



A Hierarchical Speaker Representation Framework for One-shot Singing Voice Conversion

Xu Li, Shansong Liu, Ying Shan

ARC Lab, Tencent PCG

{nelsonxli, shansongliu, yingsshan}@tencent.com

Abstract

Typically, singing voice conversion (SVC) depends on an embedding vector, extracted from either a speaker lookup table (LUT) or a speaker recognition network (SRN), to model speaker identity. However, singing contains more expressive speaker characteristics than conversational speech. It is suspected that a single embedding vector may only capture averaged and coarse-grained speaker characteristics, which is insufficient for the SVC task. To this end, this work proposes a novel hierarchical speaker representation framework for SVC, which can capture fine-grained speaker characteristics at different granularity. It consists of an up-sampling stream and three down-sampling streams. The up-sampling stream transforms the linguistic features into audio samples, while one down-sampling stream of the three operates in the reverse direction. It is expected that the temporal statistics of each down-sampling block can represent speaker characteristics at different granularity, which will be engaged in the up-sampling blocks to enhance the speaker modeling. Experiment results verify that the proposed method outperforms both the LUT and SRN based SVC systems. Moreover, the proposed system supports the one-shot SVC with only a few seconds of reference audio.

Index Terms: one-shot singing voice conversion, hierarchical speaker representation

1. Introduction

Singing voice conversion (SVC) is a technology that transfers the voice of a singing signal to a voice of a target singer while maintaining the underlying content and melody. It has been paid more and more attention due to its potential applications in human-computer interaction, entertainment, etc.

Singing signals contain many components mixed together, such as content, pitch, loudness, and speaker identity. An SVC system should be able to disentangle these components, then reconstruct the converted singing with the target speaker information while keeping other components unchanged. Prior work has explored adequately in disentangling singing components by separately encoding each component. For the content component, Zhang et al. [1] encodes the lyrics into hidden features to represent the content, while other approaches directly extract the content information from the source singing signal, e.g. extracting linguistic features by a pre-trained automatic speech recognition (ASR) model [2–4], applying adversarially trained content encoders [5, 6] and utilizing the vector quantized variational autoencoder (VQVAE) [7, 8]. In terms of the pitch extraction, the log-F0 [2, 9] and sine-excitation [10, 11] are two common methods to capture the pitch variation in singing signals. Moreover, the A-weighting mechanism of the power spectrum [12] and the root mean square energy [1] are also adopted to measure the loudness for more precise SVC.

The speaker modeling in singing signals is non-trivial. For a single singer conversion system, a.k.a. any-to-one conversion, it implicitly acquires the target speaker’s characteristics from the training data [2, 13]. For a multi-singer system, speaker modeling is necessary for distinguishing different singers and facilitating singing voice conversion. One of the commonly adopted methods in the multi-singer system training is the speaker lookup table (LUT) [5, 7, 12]. It assigns a fixed-dimensional vector to each singer in the training data. The assigned vectors are learnable so as to adaptively capture the corresponding speaker characteristics during training. The LUT-based approaches are typically leveraged in the any-to-many SVC scenario, however, it fails to perform SVC towards an unseen singer which is outside of the training data. Other approaches [1, 3, 11] leverage a speaker recognition network (SRN) to distill speaker information by extracting speaker embeddings from a target singer’s reference audio. This enables the one-shot SVC scenario where the target singer is not observed in the training data.

However, singing voice can be much more expressive than conversational speech. The speaker characteristics in singing can greatly fluctuate even within a few seconds. Probably the speaker information of a speech clip can, to some extent, be condensed into a compact speaker embedding. But it is highly questionable that a single speaker embedding, extracted from either LUT or SRN, is adequate to capture the dynamically varying speaker characteristics of a singing utterance. Thereby, this work proposes a novel SVC framework on exploring a hierarchical speaker representation to capture fine-grained speaker characteristics at different granularity. Specifically, it adopts an up-sampling stream and three down-sampling streams. The up-sampling stream contains several up-sampling blocks and transforms linguistic features into audio samples in an end-to-end manner. Reversely, one down-sampling stream of the three alternates down-sampling blocks with instance normalization (IN) [14] to transform audio samples back to linguistic features. The IN module after each down-sampling block performs sample-wise normalization over the time dimension and generates temporal statistics for each down-sampling block, which can be utilized to represent speaker characteristics. The temporal statistics of all the blocks form a hierarchical speaker representation at different granularity, which is fused into the up-sampling blocks to provide the speaker condition.

