbf{o}}_{1:T}^{(p)}$  and a noise branch that converts a noise sequence  $\mathbf{a}_{1:T}$  into  $\hat{\mathbf{o}}_{1:T}^{(a)}$ . The two signals  $\hat{\mathbf{o}}_{1:T}^{(p)}$  and  $\hat{\mathbf{o}}_{1:T}^{(a)}$  are then filtered by low-pass and high-pass sinc filters, respectively, and summed as the output waveform  $\hat{\mathbf{o}}_{1:T}$ . The time-variant cut-off frequency of the sinc filters, or maximum voiced frequency (MVF) [22], is predicted from the input acoustic features. The model is trained by minimizing spectral amplitude losses between  $\hat{\mathbf{o}}_{1:T}$  and  $\mathbf{o}_{1:T}$ .

The source signal  $\mathbf{e}_{1:T} = \{e_1, \dots, e_T\}$  in the harmonic branch is based on sine waveforms and can be generated in two steps. First, with the up-sampled F0  $\mathbf{f}_{1:T}$ , a sine waveform

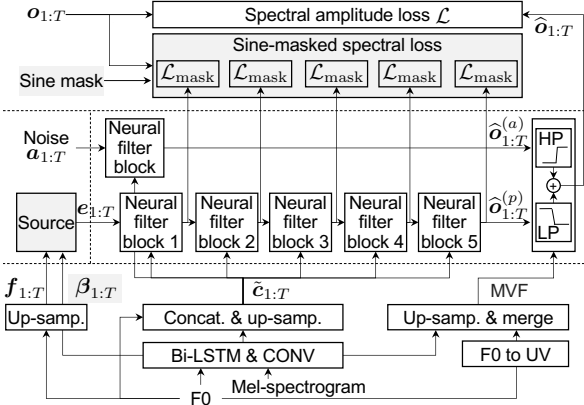


Figure 1: *Hn-sinc-NSF* during training. *LP* and *HP* denote low- and high-pass digital filters, respectively. Up-sampled feature sequence is smoothed except  $\mathbf{f}_{1:T}$ . Compared with [14],  $\tilde{\mathbf{c}}_{1:T}$  and *MVF* are produced from the same *Bi-LSTM* and *CONV* layers to simplify condition module. Shaded components are those proposed in this paper.

$e_{1:T}^{<h>}$  that carries the F0 or the  $h$ -th harmonic is generated by

$$e_t^{<h>} = \begin{cases} \alpha \sin\left(\sum_{k=1}^t 2\pi \frac{hf_k}{N_s} + \phi\right) + n_t, & \text{if } f_t > 0 \\ \frac{\alpha}{3\sigma} n_t, & \text{if } f_t = 0 \end{cases}, \quad (1)$$

where  $f_t = 0$  denotes being unvoiced at the  $t$ -th time step,  $\phi \in [-\pi, \pi]$  is a random initial phase,  $N_s$  is the waveform sampling rate, and  $n_t \sim \mathcal{N}(0, \sigma^2)$  is a Gaussian noise. Then, for each  $t \in \{1, \dots, T\}$ ,  $e_t$  is computed as

$$e_t = \tanh\left(\sum_{h=1}^H w_h e_t^{<h>} + w_b\right), \quad (2)$$

where  $H$  is the number of harmonics and  $\{w_1, \dots, w_H, w_b\}$  are the weights of a trainable feedforward layer with a tanh activation function. In our implementation, we used  $H = 8$ ,  $\sigma = 0.003$ , and  $\alpha = 0.1$  as the hyper-parameters [13]. Note that  $n_t$  in Eq.(1) is independent of the noise source signal  $a_t$  for the noise branch of the *hn-sinc-NSF*, even though both of them can be sampled from the same Gaussian distribution.

