Speaker-Adaptive Neural Vocoders for Parametric Speech Synthesis Systems

Eunwoo Song[1], Jinseob Kim[1], Kyungguen Byun[2] and Hong-Goo Kang[2]

[1]NAVER Corp., Seongnam, Korea

[2]Department of Electrical and Electronic Engineering, Yonsei University, Seoul, Korea

eunwoo.song@navercorp.com

Abstract

This paper proposes speaker-adaptive neural vocoders for statistical parametric speech synthesis (SPSS) systems. Recently proposed WaveNet-based neural vocoding systems successfully generate a time sequence of speech signal with an autoregressive framework. However, it remains a challenge to build high-quality speech synthesis systems when the amount of a target speaker's training data is insufficient. To generate more natural speech signals with the constraint of limited training data, we propose a speaker adaptation task with an effective variation of neural vocoding models. In the proposed method, a speaker-independent training method is applied to capture universal attributes embedded in multiple speakers, and the trained model is then optimized to represent the specific characteristics of the target speaker. Experimental results verify that the proposed SPSS systems with speaker-adaptive neural vocoders outperform those with traditional source-filter model-based vocoders and those with WaveNet vocoders, trained either speaker-dependently or speaker-independently.

Index Terms: statistical parametric speech synthesis, WaveNet, ExcitNet, neural vocoder, speaker adaptation

1. Introduction

Waveform generation systems using WaveNet have attracted a great deal of attention in the speech signal processing community thanks to their high quality and ease of use in various applications [1, 2]. In a system of this kind, the time domain speech signal is represented as a sequence of discrete symbols, and its distribution is autoregressively modeled by stacked convolutional neural networks (CNNs). By appropriately conditioning the acoustic features to the input, WaveNet-based systems have also been successfully adopted in a neural vocoder structure for statistical parametric speech synthesis (SPSS) systems [3–7].

To further improve the perceptual quality of the synthesized speech, more recent neural excitation vocoders (e.g. ExcitNet [8]) take advantages of the merits from both the parametric LPC vocoder and the WaveNet structure [9–13]. In this framework, an adaptive predictor is used to decouple the formant-related spectral structure from the input speech signal, and the probability distribution of its residual signal (i.e. the excitation signal) is then modeled by the WaveNet network. As variation in the excitation signal is only constrained by vocal cord movement, the training and generation processes become much more efficient. As such, SPSS systems with neural excitation vocoders reconstruct more accurate speech signals than conventional parametric or WaveNet vocoders [8].

However, this approach still requires large amounts of training data to faithfully represent the complex mechanics of human speech production. As a result, unnatural outputs are generated when the training data for the target speaker is insufficient (e.g. a database comprising less than ten minutes' speech).

The speaker-independent training method that utilizes multiple speakers for a single unified network shows the feasibility of generating diverse characteristics of voices by conditioning the target speaker's acoustic features [4]. However, our preliminary experiments verify that this approach still generates discontinuous speech segments if the target speaker's data is not included in the training process. This problem is more prominent under an SPSS framework where prediction errors in estimating auxiliary parameters are inevitable; prediction errors are propagated throughout the autoregressive generation process.

To alleviate this problem, we propose a speaker-adaptive training method for neural vocoding systems. In this framework, to address the lack of speaker-specific information caused by limited training data for a target speaker, a model is trained independently of the target speaker such that it extracts universal attributes from multiple speakers [4]. This model is then used to initialize the training model of the target speaker, and all weights are fine-tuned to represent the distinctive characteristics within the target's database. Because this adaptation process helps the CNNs capture speaker-specific characteristics, it is also advantageous in reducing the discontinuity problems that occur in conventional speaker-independent models.

We investigate the effectiveness of the proposed method by conducting objective and subjective evaluations with systems designed both dependently and independently of the target speaker. The merits of the proposed method can be found in its robust performance in a pitch modification task because its initial model shares the diverse characteristics extracted from multiple speech databases. Furthermore, in speech analysis and synthesis, and under SPSS conditions, the experimental results show that the speaker-adaptive neural vocoder significantly improves the perceptual quality of synthesized speech compared to conventional methods.