A concurrent work [15] on the text-to-speech task expresses a similar idea to ours, i.e. incorporating the IN module within a U-net structure to perform voice cloning. The major differences between our present study and [15] are: 1) [15] absorbs both speaker and speaking style information, e.g. pitch, from the temporal statistics, while this work extracts only speaker information from the temporal statistics. The pitch and loudness information are explicitly fed into the system as done in [12]. This can help disentangle speaker information from pitch and loudness; 2) the U-net structure in [15] transforms features be-

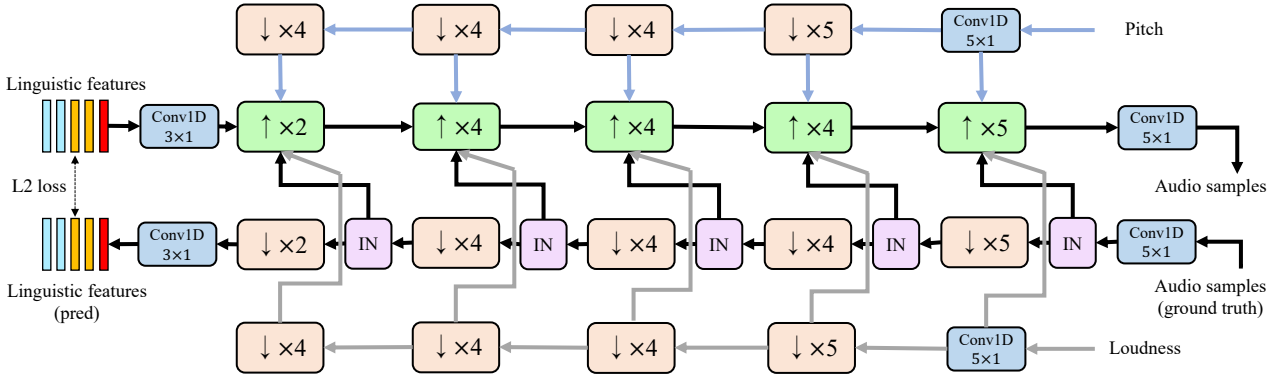


Figure 1: Schematic diagram of the proposed model architecture. In the figure, the “↑” and the “↓” represent the up-sampling and down-sampling blocks, respectively. The “IN” denotes instance normalization.

tween the content representations and the Mel-spectrograms. It utilizes a pre-trained MelGAN vocoder [16] for audio synthesis. While this work directly operates between content representations and audio samples in an end-to-end manner, which can alleviate the cumulative error in the multi-stage SVC [3]; 3) Both this work and [15] utilize a U-net-like architecture, but the block structures are different. [15] designs the blocks as the AdaIN layers [17] alternated with residual layers, while this work additionally adopts the feature-wise affine module [18] to fuse other conditions into the up-sampling stream.

This work has the following contributions: 1) Proposing a novel SVC framework that captures a hierarchical speaker representation at different granularity; 2) Demonstrating the superior performance of the proposed system over the LUT and speaker embedding based systems; 3) Showing our system supporting the one-shot conversion with only a few seconds of reference audio ¹.

The rest of this paper is organized as follows: Section 2 illustrates the proposed system architecture. Experimental setup and experiment results are demonstrated in Section 3 and 4, respectively. We conclude this work in Section 5.

2. Proposed Method

2.1. The proposed framework

The proposed SVC framework is illustrated in Fig. 1. It is modified from the architecture of FastSVC [12], which is an end-to-end SVC framework with high efficiency. The difference is that this work presents a hierarchical speaker modeling at different granularity to provide speaker information, rather than utilizing a speaker LUT as that in [12]. As shown in Fig. 1, the system consists of an up-sampling stream together with three down-sampling streams. The up-sampling stream transforms linguistic features block by block into the audio samples. The down-sampling stream with black arrows alternates down-sampling blocks with instance normalization (IN) [14] to transform audio samples back to linguistic features. The IN module has been shown promising in vision and speech tasks [14, 17], it performs sample-wise normalization over width/height/time dimensions to extract sample-wise characteristics. In this work, it performs sample-wise normalization over the time dimension to extract global temporal statistics, which can be utilized to represent speaker characteristics [17]. The temporal statistics of