### 3. Proposed methods

#### 3.1. Cyclic noise-based source signal

Although a sine-based source is useful for generating periodic waveforms [13], the sine shape may cause perceptual artifacts. For example, it may be inappropriate for creaky voices with irregular glottal pulses [16, 23] and low-pitch voices where the shape of the source signal has a greater perceptual effect [24, 25]. We hypothesize that a better NSF source signal for voiced sounds may contain a certain degree of randomness in the short term while preserving long-term periodicity. Although source signals for classical speech vocoders [26, 27, 28, 29] may be used, we focus on source signals that have a simple parametric form and require no additional analysis loop.

Inspired by an experimental study on speech perception [25], we propose a source signal called *cyclic noise*, which has a simple parametric form and can preserve both periodicity and

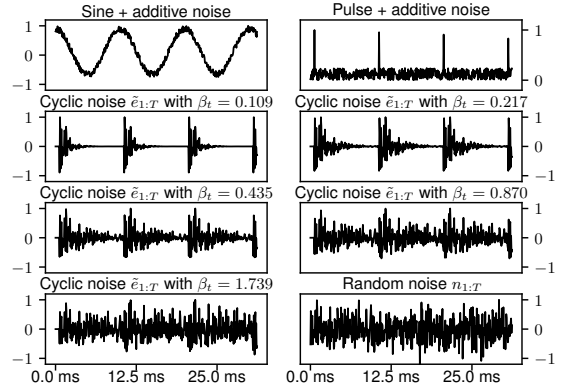


Figure 2: Example of different source signals. Maximum signal amplitude is scaled to 1 for visualization.

randomness. Let  $\mathbf{p}_{1:T} = [p_1, \dots, p_T]$  denote a unit impulse train corresponding to the up-sampled F0  $\mathbf{f}_{1:T}$ , and let  $\mathbf{n}_{1:T}$  denote a Gaussian noise sequence. The proposed source signal  $e_t$  at the  $t$ -th time step is computed in two steps:

$$\tilde{e}_t = \begin{cases} \sum_{k=0}^{t-1} n_{k+1} \exp\left(-\frac{kf_t}{\beta_t N_s}\right) p_{t-k}, & \text{if } f_t > 0 \\ n_t, & \text{if } f_t = 0 \end{cases}, \quad (3)$$

and

$$e_t = \tanh(w_1 \tilde{e}_t + w_b). \quad (4)$$

Equation (4) is similar to Eq. (2) but uses only  $\tilde{e}_t$  rather than multiple signals with higher harmonics. While  $\tilde{e}_t$  in the unvoiced regions is still random noise,  $\tilde{e}_t$  in the voiced regions is the convolution of  $\mathbf{p}_{1:T}$  and an exponentially decaying noise whose value is  $n_{k+1} \exp(-\frac{kf_t}{\beta_t N_s})$ ,  $\forall k \in [0, t-1]$ . The factor  $\frac{f_t}{N_s}$  makes the decaying rate independent of the period length or F0 — no matter how long the interval between two pulses is, the noise value is scaled by  $\exp(-\frac{1}{\beta_t})$  when  $k$  is equal to the period length  $\frac{N_s}{f_t}$ . In such a manner,  $\beta_t$  becomes the main parameter that controls the decaying rate.

By adjusting  $\beta_t$ , we can change the randomness and periodicity of the cyclic noise-based source signal for voiced sounds. Figure 2 illustrates a few source signals with different values of  $\beta_t$  for  $t \in [1, T]$ . Note that when  $\beta_t = 0.435$ , the noise amplitude at  $t = \frac{N_s}{0.435}$  is scaled by  $\exp(-\frac{1}{0.435}) \approx 0.1$ . When  $\beta_t$  is smaller, the random noise decays more quickly, and the source signal is more similar to the pulse train. When  $\beta_t$  is increased, the source signal becomes noisy and less periodic.

We produce  $\mathbf{p}_{1:T}$  by setting  $p_t = 1$ , where  $t$  corresponds to the local maxima of the sine waveform that carries the F0. Otherwise,  $p_t = 0$ . The noise sequence  $\mathbf{n}_{1:T}$  is drawn from the same Gaussian distribution as that in Eq. (1), and it changes for every input training sample.