2. Relationship to prior work

The idea of using WaveNet in neural vocoders for speech synthesis frameworks is not new. By effectively representing sample-by-sample correspondence between acoustic features and speech waveform, the WaveNet vocoder has successfully replaced the role of traditional parametric vocoders [3–7]. However, the speech signals generated by neural vocoders often suffer from unnatural outputs when the amount of training data for the target speaker is insufficient. Although employing a multi-speaker training method enables the general speech signal characteristics to be represented [4], capturing the specific nature of the target speaker's own speech is limited.

To ameliorate these issues, our aim is to use an adaptation task to improve the modeling accuracy of WaveNet-based neural vocoders. Although prior works in using speaker-adaptive WaveNet vocoders in voice conversion applications have been undertaken [14, 15], our research differs from these studies in

Figure 1: *ExcitNet vocoder framework for an SPSS system: (a) training and (b) synthesis.*

Figure 2: *Negative log-likelihood (NLL) obtained during the training process with (w/) and without (w/o) adaptation process.*

several ways: First, we focus on the effect of the speaker adaptation in SPSS tasks. In this study, we verify the effectiveness of the proposed method not only in speech analysis/synthesis but also within an SPSS framework. Second, our experiments seek to verify the superior performance of speaker-adaptive training methods over conventional speaker-dependent and speaker-independent approaches. Furthermore, the synthesis quality of each training method is investigated across various types of vocoder, for example, a plain WaveNet framework and a neural excitation vocoder, namely ExcitNet. Both the objective and subjective test results provide helpful guidelines for the design of similarly configured vocoding systems. Third, in terms of the perceptual quality of the vocoding, the proposed method shows superiority over the other approaches under the same SPSS model structure. And finally, we also explore the effectiveness of the proposed method in a fundamental frequency (F0) modification task. Experiments in arbitrary changes to F0 contours confirm that the proposed speaker-adaptive training method synthesizes the modified F0 sound very reliably compared to conventional speaker-dependent approaches.

3. Neural vocoders

3.1. WaveNet-based neural vocoding frameworks

The basic WaveNet framework is an autoregressive network which generates a probability distribution of waveforms from a fixed number of past samples [2]. Recent *WaveNet vocoders* directly utilize acoustic features as the conditional input where these features are extracted from conventional parametric vocoders [3–7]. This enables the WaveNet system to automatically learn the relationship between acoustic features and speech samples which results in superior perceptual quality over traditional parametric vocoders [3, 16].

However, due to the inherent structural limitations of CNNs in terms of capturing the dynamic nature of speech signals, the WaveNet approach often generates noisy outputs caused by distortion in the spectral valley regions. To improve the perceptual quality of synthesized speech, several frequency-dependent noise-shaping filters have been proposed [8–13]. In particular, the neural excitation vocoder *ExcitNet* (described in Figure 1a) exploits a linear prediction (LP)-based adaptive predictor to decouple the spectral formant structure from the input speech signal. The WaveNet-based generation model is then used to train the residual LP component (i.e. the excitation signal). As variation in the excitation signal is only constrained by vocal cord movement, the training process becomes much more effective.

In the speech synthesis step shown in Figure 1b, the acoustic parameters of the given input are first generated by an acoustic model designed with a conventional deep learning SPSS system [17]. Those parameters are used as auxiliary conditional

features for the WaveNet model to generate the corresponding time sequence of the excitation signal. Ultimately, the speech signal is reconstructed by passing the generated excitation signal through the LP synthesis filter. In this way, the quality of the synthesized speech signal is further improved because the spectral component is well represented by the deep learning framework and the residual component is efficiently generated by the WaveNet framework.

3.2. Speaker-adaptive neural vocoders

The superiority of neural vocoding systems over traditional parametric vocoders has been explained above but it is still challenging to build a high-quality speech synthesis system when the training data for a target speaker is insufficient, for example with just ten minutes of speech.