all the blocks form a hierarchical speaker representation at different granularity, which is fused into the up-sampling blocks to provide the fine-grained speaker information. This work utilizes the mean as the temporal statistics for speaker representation, which will be explained in Section 2.3. Moreover, the other two down-sampling streams operate on the pitch and loudness features, respectively. The pitch and loudness features are explicitly extracted from the source singing voice, then up-sampled to match the rate of audio samples. After that, they are down-sampled accordingly to match the rate of each up-sampling block, and fused into the up-sampling blocks. Section 2.2 will detail the feature extraction process, while the structure of the building blocks will be demonstrated in Section 2.3. Finally, the training loss will be discussed in Section 2.4.

2.2. Feature extraction

Linguistic feature extraction: A hybrid CTC-attention ASR system [19] is adopted as the content extractor, which achieves state-of-the-art ASR performance. It consists of a Conformer [19] based encoder, a CTC decoder and an attention decoder, and adopts 80-dimensional log Mel-spectrograms as the input features. We follow the large version of the training configuration presented in Table 1 of [19] to train the system on the LibriSpeech corpus. Both decoders are discarded after training, and 512-dimensional vectors are extracted from the last output layer of the encoder as the linguistic features.

Pitch extraction: The Log-F0 [2, 9] and the sine-excitation [10, 11] are two common methods for pitch representation. Previous studies [4, 11] have shown that harmonic sine-excitation signals can improve the smoothness and continuity of the generated singing voice, which are thus adopted in this work. We extend the single sinusoid sine-excitation adopted in [12] to multiple harmonics [10] for more robust pitch representation.

Firstly, the frame-rate F0 values are extracted from the source singing audio by Parselmouth ², then up-sampled by linear interpolation to match the rate of audio samples, denoted by $f_{1:T}$. Finally, the harmonic sine-excitation values F_{it} are derived from $f_{1:T}$ as Eq.1:

$$F_{it} = \begin{cases} 0.1 \sin(\sum_{k=1}^t 2\pi \frac{i \times f_k}{f_s} + \phi_i) + n_{it}, & f_t > 0 \\ 100n_{it}, & f_t = 0 \end{cases} \quad (1)$$

where $i \in \{1, 2, \dots, K\}$, $t \in \{1, 2, \dots, T\}$ are the indexes of harmonics and time, respectively. F_{it} represents the excitation

¹Audio samples: <https://lixucuhk.github.io/Unet-SVC-Demo/>

²<https://github.com/YannickJadoul/Parselmouth>

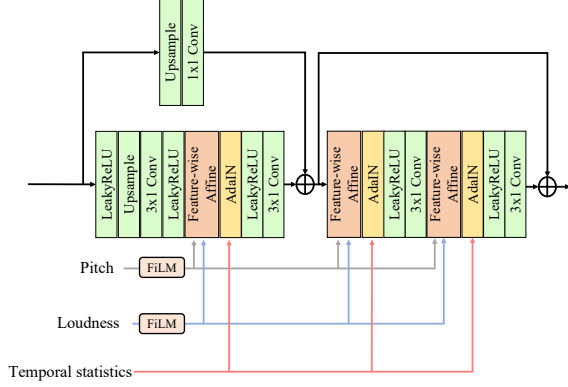


Figure 2: The structure of the up-sampling block, where the “FiLM” is a feature-wise linear modulation adopted in [12, 20].

value of the i -th harmonic at time t . This work sets $K = 8$ to use 8 harmonics. f_s is the audio sampling rate. ϕ_i is a random initial phase of the i -th harmonic. In unvoiced regions where $f_t = 0$, we set F_{it} as a Gaussian noise $n_{it} \sim \mathcal{N}(0, 0.003^2)$.

Loudness extraction: This work adopts the A-weighting mechanism of the power spectrum [21] as the loudness feature. It assigns greater emphasis to higher frequencies since human ears are less sensitive to low audio frequencies. The loudness features are computed with the configuration in [22], and the hop-size is set as 64. It is linearly interpolated to match the rate of audio samples before feeding into the up-sampling stream.

2.3. The structure of the building blocks

The building blocks are adapted from [12]. We make some modifications to the up-sampling block, while the down-sampling block remains unchanged. Fig. 2 illustrates the structure of the modified up-sampling block, where the feature-wise affine module [18] is used to fuse the pitch and the loudness features to the up-sampling block, and the AdaIN [23] module is fed with the temporal statistics.