#### 3.2. Predicting $\beta_t$ from acoustic features

We may treat the decaying rate  $\beta$  as a hyper-parameter and decide its value through trial and error. Alternatively, we can use a small neural network to predict the value of  $\beta_t$  from the input acoustic features for each time step  $t$ .

Our implementation is illustrated in Figure 1. We use the *Bi-LSTM* and *CONV* layers in the condition module to predict a one-dimensional signal at the frame level. This signal is then up-sampled and smoothed as  $\beta_{1:T} = \{\beta_1, \dots, \beta_T\}$ .

The smoothing is conducted twice by moving averaging over a window of size 320. During training, an  $L_1$  norm  $|\beta_t - 0.870|_1$  with weight 0.01 is added to prevent  $\beta_t$  from going wild. The value 0.870 is used because  $\beta_t = 0.870$  keeps periodicity and randomness, as shown in Figure 2.

### 3.3. Additional masked spectral loss to stabilize pitch

In our experiments, we observed that the randomness in the cyclic noise made it more difficult to produce a waveform with stable pitch. Since the source signal is gradually morphed by the NSF neural filter blocks, it may be helpful if additional loss can be defined to penalize the mismatch of the harmonic structure between the natural waveform and the output of the filter blocks.

We propose a masked spectral loss for this purpose. Let  $\mathbf{y}^{(n)} = [y_1^{(n)}, \dots, y_K^{(n)}]^\top \in \mathbb{C}^K$  denote the complex-valued spectrum of the  $n$ -th frame of the natural waveform calculated using  $K$ -points FFT. Similarly, let  $\hat{\mathbf{p}}^{(n)}$  and  $\mathbf{m}^{(n)}$  denote the spectrum calculated on a neural filter block’s output signal and a mask signal, respectively. The proposed loss is defined as

$$\mathcal{L}_{\text{mask}} = \frac{1}{2NK} \sum_{n=1}^N \sum_{k=1}^K \left[ \log \frac{|y_k^{(n)}|^2 |m_k^{(n)}|^2 + \eta}{|\hat{p}_k^{(n)}|^2 |m_k^{(n)}|^2 + \eta} \right]^2, \quad (5)$$

where  $N$  is the number of frames and  $\eta = 1e-5$  is a constant to ensure numerical stability. The mask signal is defined as  $\frac{1}{H} \sum_{h=1}^H e_{1:T}^{<h>}$ , where  $e_{1:T}^{<h>}$  is the sine waveform that carries the  $i$ -th harmonics of the input F0 (see Eq. (1)). The spectrum amplitude of such a mask signal, i.e.,  $|m^{(n)}|^2$ , has peaks of roughly equal height at the harmonic frequencies, and  $\mathcal{L}_{\text{mask}}$  thus penalizes the mismatch of the harmonic structure rather than the whole spectral envelope, as Figure 3 illustrates.

The masked spectral loss  $\mathcal{L}_{\text{mask}}$  can be computed using multiple STFT configurations in the same manner as the unmasked spectral amplitude loss  $\mathcal{L}$  between  $\mathbf{o}_{1:T}$  and  $\hat{\mathbf{o}}_{1:T}$ . We used the STFT configurations in our previous work (Table III in [13]) for both  $\mathcal{L}$  and  $\mathcal{L}_{\text{mask}}$ . Further, as Figure 1 plots,  $\mathcal{L}_{\text{mask}}$  is measured between the natural waveform and the outputs of the five neural filter blocks (i.e., block 1 to block 5 in Figure 1).

## 4. Experiments

### 4.1. Data and feature configuration

We used the data of four US English speakers from the CMU ARCTIC database [30] for our experiments: SLT (female), CLB (female), BDL (male), and RMS (male). Each speaker provided 1000, 64, and 64 utterances for the training, validation, and test sets, respectively, and the data were pooled together to train the experimental models in a speaker-independent manner.

The acoustic features input to the experimental models included the Mel-spectrogram of 80 dimensions and the F0 sequence, both of which were extracted with a frame shift of 5 ms. The target speech waveforms had a sampling rate of 16 kHz and were equalized to -26 dBov using SV56 [31]. During the test stage, the natural acoustic features were used to produce the waveforms. Speaker vectors were not used.

