



Speaker-independent neural formant synthesis

Pablo Pérez Zarazaga¹, Zofia Malisz¹, Gustav Eje Henter¹, Lauri Juvela²

¹Division of Speech, Music and Hearing, KTH Royal Institute of Technology, Sweden

²Department of Information and Communications Engineering, Aalto University, Finland

{pablopez, malisz, ghe}@kth.se, lauri.juvela@aalto.fi

Abstract

We describe speaker-independent speech synthesis driven by a small set of phonetically meaningful speech parameters such as formant frequencies. The intention is to leverage deep-learning advances to provide a highly realistic signal generator that includes control affordances required for stimulus creation in the speech sciences. Our approach turns input speech parameters into predicted mel-spectrograms, which are rendered into waveforms by a pre-trained neural vocoder. Experiments with WaveNet and HiFi-GAN confirm that the method achieves our goals of accurate control over speech parameters combined with high perceptual audio quality. We also find that the small set of phonetically relevant speech parameters we use is sufficient to allow for speaker-independent synthesis (a.k.a. universal vocoding).

Index Terms: speech synthesis, formant synthesis, neural vocoding

1. Introduction

The quality of synthetic speech audio has advanced dramatically in the last decade, and can now rival recorded natural speech in terms of realism [1]. This development has largely come to pass by replacing methods based on signal processing with data-driven approaches such as neural vocoders [2, 3].

Not all applications of speech technology have benefited from these advances, however. A case in point is speech science, e.g., phonetics or psycholinguistics. In these disciplines, synthetic stimuli are used in e.g., experiments that disentangle perceptual effects of different acoustic speech properties. To build the stimuli, speech parameters such as pitch and formant frequencies, need to be carefully controlled [4, 5]. Unfortunately, contemporary signal generators based on deep learning do not offer the ability to accurately control such low-level signal aspects. As a result, phonetics and other speech sciences remain reliant on legacy tools such as formant synthesis [6, 7] or overlap-add techniques [8, 9] for stimulus creation. However, speech from such systems is not convincingly realistic: it lacks many of the cues present in natural speech and is perceived differently by humans as evidenced in slower cognitive processing times [10, 4]. This casts some doubts over research findings based on stimuli generated by such legacy tools [10].

This paper aims to show a way in which advances in neural vocoding can benefit speech scientists by providing

speech signals that are both realistic and controllable. Concretely, we describe a system that predicts acoustic features (mel-spectrograms) for waveform synthesis from phonetically meaningful speech parameters. We use the term *neural formant synthesis* for the idea of mapping of speech parameters to waveforms using deep learning. Further goals are: 1) speaker-independence, to synthesise speech in any voice, and 2) use of pre-trained systems, to make the approach easy to adopt. We validate the proposed approach in objective experiments on control accuracy, as well as through subjective listening tests comparing to conventional neural vocoders¹. These experiments show that speaker-independent neural formant synthesis leveraging HiFi-GAN reaches better subjective quality and more accurate speech-parameter reproduction than a WaveNet vocoder driven by conventional mel-spectrograms.

2. Related work

2.1. Neural vocoders

Beginning with WaveNet [11, 2], recent years have seen an explosion of *neural vocoders* - speech signal generators based on neural networks. These have substantially advanced the realism of synthetic speech signals. Virtually every generative deep-learning paradigm has been investigated for neural vocoding: autoregressive density models using dilated CNNs [11] or RNNs [12, 13, 14], GANs [15, 16, 3], normalising flows [17, 18, 19], and denoising diffusion models [20, 21]. Although neural vocoders in principle can be trained to generate waveforms from any audio representation, the most common choice is to generate output from mel-spectrogram acoustic features, e.g., as in [22, 17, 3]. These features have proven useful for text-to-speech systems, but they do not provide an intuitive way to control important phonetic properties of speech. There are a few neural vocoders, e.g., [15, 14, 16, 19], that leverage the source-filter model of speech production [23]. These offer a better starting point for pitch manipulation, but not necessarily other speech properties such as formant frequencies.