In the FiLM module, given one condition feature, i.e. pitch or loudness, it generates a scale vector γ and a shift vector ξ , which are used to linearly transform the hidden features in the up-sampling block. Suppose we have the sine-excitation signals F , the loudness features L and the hidden features U , the feature-wise affine module will transform U into \tilde{U} by Eq. 2:

$$\tilde{U} = (\gamma_F + \gamma_L) \odot U + \xi_F + \xi_L \quad (2)$$

where γ_F and ξ_F are the affine parameters associated with F , γ_L and ξ_L are the ones associated with L , and \odot denotes the Hadamard product.

The AdaIN module first normalizes the hidden features by instance normalization, then utilizes the given temporal statistics to linearly transform the hidden features. This process removes the speaker characteristics of the up-sampling stream, then adapts the hidden features with the new speaker characteristics, i.e. the temporal statistics, generated from the down-sampling stream. In the conversion stage, the down-sampling stream operates on a target reference audio, then the temporal statistics containing the target speaker characteristics are fused to the up-sampling blocks to perform SVC. This work adopts all “IN” modules as the mean normalization and the variance statistics are not computed or used. Because in our experiments, we observed a quality degradation of the synthesized audio when

the mean-variance normalization is applied. One possible explanation is that the variance statistics could be largely affected by linguistic information, and it may be less helpful when transferring the variance from one audio to another that has mismatched linguistic contents, e.g. at the inference stage.

2.4. Training loss

The system is trained with a combination of three losses: the least-squares GAN [24] loss, a multi-scale STFT loss [25] and the reconstruction loss on linguistic features.

The discriminator adopts a multi-scale discriminator architecture [26]. Given the three sub-discriminators D_k ($k \in \{1, 2, 3\}$), the ground-truth audio x and the predicted audio \hat{x} , the adversarial loss of the generator can be expressed as:

$$\mathcal{L}_{adv} = \frac{1}{3} \sum_{k=1}^3 \|1 - D_k(\hat{x})\|_2 \quad (3)$$

The discriminator loss \mathcal{L}_D is formulated by Eq. 4:

$$\mathcal{L}_D = \frac{1}{3} \sum_{k=1}^3 (\|1 - D_k(x)\|_2 + \|D_k(\hat{x})\|_2) \quad (4)$$

The multi-scale STFT loss is computed as:

$$\mathcal{L}_{stft} = \frac{1}{|M|} \sum_{m \in M} \left(\frac{\|S_m - \hat{S}_m\|_2}{\|S_m\|_2} + \frac{\|\log S_m - \log \hat{S}_m\|_1}{N} \right) \quad (5)$$

where S_m and \hat{S}_m are the STFT magnitudes computed from x and \hat{x} respectively, with FFT sizes of $m \in M = \{2048, 1024, 512, 256, 128, 64\}$ and with 75% overlap. N is the total number of elements in S_m .

The reconstruction loss is the mean squared error between the ground-truth linguistic features c and the predicted ones \hat{c} :

$$\mathcal{L}_{mse} = \frac{1}{K} \sum_{i=1}^K (c_i - \hat{c}_i)^2 \quad (6)$$

where K is the number of elements in c . This constraint guides the down-sampling network to produce speaker-invariant \hat{c} .

Finally, the generator loss \mathcal{L}_G is formulated as:

$$\mathcal{L}_G = \mathcal{L}_{stft} + \alpha \mathcal{L}_{adv} + \beta \mathcal{L}_{mse} \quad (7)$$

where α and β are both experimentally set as 2.5 in this work.

3. Experimental Setup

The proposed model is evaluated in both in-set and out-set conditions. The in-set evaluation performs SVC among training speakers, while the out-set evaluation alters the voice of singing signals to the voice of out-set speakers which are unseen during training. The commonly used NUS-48E [27] corpus is adopted to train the model and perform the in-set evaluation. It contains 12 singers with around 15 minutes of data per singer. The dataset is randomly divided into train-validation-test sets according to a 90%-5%-5% split. For the out-set evaluation, we randomly select 2 male and 2 female singers from the NHSS [28] corpus, then use the models trained on the NUS-48E corpus to perform SVC to these unseen singers. All audios are down-sampled to 16kHz in our experiments.