### 4.2. Experimental models

We utilized the hn-sinc-NSF models listed in Table 1. Except the source signal and  $\mathcal{L}_{\text{mask}}$ , all these models used the network architecture detailed in our previous work (i.e., sinc1-h-NSF in [14]): five and one neural filter blocks in the harmonic and noise branches, respectively, ten dilated-CONV layers in each




Figure 3: Sine-masked spectral loss.  $\mathbf{y}^{(n)}$ ,  $\hat{\mathbf{p}}^{(n)}$ , and  $\mathbf{m}^{(n)}$  denote the spectrum of natural waveform, output signal from a neural filter block, and sine mask, respectively, at  $n$ -th frame.

block, and sinc filters of order 31. The  $k$ -th dilated-CONV layer in each neural filter block had a dilation size of  $2^{k-1}$ . The experiments included the natural waveforms (Natural) and WaveNet (WaveNet) as a reference. WaveNet used 30 dilated CONV layers, where the  $k$ -layer had a dilation size of  $2^{\text{mod}(k-1,10)}$ . It used the Mel-spectrogram as the input features and generated 10-bit  $\mu$ -law quantized waveform values.

All the experimental models were trained for 100 epochs using an ADAM optimizer with a learning rate 0.0003,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and  $\epsilon = 10^{-8}$  [32]. The model parameters from the epoch with the lowest validation error were saved.

### 4.3. Results and analysis

Each trained experimental model generated 264 test set waveforms (64 utterances per speaker). The quality of the generated waveforms was evaluated by means of a large-scale crowd-sourced listening test. All the waveforms were converted to a 16-bit PCM format for the test<sup>1</sup>.

Each test round, an evaluator was asked to complete 22 test pages. The evaluator listened to one waveform on each page and rated its quality on a 1-to-5 mean opinion score (MOS) scale. The 22 test pages contained two randomly selected natural waveforms and two generated waveforms from each of the experimental models from the same speaker. The page order was randomly shuffled, and the waveform on each page could be played multiple times. Around 710 evaluators participated in the listening test and 792 test rounds were conducted. Accordingly, each experimental model received 1584 MOS scores, the boxplots of which are shown in Figure 5.

The results reveal interesting cross-gender differences. As the scores of Sin, Pul, and Rno demonstrate, the sine and pulse were better than the random noise for female speakers while the random noise outperformed the sine and pulse for male speakers. The differences in both cases were statistically significant (Mann-Whitney two-sided test,  $p < 0.01$ )<sup>2</sup>. While cross-gender differences have also been observed on the source signals for classical speech vocoders [25, 33, 34], new findings can be observed here. For male voices, for example, Rno produced periodic voiced sounds, and they were perceptually better than those using sine or pulse. The performance of Rno suggests that the neural filter blocks can transform random noise into periodic waveforms with long periodic cycles, which is difficult for linear filters in classical vocoders. For the female voice, we found that it was still difficult for neural filter blocks to produce periodic waveforms with a short period cycle, so it is better to use periodic source signals.

<sup>1</sup>Code, models, and samples: <https://nii-yamagishilab.github.io/samples-nsf/nsf-v4.html>.

<sup>2</sup>Results of significant tests are uploaded to the website above.

Table 1: *Experimental models based on hn-sinc-NSF.*

Model	Type of source signal	$\mathcal{L}_{mask}$
Sin	Sine-based waveform (Eq. (1-2))	-
Pul	Pulse with noise (Gaussian)	-
Rno	Random noise (Gaussian)	+
Cno $_{\beta_1}$	Cyclic noise, $\beta_t = 0.435$	+
Cno $_{\beta_2}$	Cyclic noise, $\beta_t = 0.870$	+
Cno $_{\beta_3}$	Cyclic noise, $\beta_t = 1.739$	+
Cno $_{\beta_t}$	Cyclic noise, trainable $\beta_t$	+
Rno $_{no\mathcal{L}_{mask}}$	Random Gaussian noise	-
Cno $_{\beta_2no\mathcal{L}_{mask}}$	Cyclic noise, $\beta_t = 0.870$	-




