



Transferring Source Style in Non-Parallel Voice Conversion

Songxiang Liu^{1,*}, Yuwen Cao^{1,*}, Shiyin Kang², Na Hu²,
Xunying Liu¹, Dan Su², Dong Yu², Helen Meng¹

¹Human-Computer Communications Laboratory,
The Chinese University of Hong Kong,
²Tencent AI Lab, Tencent, Shenzhen

{sxliu, ywcao, xyliu, hmmeng}@se.cuhk.edu.hk, {shiyinkang, ninahu, dansu, dyu}@tencent.com

Abstract

Voice conversion (VC) techniques aim to modify speaker identity of an utterance while preserving the underlying linguistic information. Most VC approaches ignore modeling of the speaking style (e.g. emotion and emphasis), which may contain the factors intentionally added by the speaker and should be retained during conversion. This study proposes a sequence-to-sequence based non-parallel VC approach, which has the capability of transferring the speaking style from the source speech to the converted speech by explicitly modeling. Objective evaluation and subjective listening tests show superiority of the proposed VC approach in terms of speech naturalness and speaker similarity of the converted speech. Experiments are also conducted to show the source-style transferability of the proposed approach.

Index Terms: voice conversion, style transfer

1. Introduction

Human speech conveys a wide range of information, among which the linguistic information and speaker identity are the most important. Voice conversion (VC) aims to modify speech characteristics mainly targeting the speaker identity (i.e. voiceprint) of the source speech, while the linguistic information should be modified during the VC process, but other spectral characteristics related to spoken behavior (e.g. emotions, emphasis, etc.) should be retained. This paper will refer to such spoken behavior as speaking style.

Various VC approaches have been proposed and most of them ignore speaking style during conversion. One class of techniques focus only on speaker identity conversion, such as the VC approaches based on Gaussian-mixture model (GMM) [1, 2], bidirectional LSTM [3], phonetic-posteriorgram (PPG) [4, 5], variational auto-encoder (VAE) [6, 7], generative adversarial network (GAN) [8, 9, 10, 11, 12] and etc. These approaches model the source-target speech features frame-by-frame and have limited capability in conducting time-scale modification on the source speech. This may lead to degradation on conversion performance in terms of naturalness of the speech output and its speaker similarity. Another class of techniques uses sequence-to-sequence models and can convert prosodic features such as F0 contour and duration [13], resulting in output that sounds more natural and more similar to the target speaker. However, these approaches inevitably change the speaking style of the source speech during conversion. For example, the source speech may sound excited but the converted speech does not. We consider that voice conversion tech-




Figure 1: (a) Graphical model for speech generation process. (b) Graphical model for voice conversion process. X and X_* represent speech, S represents speaker identity, Z represents speaking style, Y represents linguistic information and R represents rhythm.

nologies should preserve such speaking styles and transfer them from the input to the output.

To this end, we propose a sequence-to-sequence based non-parallel VC approach, which has the capability of transferring speaking style from the source speech to the converted speech by explicitly modeling. In this study, we regard the speech generation process as shown in Figure 1(a), where speech X is generated while conditioning on linguistic information Y , speaker identity S , rhythm R and speaking style Z . Rhythm characterizes the duration distribution for realizations of different phonetic structures, which is closely related to the speaking style. We explicitly model rhythm by incorporating a rhythm module into the speech generation model. The conversion process is shown in Figure 1(b). To generate the converted speech X_{VC} , the four components (Y , S , Z and R) are required. S represents speaker identity of a desired target speaker. Linguistic information Y is recognized from the source speech X_{src} , while the speaking style Z is inferred from a reference speech X_{ref} . When $X_{ref} = X_{src}$, we expect to transfer the speaking style from the source to the converted speech.

The current paper presents an approach which is among the first to explicitly model speaking style throughout the conversion process. The rest of the paper is organized as follows: Section 2 will present the proposed approach. Section 3 will describe baseline approaches for comparison. Experiments and results are presented in Section 4 and Section 5 concludes this paper.

2. Proposed approach

The proposed VC system is decomposed into a linguistic content recognizer and a speech generator, which are introduced in the following subsections. Since the recognizer and the generator can be optimized independently, the proposed VC approach does not require parallel speech data between a source speaker

*Work done during internship at Tencent AI Lab.




Figure 2: Schematic diagram of speech generator training process.