To generate a more natural speech signal with limited training data, we employ an adaptation task in training the neural vocoders. In the proposed framework, a speaker-independently trained multi-speaker model is used as the initializer, and then all weights are updated in training the target speaker's model. As the initial model already represents global characteristics embedded in the multiple speakers quite well [4], the fine-tuning mechanism only needs to capture speaker-specific characteristics from the target's data set. Consequently, the entire learning process becomes much more effective. Fig. 2 shows the negative log-likelihood obtained during the training phase, of which results confirm that the proposed method significantly reduces both training and development errors as compared to the system without having an adaptation process.

4. Experiments

4.1. Experimental setup

To investigate the effectiveness of the proposed algorithm, we trained neural vocoding models using three different methods:

- **SD**: speaker-dependent training model
- **SI**: speaker-independent training model
- **SA**: speaker-adaptive training model

In the SD and SA models, speech corpora recorded by Korean male and Korean female speakers were used. The speech signals were sampled at 24 kHz, and each sample was quantized by 16 bits. Table 1 shows the number of utterances in each set. To train the SI model, speech corpora recorded by five Korean male and five Korean female speakers not included in training the SD and SA models were used. For this, 6,422 (10 h) and 1,080 (1.7 h) utterances were used for training and development, respectively. The testing set in the SD and SA models was also used to evaluate the SI model.

To compose the acoustic feature vectors needed for auxiliary input information, the spectral and excitation parameters were extracted using a previously proposed parametric ITFTE vocoder [17]. In this way, 40-dimensional line spectral frequencies (LSFs), 32-dimensional slowly evolving waveform (SEW), 4-dimensional rapidly evolving waveform (REW), the F0, gain, and v/uv were extracted. The frame and shift lengths were set to 20 ms and 5 ms, respectively.

In the WaveNet training step, all acoustic feature vectors were duplicated from a frame to the samples to match the length of the input speech signals [3]. Before training, they were normalized to have zero mean and unit variance. The corresponding speech signal was normalized in in a range between -1.0 and 1.0 and encoded by 8-bit-μ compression. The WaveNet architecture comprised of three convolutional blocks, each with ten dilated convolution layers with dilations of 1, 2, 4, and so on up to 512. The number of channels of dilated causal convolution and the 1×1 convolution in the residual block were both set to 512. The number of 1×1 convolution channels between the skip-connection and the softmax layer was set to 256. The learning rate was set to 0.0001, and the batch size was set to 30,000 (1.25 sec).

To train the SI-WaveNet model, all data from the multiple number of different speakers were used; the sequence of each batch was randomized across all speakers before input to the training process. The weights were initialized using *Xavier* initialization and *Adam* optimization was used [18, 19]. The training methods of the SD- and SA-WaveNets were similar but the initialization process was different in each case the SD model was initialized by Xavier initialization whereas the SA-WaveNet was initialized using the SI-WaveNet model whose weights were optimized toward the target speaker's database to represent speaker-specific characteristics.

To construct a baseline SPSS system, we employed a *shared hidden layer (SHL) acoustic model* [20, 21]. The linguistic input feature vectors were 356-dimensional contextual information consisting of 330 binary features of categorical linguistic contexts and 26 features of numerical linguistic contexts. The output vectors consisted of all the acoustic parameters together with their time dynamics [22]. Before training, both input and output features were normalized to have zero mean and unit variance. The SHL consisted of three feedforward layers with 1,024 units and one long short-term memory layer with 512 memory blocks. The weights were trained using a *backpropagation through time* algorithm with Adam optimization [23].

In the synthesis step, the means of all acoustic features were predicted by the SHL model first, then a speech parameter generation algorithm was applied with the pre-computed global variances [24, 25]. To enhance spectral clarity, an LSF-sharpening filter was also applied to the spectral parameters [17]. To reconstruct the speech signal, the generated acoustic features were used to compose the input auxiliary features. By conditioning these features, the WaveNet generated discrete symbols corresponding to the quantized speech signal, and its dynamic was recovered via μ-law expansion.

The setups for training the SI-, SD-, and SA-ExcitNets were the same as those for the WaveNets but the ExcitNet-based framework predicted the distribution of the excitation signal, obtained by passing the speech signal through the LP analysis filter. Similar to the WaveNet vocoder, the ExcitNet vocoder generated the excitation sequence in the synthesis step. Ultimately, the speech signal was reconstructed through an LP synthesis filter.