Figure 4: *Waveform sample from BDL (arctic\_b0477).*

The cyclic noise worked well for both the female and male voices when a good  $\beta$  was selected. For the female speakers SLT and CLB, the cyclic noises in Cno $_{\beta_2}$  and Cno $_{\beta_3}$  were quasi-periodic (as Figure 2 illustrates) and may be sufficient for generating periodic waveforms. For male speakers RMS and BDL, although the difference is not always statistically significant, Cno $_{\beta_2}$  and Cno $_{\beta_3}$  were perceptually better than Sin on many test utterances because the generated waveforms in each periodic cycle had one obvious spike in the time domain, which is similar to how natural waveforms behave (as Figure 4 shows). In contrast, the waveform from Sin scattered within each periodic cycle in the time domain. Cno $_{\beta_1}$  performed worse on male voices and its generated speech was perceptually similar to those from Sin.

Interestingly, compared with the LPC analysis on low pitch voices (e.g., M4 in [25]), the waveforms from Cno $_{\beta_2}$  and Cno $_{\beta_3}$  were perceptually similar to those of using a zero-phase source signal in LPC, while the waveforms from Sin sounded similar to the LPC-produced waveforms using cyclostationary random phase source signals, even though the source signals in Cno $_{\beta_2}$  and Cno $_{\beta_3}$  were cyclostationary rather than zero-phased. The discrepancy from the LPC analysis again may be due to the difference between NNs and linear filters in LPC.

Although we used Cno $_{\beta_t}$  to infer time-variant and speaker-dependent  $\beta_t$ , the predicted  $\beta_t$  was roughly equal to  $\beta = 0.86$  for all  $t$  and all speakers. This may be due to the difficulty of optimization. Another reason could be that a better spectral loss is necessary to evaluate the difference caused by the varying  $\beta$ .

As a final remark, Cno $_{\beta_2}$  outperformed Cno $_{\beta_2no\mathcal{L}_{mask}}$  for all four speakers, and the difference was statistically significant. This result indicates the usefulness of the masked spectral loss  $\mathcal{L}_{mask}$  for Cno $_{\beta_2}$ . However,  $\mathcal{L}_{mask}$  was less effective for Rno, especially on female speakers. We conclude that  $\mathcal{L}_{mask}$  is insufficient to guide the model to reproduce stable pitch from noise source signals for female speakers.

## 5. Conclusion

We proposed a new type of source signal called cyclic noise for NSF models. Cyclic noise has a simple parametric form based on the convolution of a pulse train and an exponentially decaying noise sequence. We demonstrated through experiments that an appropriate  $\beta$  can be specified