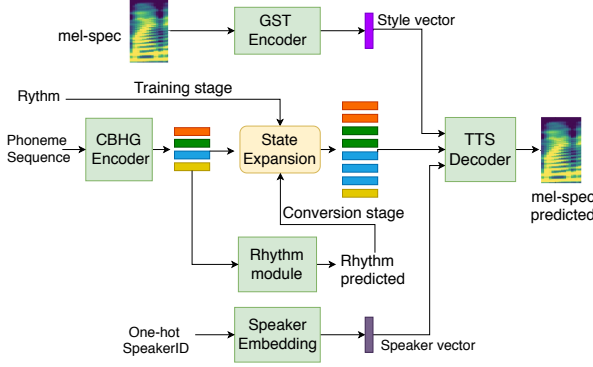


Figure 3: Model architecture of the speech generator.

and a target speaker.

2.1. Linguistic content recognizer

In this study, we use text transcriptions in supervised learning of a linguistic content recognizer. An attention-based sequence-to-sequence automatic speech recognition (ASR) model is used to recognize the phoneme sequence from the speech signal. To speed up training, we incorporate a connectionist-temporal-classification (CTC) [14] module into the ASR encoder model. The optimization objective during training is:

$$\log P(Y|X; \Theta_{rec}) = \lambda \log P_{CTC}(Y|X; \Theta_{rec}) + (1 - \lambda) \log P_{Att}(Y|X; \Theta_{rec}) \quad (1)$$

where P_{CTC} is the CTC objective, P_{Att} is the attention decoder objective and λ is a hyper-parameter weighting these two terms. Log-Mel-spectrograms are used to represent speech X . Θ_{rec} is model parameters.

2.2. Speech generator

We implement the speech generator based on a multi-speaker end-to-end text-to-speech (TTS) model. The training process is shown in Figure 2, where we use the phoneme sequence as the linguistic representation Y and log-Mel-spectrograms as speech features X . We use one-hot representation as the speaker identity input S and learn a speaker embedding table during training. Since the speaking style representation Z is hard to formulate, human annotation is difficult. Hence, for the purpose of our investigation, we infer speaking style from acoustic features X in an unsupervised way. While various techniques have been proposed to learn latent style representations from speech, we follow [15] and use a global style token (GST) encoder to compute style representation from the acoustic features. We incorporate a rhythm module to predict rhythm R into the speech generator, which functions similarly to the duration predictors

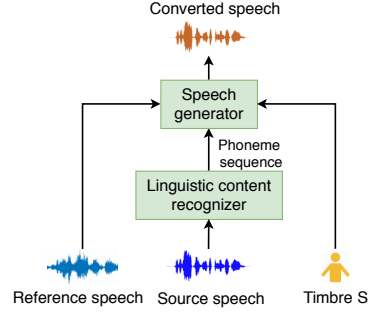


Figure 4: Schematic diagram of conversion process.

studied in [16, 17]. Auto-regressive (AR) TTS models such as Tacotron [18, 19] suffer from exposure bias [20], which leads to repeating/skipping words and early stopping phenomena in the generated speech. We discover that adding the rhythm module mitigates this issue. There are many ways to represent rhythm, e.g. number of phonemes uttered per second. Here we use an integer sequence for R , whose elements correspond to the duration of each phoneme in Y . Denoting the rhythm space as \mathcal{R} and speaking style space as \mathcal{Z} , during training we want to maximize:

$$P(X|Y, S; \Theta_{gen}) = \sum_{R \sim \mathcal{R}} \int_{Z \sim \mathcal{Z}} P(X, R, Z|Y, S; \Theta_{gen}) dZ \approx \sum_{R \sim \mathcal{R}} P(X|Y, S, R, Z^*; \Theta_{gen}) \cdot P(R|Y, S, Z^*; \Theta_{gen}) \quad (2)$$

where Θ_{gen} is the model parameter and Z^* is the computed style representation from X by the GST encoder, i.e., $Z^* = \text{Encoder}_{GST}(X)$. Since there are infinitely many integer sequences $R \sim \mathcal{R}$, the computation of the summation in Equation (2) is intractable. We resort to variational inference, where $P(R|Y, S, Z^*; \Theta_{gen})$ is approximated using a proposed distribution $q(R|X, Y)$, leading to the following evidence lower bound (ELBO):

$$\log P(X|Y, S; \Theta_{gen}) \geq \mathbb{E}_{R \sim q(R|X, Y)} [\log P(X|Y^*, S, Z^*; \Theta_{gen}) + \log P(R|Y, S, Z^*; \Theta_{gen})] \quad (3)$$

where Y^* is the expanded Y along the time axis according to R . The first term in the right-hand side of Equation (3) is related to the reconstruction loss of the mel-spectrograms while the second term is related to the rhythm prediction loss. While using an AR teacher model to sample R for ELBO optimization in Equation (3) is possible as shown in [17], we use an HMM-GMM-based forced-aligner to sample R . A LSTM-based AR decoder is adopted to model $\log P(X|Y^*, S, Z^*; \Theta)$ as

$$\log P(X|Y^*, S, Z^*; \Theta_{gen}) = \sum_{t=1}^{R_{sum}} \log P(X_t|X_{<t}, Y_{<t}^*, S, Z^*; \Theta_{gen}) \quad (4)$$

where R_{sum} is the summation of entries in R .