Table 1: *Number of utterances in different sets for the Korean male (KRM) and the Korean female (KRF) speakers (SPKs).*

SPK	Training	Development	Test
KRM	55 (10 min)	25 (5 min)	80 (15 min)
KRF	90 (10 min)	40 (5 min)	130 (15 min)

Table 2: *LSD (dB) and F0 RMSE (Hz) results for the Korean male (KRM) and the Korean female (KRF) speakers (SPKs): the smallest errors are in bold.*

SPK	System	WaveNet		ExcitNet	
		LSD	F0 RMSE	LSD	F0 RMSE
KRM	SD	4.37	21.30	3.93	14.83
	SI	4.06	14.76	3.86	14.39
	SA	**4.03**	**14.16**	**3.82**	**14.03**
KRF	SD	4.78	48.75	4.50	39.14
	SI	4.51	35.53	4.42	36.28
	SA	**4.45**	**35.45**	**4.36**	**35.47**

Table 3: *Objective test results in the large-scale (7 hours) adaptation: the smallest errors are in bold.*

SPK	System	WaveNet		ExcitNet	
		LSD	F0 RMSE	LSD	F0 RMSE
KRM	SD	3.63	11.28	3.37	12.08
	SI	3.66	12.08	3.40	11.65
	SA	**3.58**	**11.04**	**3.33**	**10.93**
KRF	SD	4.08	28.75	3.95	28.76
	SI	4.13	28.77	4.00	28.84
	SA	**4.03**	**28.57**	**3.90**	**28.36**

4.2. Objective test results

4.2.1. Statistical parametric speech synthesis

To verify the performance of the proposed method, we measured distortions between the original speech and the synthesized speech with log-spectral distance (LSD; dB) and F0 root mean square error (RMSE; Hz) measures. Table 2 presents the test results with respect to the different types of training methods. The findings can be outlined as follows: (1) The proposed SA training method reconstructs much more accurate speech signals than the SD and SI models in both WaveNet and ExcitNet vocoders. (2) Among the different vocoding systems, the ExcitNet-based framework performed significantly better than the WaveNet-based one in terms of spectral distortion because the adoption of an adaptive spectral filter for the ExcitNet vocoder is beneficial for more accurately modeling the target speech signals.

To verify the effectiveness of the proposed algorithm in a *large amount* of training database condition, additional experiments were conducted by changing the adaptation data size from 10 minutes to 7 hours. For the comparison, the amount of database to train the SD model and the SHL acoustic model was also increased to 7 hours. Table 3 shows the test results, which confirms that adapting the vocoding model toward target speaker's database is still advantageous to improve the modeling accuracy regardless of the amount of training data set.

4.2.2. Speech modification

To further verify the effectiveness of the proposed SA training method, we investigated the performance variation of neural vocoders when F0 is manually modified. It has already been

Figure 3: *F0 RMSE (Hz) results with respect to different values of the scaling factor: (a) Korean male and (b) Korean female speakers.*
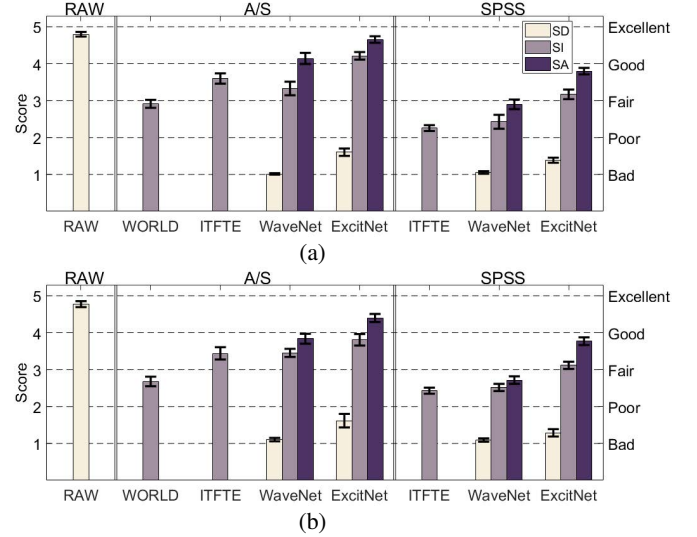

Figure 4: *MOS results with 95% confidence intervals. Acoustic features extracted from recorded speech and generated from an acoustic model were used to compose the input auxiliary features in the A/S and SPSS tests, respectively: (a) Korean male and (b) Korean female speakers.*

shown that the SI model effectively generates pitch-modified synthesized speech [4]. As the entire network of the SA approach in the present study was adapted from an SI model, it was expected to further improve performance compared to conventional SD approaches.