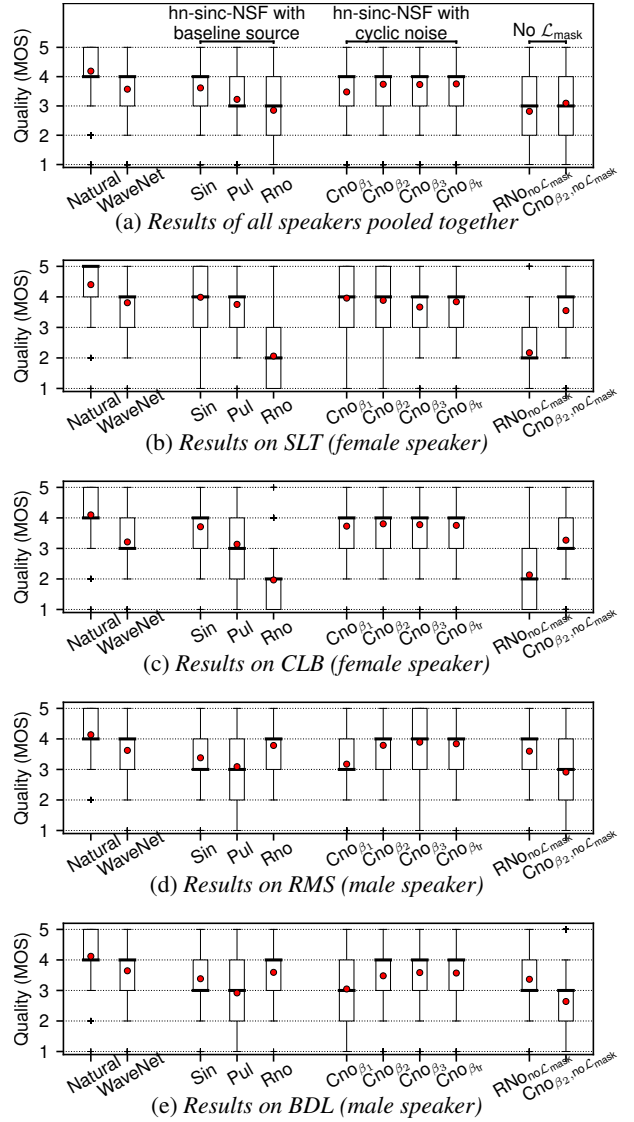


Figure 5: *Boxplots of MOS. Dot denotes mean value of MOS.*

to maintain the periodicity and introduce randomness to the cyclic noise signal. Compared with the sine-based source signal, which worked well for female speakers but not for male speakers, the proposed cyclic noise is suitable for both female and male voices.

Although the findings in this paper were derived from NSF models, they showed interesting differences from those on the classical vocoders. Future work will examine the link between the perceptual difference and the phase [25, 34, 35].

## 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 Google AI for Japan program. The experiments partially were conducted on TSUBAME 3.0 of Tokyo Institute of Technology. We thank Mr. Yusuke Yasuda for setting up the server for listening test. We also thank Dr. Heiga Zen and Dr. Yotaro Kubo for comments and discussions.