2.3. Implementation details

We adopt the transformer-based hybrid CTC-attention ASR model [21] for the linguistic content recognizer. The input

features are 80-dimensional log-mel-spectrograms, on which we conduct utterance-level mean-variance normalization before feeding into the ASR encoder. The Encoder first sub-samples the input features by 4 times in time-scale using two VGG-like max pooling layers. Then the hidden features go through 12 blocks of 8-head multi-head self-attention modules. The ASR decoder consists of 6 blocks of 8-head multi-head self-attention modules. λ in Equation (1) is set as 0.3.

The model architecture of the speech generator is as shown in Figure 3. The CBHG Encoder has the same structure as in [18], which takes phoneme sequence as input. The encoded phoneme sequence is expanded by repetitions along the time axis according to the duration sequence obtained from an HMM-GMM-based forced-aligner. The GST encoder employs the same network structure as in [15] and produces a 256-dimensional style embedding vector from a mel-spectrogram. One-hot vector is used for speaker identity representation and goes through a speaker embedding layer to get a 256-dimensional speaker vector. The TTS decoder consists of two 128-unit fully-connected (FC) layers, two 512-unit LSTM layers and a post-process module as in [19]. The expanded encoder output first goes through the FC layers and then concatenates with the style embedding and speaker vector at every time step. The concatenated features are then fed into the remaining layers of the TTS decoder to obtain the predicted mel-spectrogram. The rhythm module consists of three 512-unit bidirectional LSTM layers, which is trained to predict rhythm (i.e. duration sequence in this study). What is not shown in Figure 3 for simplicity is that the rhythm module also takes style embedding and speaker vector as additional input, which are concatenated with CBHG encoder output at every frame. We use L2 loss for both mel-spectrogram prediction and rhythm prediction, which correspond to terms 1 and 2 in the right-hand side of Equation (3).

2.4. Conversion process

The conversion process is shown in Figure 4. The linguistic content recognizer first predicts the phoneme sequence from the source speech. Speaking style representations are extracted from reference speech, which can be the source speech if we want to transfer the source style into the converted speech. Speaker identity S is a one-hot vector representing the desired target speaker. Conditioning on the obtained phoneme sequence, speaking style and speaker identity on the speech generator, we get the converted speech, where state expansion operation in Figure 3 uses predicted duration sequence from the rhythm module. In this study, we use a WaveRNN vocoder [22] to synthesize waveforms from log-mel-spectrograms.

3. Baseline approaches

We compare the proposed approach with two strong baseline models:

PPG-VC: This baseline model is similar to the N10 system [23] in VCC 2018 [24]. The approach consists of a conversion model and a neural vocoder. The conversion model maps PPGs into log Mel-spectrograms, which consists of four 256-unit bidirectional LSTM layers. We use a WaveRNN model as the neural vocoder.

Seq2seqVC: This baseline model is proposed in [13] and we use the released implementation by the authors in this study.

Table 1: ASR-based PER results.

Model	Sub(%)	Del(%)	Ins(%)	PER(%)
Ground-truth	1.5	0.1	0.1	1.7
PPG-VC	20.0	5.3	2.4	27.7
Seq2seqVC	9.6	9.5	12.6	31.7
Proposed	4.4	0.8	0.4	5.6

4. Experiments

4.1. Experimental setup

We use an internal multi-speaker Mandarin Chinese corpus for experimental evaluation. The corpus contains 42 speakers (26 female + 16 male). In total, there are 220 hours speech data and the average is 3.5 seconds per utterance. There are few parallel sentences between any two speakers. We split the corpus into training set (240k utterances), validation set (6457 utterances) and test set (6437 utterances). The sampling rate of the audio is 24 kHz. We first conduct experiments to show that the proposed method has superior VC performance than the two baseline approaches introduced in Section 3. Then, we show that the proposed approach has better source style transfer performance than the baseline approaches.