2.2. Speech parameter manipulation

As existing neural vocoders provide little or no ability to control meaningful properties of the speech signal, they have remained largely unused in speech and language sciences, which still depend on older approaches instead. The tools used relate to the speech aspect that needs to be manipulated in a given linguistic experiment. Pitch and timing manipulation is probably implemented best e.g., using tools based on pitch-synchronous overlap-add (PSOLA) and its extensions [8, 9],

This work was supported by the Swedish Research Council grant no. 2017-02861 “Multimodal encoding of prosodic prominence in voiced and whispered speech” and partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. We acknowledge the computational resources provided by the Aalto Science-IT project.

¹Audio examples and links to the code and models can be found here: <https://perezpoz.github.io/neuralformants>




Figure 1: Block diagram of our neural formant synthesis approach. At analysis stage, we estimate a compact set of speech parameters, consisting most notably of the fundamental and formant frequencies. The synthesis stage consists of mapping these compact parameters into mel spectrograms that in turn drive a generic neural vocoder for waveform synthesis.

or via signal-processing-based vocoders such as STRAIGHT [24]. Manipulating other speech properties, such as formant frequencies, is more involved and difficult to achieve without degraded speech quality. A widely used tool for formant manipulation is Praat [25], which combines linear predictive analysis with root-finding to translate the LPC polynomial into formant frequencies and bandwidths. While minor formant modification still allows the use of the original excitation source signal, more drastic modification can create artefacts due to imperfect source-filter separation, resulting in clashing residue formant information in the source signal. For larger modification, Praat can create a resonance-free source signal using the classic impulse train method, but this simplified approach degrades the quality of the synthesis. For optimal speech signal control, one might use a conventional formant synthesiser to generate synthetic speech waveforms with formant trajectories etc. completely determined by the experimenter. However, even though formant synthesis is capable of high signal quality [26], the results achieved when using these systems in practice leave a lot to be desired in terms of speech realism [10, 1].

The most similar prior work is probably Wavebender GAN [27], a model able to generate high-quality synthetic speech from a small set of phonetically relevant core speech parameters using a pre-trained neural vocoder. We use a closely related set of speech parameters, but whereas Wavebender GAN is single speaker and requires several networks to be trained, we present a simpler architecture. We show that it successfully generates speech also with a variety of unseen voices. We also study the effect of using different neural vocoders with the approach.

3. Method

The goal of this paper is to generate speech signals from phonetically meaningful speech parameters, as in classic formant synthesis, using neural waveform generators for better speech quality. To do this, we use multispeaker data to train a small network that converts speech-parameter trajectories as input into mel-spectrograms, which then drive a pre-trained neural vocoder. See Fig. 1 for an overview. In this section, we present the speech parameters to be controlled and the approach to speaker independence, as well as detail the network used to predict mel-spectrogram acoustic features from speech parameters.

3.1. Speech parameter extraction

For the controllable neural formant synthesis in this paper, we include a core set of phonetically meaningful speech parameters. Specifically, we decided to stay close to the parameters selected for the demonstration system in Wavebender GAN [27]. The speech parameters we include are:

- A voice flag: A binary flag indicating if the current frame is

voiced or unvoiced (VUV).

- Log F_0 : log-scale fundamental frequency of the signal.
- F_1 – F_4 : Values of the first four formants as resonant frequencies of the vocal tract on the linear scale.
- Spectral tilt: Rate of decay of the spectrum’s energy towards higher frequencies, measured as the slope of the linear regression of all the points in the magnitude spectrum of a frame.
- Spectral centroid: This represents the “centre of gravity” of the linear magnitude spectrum, computed as an average of the frequencies of all frequency-bin centre frequencies, weighted by the fraction of the total spectral energy in each bin.
- Energy: Signal energy of the frame in linear scale (not used in Wavebender GAN).

The parameters are extracted for every analysed frame, using an analysis window of 0.0464 s and a hop size of 0.0116 s, which correspond to 1024 and 256 samples at the 22,050 Hz sampling rate. The parameters make a total of one binary flag and eight continuous-valued features per each analysis frame. Each parameter trajectory except the binary flag was mean and variance normalised prior to applying deep learning.