In this experiment, the F0 trajectory was first generated by the SPSS framework and then multiplied by a scaling factor (0.6, 0.8, 1.0, and 1.2) to modify the auxiliary feature vectors. Finally, the speech signal was synthesized using the neural vocoding systems. Figure 3 illustrates the F0 RMSE (Hz) test results with respect to the different values of scaling factor. The results can be analyzed as follows: (1) The proposed SA training models result in much smaller modification errors than the conventional SD approaches. (2) The performance of the SI and SA methods was not much different, but the SI method was somewhat better than the SA method for the female speaker case, especially when the modification ratio was high. (3) In all experiments, the ExcitNet-based system performed significantly better than the WaveNet-based one because the ExcitNet model was instructed to learn the variation of vocal cord movement.

4.3. Subjective test results

To evaluate the perceptual quality of the proposed system, mean opinion score (MOS) tests were performed[1]. In the tests[2], 12 native Korean listeners were asked to make quality judgments about the synthesized speech using the following five possible responses: 1 = Bad; 2 = Poor; 3 = Fair; 4 = Good; and 5 = Excellent. In total, 20 utterances were randomly selected from the test set and were then synthesized using the different neural vocoders. To verify vocoding performance, speech samples synthesized by conventional vocoders such as ITFTE and WORLD (D4C edition [26]) were also included.

---

[1]Generated audio samples are available at the following url: `https://sewplay.github.io/demos/vocoder_adaptation/`

[2]The tests were performed in an acoustically isolated room using a Sennheiser HD650 headphone.

As presented in Figure 4, the subjective test results confirm the effectiveness of each system in several ways: (1) In both the analysis and synthesis (A/S) and SPSS frameworks, the SD vocoders performed worst because it was difficult to learn the target speaker's characteristics with such a small amount of training data. (2) As the SI models could represent multiple speaker's voices, they were able to synthesize more natural speech than than the SD approaches. (3) Across all the training methods, the SA version achieved the best quality, which confirms that adapting the multi-speaker model to the target speaker's database is beneficial for the vocoding performance. (4) Comparing with the WaveNet, the ExcitNet performed much better overall, confirming that decoupling the formant component of the speech signal via an LP inverse filter significantly improves the modeling accuracy. (5) Consequently, the SPSS system with the proposed SA-ExcitNet vocoder achieved 3.80 and 3.77 MOS for the Korean male and Korean female speakers, respectively.

5. Conclusion

This paper proposed speaker-adaptive neural vocoders for statistical parametric speech synthesis systems (SPSS) when the amount of target speaker's data is insufficient. Using an initial speaker-independent trained model, the system first captured universal attributes from the waveform of multiple speakers'. This model was then fine-tuned with the target speaker's database to successfully represent speaker-specific characteristics using only ten minutes of training data. Adapting an ExcitNet framework with spectral filters also helped to improve the modeling accuracy. The experimental results verified that an SPSS system with the proposed speaker-adaptive neural vocoder performed significantly better than traditional versions with linear predictive coding-based vocoders and systems with similarly configured neural vocoders trained both speaker-dependently and speaker-independently. Future research includes integrating the entire framework into speech synthesis systems that use an end-to-end approach.

6. References

[1] A. Van Den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves *et al.*, "Conditional image generation with PixelCNN decoders," in *Proc. NIPS*, 2016, pp. 4790–4798.

[2] 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," *CoRR abs/1609.03499*, 2016.

[3] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent WaveNet vocoder," in *Proc. INTER-SPEECH*, 2017, pp. 1118–1122.

[4] 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.

[5] Y.-J. Hu, C. Ding, L.-J. Liu, Z.-H. Ling, and L.-R. Dai, "The USTC system for blizzard challenge 2017," in *Proc. Blizzard Challenge Workshop*, 2017.