4.2. VC evaluation

In this part, we choose one female (with 37,482 training utterances) and one male (with 13,863 training utterances) speaker as the target speakers, and choose another female and male speaker as the source speakers. To evaluate the VC performance of the proposed approach, we use one target utterance uttered in neutral prosodic style as the reference speech in Figure 4 during conversion. Since VC tasks often assume that there is usually a limited amount of speech data available from the target speaker, we only use 50 randomly chosen utterances (about 3 minutes) from each target speaker.

During the training process of the proposed approach, we use speech data from all speakers except the two target speakers to train the linguistic content recognizer and speech generator. Then we adapt the speech generator for each target speaker using the 50 utterances from him/her. During the training process of the PPG-VC baseline approach, we first use speech data from all speakers except the two target speakers to pretrain the PPG-to-Mel-spectrogram conversion model, and then we adapt the conversion model for each target speaker using the 50 utterances from that speaker. The training process of the Seq2seqVC baseline approach is similar. We first train the model using speech data from all speakers except the two target speakers. Then, for a given pair of source-target speakers, we use all the training data from the source speaker and the 50 utterances from the target speaker to adapt the model. In this study, we update all model parameters during adaptation. The WaveRNN vocoder is trained using all the training data.

4.2.1. Objective measure

We use an off-line trained ASR model to measure the phone error rate (PER) of the converted speech. To show that the proposed approach has more robust conversion performance than the Seq2seqVC baseline, we choose the longest 200 utterances in terms of text lengths from the test set for each source speaker to compute the PER. The results are shown in Table 1, where Sub, Del, and Ins represent substitution, deletion, and insertion errors, respectively. We can see that the proposed approach

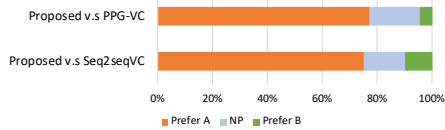


Figure 5: Speech naturalness AB listening test results.




Figure 6: Speaker similarity ABX listening test results.

achieves lower PER than both baseline models.

4.2.2. Listening tests

Two subjective evaluations are conducted: speech naturalness AB test and speaker similarity ABX test. In the AB test, paired speech samples (A and B) from the proposed approach and the baseline approaches are presented to listeners, who are asked to indicate the samples with better naturalness, or show no preference (NP). In the ABX test, X indicates the target reference sample. Paired speech samples (A and B) are presented and the listeners are asked to determine which one has closer speaker identity to the reference, or if they are equally close. Each conversion (cross-gender and intra-gender) has 20 samples for evaluation. 10 native Chinese speakers have participated in the evaluations and they are allowed to replay each sample pair as many times as necessary in both evaluations¹.

The subjective evaluation results are illustrated in Figure 5 and Figure 6. We can see that the proposed approach significantly outperforms the baseline approaches in terms of speech naturalness and speaker similarity of the converted speech.

4.3. Style transfer evaluation

In this part, we evaluate speaking style transfer performance from the source speech of the proposed approach. We choose a female as source speaker (with 7514 training utterances) and a female as target speaker. The training process is similar to that in Section 4.2, except that during adaptation, we use all the training data (6,725 utterances) from the target speaker. During conversion, we make the reference speech in Figure 4 the same as the source speech for source speaking style transfer.

4.3.1. Visualization

The F0 contour reflects prosodic variations in an utterance, which is related to speaking style. Continuously interpolated F0 contours of the converted speech of a source sample utterance are shown in Figure 7. We can see that the F0 contour of the converted speech by the proposed approach matches the source F0 contour more closely than the baseline approaches. Moreover, the converted speech by the proposed approach has greater durational similarity in the source speech than the Seq2seqVC approach. Therefore, we state that the proposed approach per-

¹Some audio samples in the listening tests can be found in “<https://liusongxiang.github.io/StyleTransferVC/>”.




Figure 7: Continuously interpolated F0 contour.




Figure 8: Source style transfer ABX listening test results.

forms better in preserving the speaking style of the source speech. The PPG-VC approach employs frame-to-frame framework for the conversion model. So the converted speech by the PPG-VC approach has exactly the same duration as the source speech. However, since the PPG-VC approach adopts a simple linear transformation in logarithmic scale on the source F0, the style transfer performance should be worse.

4.3.2. Listening test

We conduct an ABX test to compare the style transfer performance between the proposed approach and the baseline approaches. The reference X indicates the source speech. Converted samples (A and B) are presented and the listeners are asked to determine which one has closer speaking style to the reference. The listeners are asked to focus on speaking styles such as tone, stress, speaking speed, phrasing, pausing and etc. 20 testing cases were used. The ABX test results are shown in Figure 8. We can see that the proposed approach has significantly better style transfer performance than both baseline approaches.