F_0 is extracted using the algorithm implemented in Praat with a 10 ms time step, a range of 75–300 Hz for male speakers and 100–500 Hz for females. Formant values are extracted using the Burg method, with 25 ms window size, and pre-emphasis over 50 Hz. The maximum number of formants is set to 5 with the formant ceiling at 5000 and 5500 Hz for males and females respectively. The ceilings for F_1 and F_2 are then optimised using Escudero’s method [28]. Log F_0 -values were linearly interpolated in unvoiced regions. Similar interpolation was applied for missing formant values, which may happen when the formant extractor only identifies three or fewer resonant frequencies below the Nyquist frequency.

3.2. Speaker independence

Conventional neural vocoders are “universal”, i.e., the mel-spectrogram features they use as input convey enough information about the speaker’s voice to achieve good speaker similarity [29]. It is important for speech scientists to be able to generate speech stimuli in any given voice, without having to access a large speaker-specific speech database. Wavebender GAN, as a speaker-specific system requiring large amounts of audio data from the target speaker, was unable to generalise to unseen voices in preliminary experiments. We will show that our proposed method (trained on many speakers) is able to generalise to numerous unseen voices, despite using only the small selection of speech parameters listed in the previous section.




Figure 2: Log-scale plot of copy synthesis results (MSE of z -scored speech features). HiFi-GAN vocoding, neural formant synthesis with HiFi-GAN, and WaveNet vocoding are compared.

3.3. Feature-mapping neural network

The feature-mapping network takes in the frame-rate speech parameters and maps those to a sequence of mel-spectrogram acoustic features. We use a non-causal WaveNet-style gated convolutional architecture similar to [30, 31]. The network uses a kernel width of 3, and 1024 channels for residual, skip and post-net, with a total of 6 residual blocks set to follow the dilation pattern (1, 2, 4, 1, 2, 4). We train the network on the VCTK training set [32] for 99k updates using the Adam optimiser with learning rate $1e-4$ on batches of 128 sequences 46 frames long, and using mean squared error (MSE) as the loss function.

3.4. Vocoders

Once acoustic features have been predicted by the feature-mapping network, these are passed to a neural vocoder for audio waveform generation. We compare two different vocoders for this, namely WaveNet [11, 2] (autoregressive) and HiFi-GAN [3] (non-autoregressive, GAN-based). Both of these represent stronger, more realistic-sounding vocoders than the Griffin-Lim system considered in [1]. Therein, the Griffin-Lim system was found not to lead to slower cognitive processing times than in natural speech. It can, therefore, be argued that even higher synthesis quality should result in similarly low cognitive load.

The two vocoders operate at 22,050 Hz and were trained to synthesise waveforms from the same standardised acoustic representations, namely those extracted by the Nvidia implementation of Tacotron 2 [22]. While HiFi-GAN used the pre-trained V1 universal model, we trained a 30-layer WaveNet on the training portion of the dataset. The model uses 64 residual and skip channels, a kernel width of 3, and the dilation pattern grows in powers of two, reset every 10 layers. The model was trained for 1M updates with an autoregressive Gaussian density model [33]. We used the Adam optimiser at a fixed learning rate of $1e-3$ and batches of 8 one-second audio segments.

4. Evaluation

In order to assess both the accuracy of speech parameter control as well as the quality (realism) produced by our approach, we evaluate neural formant synthesis with two different neural vocoder back-ends. We conduct copy synthesis experiments as well as tests of manipulated speech. The following sections describe the experiments and their results in detail.

The presented models were trained and evaluated on data from the VCTK corpus [32], which contains 44,000 utterances uniformly divided among 110 English speakers with a variety of accents. The dataset was split into training, validation, and test sets with a ratio of 80-10-10% of the different speakers in




Figure 3: Boxplot of listening test scores of neural vocoders and neural formant synthesis using the same two vocoders.

each respective subset. We resampled the speech to 22,050 Hz and used a voice activity detector to trim extraneous silences in training and validation utterances.

In the evaluation, we compare our proposed method, i.e., feature mapping followed by a neural vocoder, to versions of WaveNet [11, 2] and HiFi-GAN [3] used as regular neural vocoders. This allows to differentiate between effects due to the neural vocoder and effects due to the feature-mapping network. Wavebender GAN, a speaker-specific system, was not included in the comparison, since we perform our evaluation on multiple speakers, none of which has enough data to feasibly train a Wavebender GAN. However, for experiments on formant and F_0 manipulation, we compare to a source-filter pipeline using Praat [25] to perform the same manipulation.