The LUT and the speaker embedding-based SVC systems are adopted as two baselines. These baselines follow exactly

the same configurations as our proposed system, except for the speaker modeling part. The LUT-based system utilizes a speaker LUT, while the speaker embedding-based system leverages the state-of-the-art speaker recognition model, ECAPA-TDNN [29], to extract speaker embeddings. The extracted speaker representations are concatenated with the linguistic features before feeding into the up-sampling block. The mean squared error in Eq. 6 is dropped in the two baseline models.

In the building blocks, since the linguistic features have a hop-size of 640 samples, 5 up-sampling blocks gradually up-sample the time dimension of the linguistic features by factors of 2, 4, 4, 4, 5 with the number of channels of 288, 192, 96, 48, 24, respectively. The dilation rates are 1, 3, 9, 27 in all up-sampling blocks. The “Conv1D 3×1” in Fig. 1 has 384 channels with dilation rate of 1. The number of channels of the down-sampling blocks matches those in the corresponding up-sampling blocks. The dilation rates are 1, 2, 4 in all down-sampling blocks. The LeakyReLU activation function uses a negative slope of 0.2.

This work adopts the ADAM optimizer [30] to train all SVC systems. Each mini-batch contains 32 audio segments of a one-second length. The learning rate is initialized as 0.001 and halved for every 100k training steps. The discriminator joins the training process after 100k steps.

4. Experiment Results

4.1. Subjective evaluation

Subjective evaluation is conducted under both in-set and out-set conditions to verify the effectiveness of the proposed system. The standard 5-scale mean opinion score (MOS) test is adopted to evaluate the naturalness and the voice similarity. In the naturalness test, each group of stimuli contains recording samples that are randomly shuffled with the ones converted by the compared systems before being presented to raters. In the voice similarity test, converted samples are directly compared with the target singers’ reference recordings. At least 12 samples are rated for the compared systems in each testing scenario. We invite 15 Chinese speakers who are also proficient in English to participate in the MOS tests.

Table 1 illustrates the in-set evaluation results. The proposed system is observed to perform the best in terms of both naturalness and voice similarity, indicating the effectiveness of the proposed hierarchical speaker modeling with different granularity. The LUT-based approach outperforms the speaker embedding-based one. It is possibly because the speaker representations learned by the LUT approach remain constant for all the audios associated with the same training speaker, while the embeddings extracted from ECAPA-TDNN vary across different reference audios. The LUT approach gives stable voice conversion only for the training speakers, however, it fails to generalize to unseen speakers. Note that the similarity score variation is larger than that in the naturalness. This is possibly because the human perception of voice similarity can be largely influenced by pitch values.

The out-set evaluation results are shown in Table 2. Since the LUT approach does not support SVC for unseen speakers, we only compare the performance between the proposed system and the speaker embedding-based system. As shown in Table 2, the proposed approach still achieves superior performance to the other system in both naturalness and voice similarity, which verifies the effectiveness of the proposed approach.

Table 1: *The in-set evaluation: mean opinion score (MOS) results with 95% confidence intervals.*

	Naturalness	Similarity
LUT	3.43 ± 0.09	3.22 ± 0.14
ECAPA-TDNN	3.36 ± 0.10	3.17 ± 0.15
Proposed	3.53 ± 0.09	3.36 ± 0.12
Recordings	4.69 ± 0.07	4.39 ± 0.12

Table 2: *The out-set evaluation: mean opinion score (MOS) results with 95% confidence intervals.*

	Naturalness	Similarity
LUT	–	–
ECAPA-TDNN	3.24 ± 0.10	3.20 ± 0.16
Proposed	3.35 ± 0.11	3.29 ± 0.17
Recordings	4.77 ± 0.05	4.64 ± 0.12

4.2. Analysis

Fig. 3 shows the spectrograms of the singing voice converted by (a) the LUT-based system, (b) the speaker embedding-based system and (c) the proposed system. It shows that the spectrogram in (c) has a clearer texture than the ones in (a) and (b). Specifically, in (a) and (b), the parts in the green box are visibly blurred, which may be caused by the insufficient speaker information within the LUT and the speaker embedding-based systems. These blurred parts are improved to be clearer in (c). This demonstrates that the designed hierarchical speaker representation is capable of providing more details to reconstruct the singing voice, which thus boosts the SVC performance.