[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 *Proc. ICASSP*, 2018, pp. 4779–4783.

[7] N. Adiga, V. Tsiaras, and Y. Stylianou, "On the use of WaveNet as a statistical vocoder," in *Proc. ICASSP*, 2018, pp. 5674–5678.

[8] E. Song, K. Byun, and H.-G. Kang, "Excitnet vocoder: A neural excitation model for parametric speech synthesis systems," in *Proc. EUSIPCO (in press)*, 2019.

[9] T. Yoshimura, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Mel-cepstrum-based quantization noise shaping applied to neural-network-based speech waveform synthesis," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 26, no. 7, pp. 1173–1180, 2018.

[10] L. Juvela, V. Tsiaras, B. Bollepalli, M. Airaksinen, J. Yamagishi, and P. Alku, "Speaker-independent raw waveform model for glottal excitation," 2018, pp. 2012–2016.

[11] J.-M. Valin and J. Skoglund, "LPCnet: Improving neural speech synthesis through linear prediction," *Proc. ICASSP (in press)*, 2019.

[12] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter-based waveform model for statistical parametric speech synthesis," *Proc. ICASSP (in press)*, 2019.

[13] K. Tachibana, T. Toda, Y. Shiga, and H. Kawai, "An investigation of noise shaping with perceptual weighting for WaveNet-based speech generation," in *Proc. ICASSP*, 2018, pp. 5664–5668.

[14] B. Sisman, M. Zhang, and H. Li, "A voice conversion framework with tandem feature sparse representation and speaker-adapted WaveNet vocoder," *Proc. INTERSPEECH*, pp. 1978–1982, 2018.

[15] L.-J. Liu, Z.-H. Ling, Y. Jiang, M. Zhou, and L.-R. Dai, "WaveNet vocoder with limited training data for voice conversion," *Proc. INTERSPEECH*, pp. 1983–1987, 2018.

[16] X. Wang, J. Lorenzo-Trueba, S. Takaki, L. Juvela, and J. Yamagishi, "A comparison of recent waveform generation and acoustic modeling methods for neural-network-based speech synthesis," in *Proc. ICASSP*, 2018, pp. 4804–4808.

[17] E. Song, F. K. Soong, and H.-G. Kang, "Effective spectral and excitation modeling techniques for LSTM-RNN-based speech synthesis systems," *IEEE/ACM Trans. Audio, Speech, and Lang. Process.*, vol. 25, no. 11, pp. 2152–2161, 2017.

[18] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proc. AISTATS*, 2010, pp. 249–256.

[19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: http://arxiv.org/abs/1412.6980

[20] Y. Fan, Y. Qian, F. K. Soong, and L. He, "Multi-speaker modeling and speaker adaptation for DNN-based TTS synthesis," in *Proc. ICASSP*, 2015, pp. 4475–4479.

[21] S. Pascual and A. Bonafonte, "Multi-output RNN-LSTM for multiple speaker speech synthesis and adaptation," in *Proc. EUSIPCO*, 2016, pp. 2325–2329.

[22] S. Furui, "Speaker-independent isolated word recognition using dynamic features of speech spectrum," *IEEE Trans. Acoust., Speech Signal Process.*, vol. 34, no. 1, pp. 52–59, 1986.

[23] R. J. Williams and J. Peng, "An efficient gradient-based algorithm for on-line training of recurrent network trajectories," *Neural computat.*, vol. 2, no. 4, pp. 490–501, 1990.

[24] H. Zen, A. Senior, and M. Schuster, "Statistical parametric speech synthesis using deep neural networks," in *Proc. ICASSP*, 2013, pp. 7962–7966.

[25] K. Tokuda, T. Yoshimura, T. Masuko, T. Kobayashi, and T. Kitamura, "Speech parameter generation algorithms for HMM-based speech synthesis," in *Proc. ICASSP*, 2000, pp. 1315–1318.

[26] M. Morise, "D4C, a band-aperiodicity estimator for high-quality speech synthesis," *Speech commun.*, vol. 84, pp. 57–65, 2016.