5. Conclusions

In this study, we have presented a source-speaking-style transferable non-parallel VC approach. We incorporate a rhythm module into the sequence-to-sequence VC model, resulting in PER of 5.6% in the converted speech. Subjective listening tests have also shown its superiority in VC performance in terms of speech naturalness and speaker similarity of the converted speech. We adopt a GST encoder for style inference from the source/reference speech. Experimental results validate the source-style transferability of the proposed approach. This study is an initial attempt to model speaking style for VC. Future work will include disentangling speaking style, linguistic information and speaker identity from the source speech effectively in an unsupervised way.

6. Acknowledgements

This project is partially supported by a grant from the HKSAR Government’s Research Grants Council General Research Fund (Project no. 14208817).

7. References

- [1] Y. Stylianou, O. Cappé, and E. Moulines, “Continuous probabilistic transform for voice conversion,” *IEEE Transactions on speech and audio processing*, vol. 6, no. 2, pp. 131–142, 1998.
- [2] T. Toda, A. W. Black, and K. Tokuda, “Voice conversion based on maximum-likelihood estimation of spectral parameter trajectory,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2222–2235, 2007.
- [3] L. Sun, S. Kang, K. Li, and H. Meng, “Voice conversion using deep bidirectional long short-term memory based recurrent neural networks,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4869–4873.
- [4] L. Sun, K. Li, H. Wang, S. Kang, and H. Meng, “Phonetic posteriorgrams for many-to-one voice conversion without parallel data training,” in *2016 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2016, pp. 1–6.
- [5] S. Liu, J. Zhong, L. Sun, X. Wu, X. Liu, and H. Meng, “Voice conversion across arbitrary speakers based on a single target-speaker utterance,” in *Proc. Interspeech*, 2018, pp. 496–500.
- [6] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, “Voice conversion from non-parallel corpora using variational auto-encoder,” in *2016 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)*. IEEE, 2016, pp. 1–6.
- [7] —, “Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks,” in *Proc. Interspeech 2017*, 2017, pp. 3364–3368. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-63>
- [8] J. chieh Chou, C. chieh Yeh, H. yi Lee, and L. shan Lee, “Multi-target voice conversion without parallel data by adversarially learning disentangled audio representations,” in *Proc. Interspeech 2018*, 2018, pp. 501–505. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1830>
- [9] Y. Gao, R. Singh, and B. Raj, “Voice impersonation using generative adversarial networks,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 2506–2510.
- [10] T. Kaneko and H. Kameoka, “Parallel-data-free voice conversion using cycle-consistent adversarial networks,” *2018 26th European Signal Processing Conference (EUSIPCO)*, 2018.
- [11] F. Fang, J. Yamagishi, I. Echizen, and J. Lorenzo-Trueba, “High-quality nonparallel voice conversion based on cycle-consistent adversarial network,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5279–5283.
- [12] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, “Stargan-vc: Non-parallel many-to-many voice conversion using star generative adversarial networks,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 266–273.
- [13] J. Zhang, Z. Ling, and L.-R. Dai, “Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2019.
- [14] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [15] Y. Wang, D. Stanton, Y. Zhang, R. J. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *ICML*, 2018, pp. 5167–5176. [Online]. Available: <http://proceedings.mlr.press/v80/wang18h.html>
- [16] C. Yu, H. Lu, N. Hu, M. Yu, C. Weng, K. Xu, P. Liu, D. Tuo, S. Kang, G. Lei *et al.*, “Durian: Duration informed attention network for multimodal synthesis,” *arXiv preprint arXiv:1909.01700*, 2019.
- [17] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech: Fast, robust and controllable text to speech,” in *Advances in Neural Information Processing Systems*, 2019, pp. 3165–3174.
- [18] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, “Tacotron: Towards end-to-end speech synthesis,” *Proc. Interspeech 2017*, pp. 4006–4010, 2017.
- [19] 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 *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [20] M. Ranzato, S. Chopra, M. Auli, and W. Zaremba, “Sequence level training with recurrent neural networks,” *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [21] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang *et al.*, “A comparative study on transformer vs rnn in speech applications,” *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019.
- [22] 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 *International Conference on Machine Learning*, 2018, pp. 2410–2419.
- [23] 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.
- [24] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, “The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods,” in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 195–202. [Online]. Available: <http://dx.doi.org/10.21437/Odyssey.2018-28>