4.1. Copy synthesis evaluation

We begin by studying the accuracy of the neural formant synthesis (henceforth NF) in reconstructing speech signals using objective and subjective measurements. First, we perform copy synthesis with each system described in 4 using speech parameters (or, for neural vocoders, acoustic features) extracted from natural speech utterances as input. In order to compare the faithfulness of speech parameter recreation, we run speech parameter extractors on the copy synthesised audio and compare the resulting parameter trajectories to those extracted from the original speech. Specifically, we subtract reference values from extracted parameters in each frame, and compute the MSE across the entire test set. We perform this error computation on z -score normalised values for each parameter quantifying the relative error for a direct comparison across speech parameters.

The results of the copy synthesis experiments on utterances from held-out speakers are shown in Fig. 2. We see that HiFi-GAN alone best approximates natural speech parameters, compared to NF+HiFi-GAN formant synthesis. However, the formant synthesiser is not expected to, and indeed does not, surpass HiFi-GAN’s performance, with the exception of random fluctuations, as it relies on the vocoder. That said, the difference between HiFi-GAN and NF+HiFi-GAN is encouragingly small. Importantly, the NF+HiFi-GAN tends to produce more faithful speech parameter trajectories than the solo WaveNet.

The median mean squared error for most speech parameters is between 10^{-4} and 10^{-2} , which we consider a good result. $\log F_0$, F_1 , and spectral centroid achieve the lowest copy synthesis error, with errors increasing with each formant. This trend is consistent with the formant extractor missing a resonant peak in the spectral envelope on occasion, affecting estimated formant frequencies above that peak. F_4 and energy display the greatest error, despite the fact that energy is straightforwardly defined and easy to manipulate without machine learning.




Figure 4: Log-scale manipulation results (MSE of z -scored speech features). Each feature is scaled in the range $[0.7, \dots, 1.3]$ and re-synthesised using neural formant synthesis with HiFi-GAN, and then compared to using Praat for the same manipulation.

To assess the subjective quality of the generated speech audio, we perform a MUSHRA-like [34] test, asking crowd-sourced listeners to “Listen to the speech samples and rate the quality of the sound”. Technical limitations of the testing platform required a discrete grading scale, and in order to provide a higher resolution than a 1-5 MOS scale, we choose an integer scale from 0 (worst) through 10 (best) for their ratings. Stimuli for the same utterance were presented side-by-side, with reference audio (the natural recorded utterance) both available as an explicit and hidden reference among the stimuli to be rated. 25 participants, all self-reported native English speakers, were recruited using the Prolific platform and ran the test in a browser, wearing headphones. The test included 10 utterances selected from the 10 held-out speakers. Three attention checks were used during the test. We only included copy synthesis stimuli in this evaluation, since the parameter manipulations would give rise to improbable and less natural parameter trajectories even if the manipulation is perfectly faithful.

The results of the listening test are presented in Fig. 3. As expected, recorded speech achieves the highest rating, whilst neural formant stimuli each score slightly below their respective vocoder, considered an upper bound. A key finding is that the perceived quality of NF+HiFi-GAN exceeds the quality of WaveNet driven by more information-rich mel-spectrograms. WaveNet has established itself as the baseline high quality neural vocoder. Being able to surpass its performance with speaker-independent neural vocoding is an impressive result that indicates the strong potential our tools have for improved stimulus creation in speech and language sciences.

4.2. Parameter-manipulation evaluation

In the second experiment, we aim to assess the control accuracy of the present approach for manipulating the selected speech parameters listed in Sec. 3.1. We study the effect of manipulating parameter trajectories fed into neural formant synthesis. Specifically, manipulated speech was created by scaling each input parameter trajectory, multiplying it with a constant factor (0.7, 0.8, 0.9, 1.1, 1.2, and 1.3) across an entire utterance. For $\log F_0$, which already has been mapped to the logarithmic domain, we added or subtracted a constant value from the trajectory instead. The scaling was performed for each utterance and for each speech parameter in isolation, keeping all other parameter trajectories as they were during the copy synthesis. In this experiment, only the HiFi-GAN vocoder was considered.

After having generated synthetic speech, we again ran automated speech parameter extraction on the resulting samples, to check if the parameter trajectories in the manipulated speech accurately hit their target values. The lower and upper limits on formants and fundamental frequency for the feature extractors were adjusted as needed, so that our target would not fall

outside the range considered by these feature extractors.