## 7. References

- [1] S. Tamura and A. Waibel, "Noise reduction using connectionist models," in *Proc. ICASSP*, 1988, pp. 553–554.
- [2] X.-M. Gao, S. J. Ovaska, and Z. O. Hartimo, "Speech signal restoration using an optimal neural network structure," in *Proceedings of International Conference on Neural Networks (ICNN'96)*, vol. 4. IEEE, 1996, pp. 1841–1846.
- [3] K. Tokuda and H. Zen, "Directly modeling speech waveforms by neural networks for statistical parametric speech synthesis," in *Proc. ICASSP*. IEEE, 2015, pp. 4215–4219.
- [4] 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.
- [5] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A Flow-based Generative Network for Speech Synthesis," in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [6] V. Kuleshov, S. Z. Enam, and S. Ermon, "Audio super-resolution using neural nets," in *ICLR (Workshop Track)*, 2017.
- [7] J. Engel, C. Resnick, A. Roberts, S. Dieleman, M. Norouzi, D. Eck, and K. Simonyan, "Neural audio synthesis of musical notes with WaveNet autoencoders," in *Proc. ICML*, 2017, pp. 1068–1077. [Online]. Available: <http://proceedings.mlr.press/v70/engel17a.html>
- [8] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, "GANSynth: Adversarial neural audio synthesis," *Proc. ICLR*, 2019.
- [9] J.-M. Valin and J. Skoglund, "LPCNet: Improving Neural Speech Synthesis Through Linear Prediction," in *Proc. ICASSP*, 2019, pp. 5891–5895.
- [10] L. Juvela, B. Bollepalli, V. Tsiaras, and P. Alku, "GlottNet - a raw waveform model for the glottal excitation in statistical parametric speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 6, pp. 1019–1030, 2019.
- [11] J. Engel, L. Hantrakul, C. Gu, and A. Roberts, "DDSP: Differentiable Digital Signal Processing," *Proc. ICLR*, 2020.
- [12] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter-based waveform model for statistical parametric speech synthesis," in *Proc. ICASSP*, 2019, pp. 5916–5920.
- [13] —, "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/>
- [14] 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, sep 2019, pp. 1–6. [Online]. Available: [http://www.isca-speech.org/archive/SSW\\_2019/abstracts/SSW10\\_O.1-1.html](http://www.isca-speech.org/archive/SSW_2019/abstracts/SSW10_O.1-1.html)
- [15] Y. Zhao, X. Wang, L. Juvela, and J. Yamagishi, "Transferring neural speech waveform synthesizers to musical instrument sounds generation," *Proc. ICASSP*, pp. 6269–6273, may 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9053047/>
- [16] M. Gordon and P. Ladefoged, "Phonation types: a cross-linguistic overview," *Journal of phonetics*, vol. 29, no. 4, pp. 383–406, 2001.
- [17] 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*. IEEE, 2020, pp. 6199–6203.
- [18] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "MelGAN: Generative adversarial networks for conditional waveform synthesis," in *Proc. NIPS*, 2019, pp. 14 910–14 921.
- [19] O. Watts, C. Valentini-Botinhao, and S. King, "Speech waveform reconstruction using convolutional neural networks with noise and periodic inputs," in *Proc. ICASSP*. IEEE, 2019, pp. 7045–7049.
- [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] L. Juvela, B. Bollepalli, J. Yamagishi, and P. Alku, "GELP: GAN-excited linear prediction for speech synthesis from Mel-spectrogram," *Proc. Interspeech*, pp. 694–699, 2019.
- [22] Y. Stylianou, "Applying the harmonic plus noise model in concatenative speech synthesis," *IEEE Transactions on speech and audio processing*, vol. 9, no. 1, pp. 21–29, 2001.
- [23] T. Drugman, J. Kane, and C. Gobl, "Data-driven detection and analysis of the patterns of creaky voice," *Computer Speech & Language*, vol. 28, no. 5, pp. 1233–1253, 2014.
- [24] A. E. Rosenberg, "Effect of glottal pulse shape on the quality of natural vowels," *The Journal of the Acoustical Society of America*, vol. 49, no. 2B, pp. 583–590, 1971.
- [25] T. Raitio, L. Juvela, A. Suni, M. Vainio, and P. Alku, "Phase perception of the glottal excitation and its relevance in statistical parametric speech synthesis," *Speech Communication*, vol. 81, pp. 104–119, 2016.
- [26] J. P. Cabral, S. Renals, K. Richmond, and J. Yamagishi, "Towards an improved modeling of the glottal source in statistical parametric speech synthesis," *Proc. SSW6*, pp. 113–118, 2007.
- [27] R. Maia, T. Toda, H. Zen, Y. Nankaku, and K. Tokuda, "An excitation model for HMM-based speech synthesis based on residual modeling," in *Proc. SSW6*, 2007, pp. 131–136.
- [28] T. Drugman and T. Dutoit, "The deterministic plus stochastic model of the residual signal and its applications," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 3, pp. 968–981, 2011.
- [29] T. Raitio, A. Suni, J. Yamagishi, H. Pulakka, J. Nurminen, M. Vainio, and P. Alku, "HMM-based speech synthesis utilizing glottal inverse filtering," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 1, pp. 153–165, 2011.
- [30] J. Kominek and A. W. Black, "The CMU Arctic speech databases," in *Fifth ISCA workshop on speech synthesis*, 2004.
- [31] R. ITU-T, "P. 56: Objective Measurement of Active Speech Level," *International Telecommunication Union*, 2011. [Online]. Available: <https://www.itu.int/rec/T-REC-P.56-201112-I/en>
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2014.
- [33] T. Drugman and T. Raitio, "Excitation modeling for HMM-based speech synthesis: breaking down the impact of periodic and aperiodic components," in *Proc. ICASSP*, 2014, pp. 260–264.
- [34] J. Skoglund, W. B. Kleijn, and P. Hedelin, "Audibility of pitch-synchronously modulated noise," in *Proc. IEEE Workshop on Speech Coding for Telecommunications Proceedings. Back to Basics: Attacking Fundamental Problems in Speech Coding*. IEEE, 1997, pp. 51–52.
- [35] H. Pobloth and W. B. Kleijn, "On phase perception in speech," in *Proc. ICASSP*, vol. 1. IEEE, 1999, pp. 29–32.