The results of manipulating fundamental frequency and formant frequencies are presented in Fig. 4, where they are compared to the effect of using Praat for performing the exact same manipulation. We see that both methods have very similar accuracy when manipulating the fundamental frequency. When manipulating formants, the neural formant synthesiser generally gives more accurate results, according to our automated analysis. This is promising, but it is not possible to tell if these numbers are affected by the fact that errors typically made by the automated speech-parameter extractors have been seen by the neural approach during training. In general, errors are larger, the greater the magnitude of the manipulation, both for our neural formant synthesis and the Praat baseline.

For every manipulated stimulus, we extracted the trajectories of all speech parameters to verify that parameter control is properly de-correlated and that changing one parameter does not increase the errors in another. There is not enough space to graph these findings here, but from inspecting the output, results are mixed: neural formant synthesis exhibits larger errors in unmanipulated features than Praat when changing fundamental frequency and F_1 , but this situation is reversed for F_2 . In general, neural formant synthesis outperforms Praat in terms of spectral centroid accuracy, whereas Praat is better at maintaining the correct frame energy. To remedy that, neural formant synthesis could be made more accurate by changing the signal gain in a post-processing step. We invite the reader to listen to the audio samples and experiment on the source code ².

5. Conclusions and future work

We have demonstrated speaker-independent neural formant synthesis, recreating speakers outside the training set using two different vocoders. Our experiments confirm that important speech parameters can be accurately controlled, and that realism is close to state-of-the-art neural vocoding despite the minimalist set of speech parameters used to drive the synthesis.

In fact, neural formant synthesis on held-out speakers with the HiFi-GAN vocoder back-end was found to provide: a) a higher perceived speech quality as well as b) a more accurate copy synthesis than a conventional WaveNet vocoder driven by mel-spectrogram acoustic features. We think this strong result paves the way for improved synthetic and manipulated speech stimuli in speech and language sciences. Important future work includes integration of explicit formant control through differentiable digital signal processing (DDSP) and validation of the approach in phonetic research, e.g., by recreating classic categorical perception experiments.

²Audio examples and links to the code and models can be found here: <https://perezpoz.github.io/neuralformants>

6. References

- [1] Z. Malisz, G. E. Henter, C. Valentini-Botinhao, O. Watts, J. Beskow, and J. Gustafson, "Modern speech synthesis for phonetic sciences: a discussion and an evaluation," in *Proc. ICPhS*, 2019, pp. 487–491.
- [2] A. Tamamori, T. Hayashi, K. Kobayashi, K. Takeda, and T. Toda, "Speaker-dependent WaveNet vocoder," in *Proc. Interspeech*, 2017, pp. 1118–1122.
- [3] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *Proc. NeurIPS*, 2020, pp. 17 022–17 033.
- [4] A. J. Van Hessen and M. E. H. Schouten, "Categorical perception as a function of stimulus quality," *Phonetica*, vol. 56, no. 1-2, pp. 56–72, 1999.
- [5] L. Lisker and A. S. Abramson, "The voicing dimension: Some experiments in comparative phonetics," in *Proc. ICPhS*, 1970, pp. 563–567.
- [6] D. H. Klatt and L. C. Klatt, "Analysis, synthesis, and perception of voice quality variations among female and male talkers," *Journal of the Acoustical Society of America*, vol. 87, no. 2, pp. 820–857, 1990.
- [7] R. Carlson, B. Granström, and S. Hunnicutt, "A multi-language text-to-speech module," in *Proc. ICASSP*, vol. 7, 1982, pp. 1604–1607.
- [8] F. Charpentier and M. Stella, "Diphone synthesis using an overlap-add technique for speech waveforms concatenation," in *Proc. ICASSP*, vol. 11, 1986, pp. 2015–2018.
- [9] T. Dutoit, V. Pagel, N. Pierret, F. Bataille, and O. Van der Vrecken, "The MBROLA project: Towards a set of high quality speech synthesizers free of use for non commercial purposes," in *Proc. ICSLP*, 1996, pp. 1393–1396.
- [10] S. J. Winters and D. B. Pisoni, "Perception and comprehension of synthetic speech," *Research on Spoken Language Processing Progress Report*, no. 26, pp. 95–138, 2004.
- [11] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and et al., "WaveNet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [12] 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 *Proc. ICLR*, 2017.
- [13] 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.
- [14] J.-M. Valin and J. Skoglund, "LPCNet: Improving neural speech synthesis through linear prediction," in *Proc. ICASSP*, 2019, pp. 5891–5895.
- [15] B. Bollepalli, L. Juvela, and P. Alku, "Generative adversarial network-based glottal waveform model for statistical parametric speech synthesis," in *Proc. Interspeech*, 2017, pp. 3394–3398.
- [16] L. Juvela, B. Bollepalli, J. Yamagishi, and P. Alku, "GELP: GAN-excited linear prediction for speech synthesis from mel-spectrogram," in *Proc. Interspeech*, 2019, pp. 694–698.
- [17] R. Prenger, R. Valle, and B. Catanzaro, "WaveGlow: A flow-based generative network for speech synthesis," in *Proc. ICASSP*, 2019, pp. 3617–3621.
- [18] S. Kim, S.-G. Lee, J. Song, J. Kim, and S. Yoon, "FloWaveNet: A generative flow for raw audio," in *Proc. ICML*, 2019, pp. 3370–3378.
- [19] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter waveform models for statistical parametric speech synthesis," *IEEE/ACM T. Audio Speech*, vol. 28, pp. 402–415, 2019.
- [20] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "DiffWave: A versatile diffusion model for audio synthesis," in *Proc. ICLR*, 2021.
- [21] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "WaveGrad: Estimating gradients for waveform generation," in *Proc. ICLR*, 2021.
- [22] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, and et al., "Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions," in *Proc. ICASSP*, 2018, pp. 4799–4783.
- [23] G. Fant, *Acoustic Theory of Speech Production*. The Hague: Mouton, 1960.
- [24] H. Kawahara, "STRAIGHT, exploitation of the other aspect of VOCODER: Perceptually isomorphic decomposition of speech sounds," *Acoust. Sci. Technol.*, vol. 27, no. 6, pp. 349–353, 2006.
- [25] P. Boersma and D. Weenink, "Praat: doing phonetics by computer (computer software)," *Version 6.3.09*, retrieved 2 March 2023 from <http://www.praat.org/>, 2023.
- [26] M. J. Hunt and C. E. Harvenberg, "Generation of controlled speech stimulating by pitch-synchronous LPC analysis of natural utterances," in *Proc. ICA*, 1986, pp. A4–2.
- [27] G. T. D. Beck, U. Wennberg, Z. Malisz, and G. E. Henter, "Wavebender GAN: An architecture for phonetically meaningful speech manipulation," in *Proc. ICASSP*, 2022, pp. 6187–6191.
- [28] P. Escudero, R. Boersma, A. S. Rauber, and R. A. Bion, "A cross-dialect acoustic description of vowels: Brazilian and European Portuguese," *JASA*, vol. 126, no. 3, pp. 1379–1393, 2009.
- [29] J. Lorenzo-Trueba, T. Drugman, J. Latorre, T. Merritt, B. Putrycz, R. Barra-Chicote, A. Moinet, and V. Aggarwal, "Towards achieving robust universal neural vocoding," *arXiv preprint arXiv:1811.06292*, 2018.
- [30] D. Rethage, J. Pons, and X. Serra, "A WaveNet for speech denoising," in *Proc. ICASSP*, 2018, pp. 5069–5073.
- [31] M. Airaksinen, L. Juvela, P. Alku, and O. Räsänen, "Data augmentation strategies for neural network F0 estimation," in *Proc. ICASSP*, 2019, pp. 6485–6489.
- [32] J. Yamagishi, C. Veaux, K. MacDonald *et al.*, "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)," *University of Edinburgh. The Centre for Speech Technology Research (CSTR)*, 2019.
- [33] W. Ping, K. Peng, and J. Chen, "ClariNet: Parallel wave generation in end-to-end text-to-speech," in *Proc. ICLR*, 2019.
- [34] International Telecommunication Union, Radiocommunication Sector, "Method for the subjective assessment of intermediate quality levels of coding systems," ITU Recommendation ITU-R BS.1534-3, 2015.