Figure 3: *The spectrograms of the converted singing voice by (a) the LUT based system, (b) the speaker embedding based system and (c) the proposed system. (Zoom in for better view.)*

5. Conclusions

This work proposes a novel SVC framework that extracts a hierarchical speaker representation to capture fine-grained speaker characteristics. It consists of an up-sampling stream and three down-sampling streams. The temporal statistics within one down-sampling stream form a hierarchical speaker representation at different granularity, which is fused into the up-sampling block to provide the speaker condition. Experiment results demonstrate the superior performance of the proposed system over both the LUT and the speaker embedding-based systems. Moreover, the proposed system supports one-shot SVC with only a few seconds of a target reference audio.

6. References

- [1] L. Zhang, C. Yu, H. Lu, C. Weng, C. Zhang, Y. Wu, X. Xie, Z. Li, and D. Yu, "Durian-sc: Duration informed attention network based singing voice conversion system," *INTERSPEECH*, 2020.
- [2] S. Liu, Y. Cao, D. Su, and H. Meng, "Diffsvc: A diffusion probabilistic model for singing voice conversion," *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021.
- [3] H. Guo, H. Lu, N. Hu, C. Zhang, S. Yang, L. Xie, D. Su, and D. Yu, "Phonetic posteriorgrams based many-to-many singing voice conversion via adversarial training," *arXiv preprint arXiv:2012.01837*, 2020.
- [4] A. Polyak, L. Wolf, Y. Adi, and Y. Taigman, "Unsupervised cross-domain singing voice conversion," *INTERSPEECH*, 2020.
- [5] C. Deng, C. Yu, H. Lu, C. Weng, and D. Yu, "Pitchnet: Unsupervised singing voice conversion with pitch adversarial network," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7749–7753.
- [6] E. Nachmani and L. Wolf, "Unsupervised singing voice conversion," *INTERSPEECH*, 2019.
- [7] N. Takahashi, M. K. Singh, and Y. Mitsufuji, "Hierarchical disentangled representation learning for singing voice conversion," in *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021, pp. 1–7.
- [8] D. Wang, L. Deng, Y. T. Yeung, X. Chen, X. Liu, and H. Meng, "Vqmvic: Vector quantization and mutual information-based unsupervised speech representation disentanglement for one-shot voice conversion," *INTERSPEECH*, 2021.
- [9] J. Lu, K. Zhou, B. Sisman, and H. Li, "Vaw-gan for singing voice conversion with non-parallel training data," in *2020 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2020, pp. 514–519.
- [10] F. Chen, R. Huang, C. Cui, Y. Ren, J. Liu, and Z. Zhao, "Singgan: Generative adversarial network for high-fidelity singing voice generation," *arXiv preprint arXiv:2110.07468*, 2021.
- [11] H. Guo, Z. Zhou, F. Meng, and K. Liu, "Improving adversarial waveform generation based singing voice conversion with harmonic signals," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [12] S. Liu, Y. Cao, N. Hu, D. Su, and H. Meng, "Fastsvc: Fast cross-domain singing voice conversion with feature-wise linear modulation," in *2021 IEEE International Conference on Multimedia and Expo (ICME)*. IEEE, 2021, pp. 1–6.
- [13] X. Chen, W. Chu, J. Guo, and N. Xu, "Singing voice conversion with non-parallel data," in *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE, 2019, pp. 292–296.
- [14] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," *arXiv preprint arXiv:1607.08022*, 2016.
- [15] R. Li, D. Pu, M. Huang, and B. Huang, "Unet-tts: Improving unseen speaker and style transfer in one-shot voice cloning," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [16] G. Yang, S. Yang, K. Liu, P. Fang, W. Chen, and L. Xie, "Multi-band melgan: Faster waveform generation for high-quality text-to-speech," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 492–498.
- [17] J.-c. Chou, C.-c. Yeh, and H.-y. Lee, "One-shot voice conversion by separating speaker and content representations with instance normalization," *INTERSPEECH*, 2019.
- [18] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan, "Wavegrad: Estimating gradients for waveform generation," *The International Conference on Learning Representations (ICLR)*, 2021.
- [19] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *INTERSPEECH*, 2020.
- [20] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "Film: Visual reasoning with a general conditioning layer," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [21] C. Meyer-Bischoff, "Measuring noise," *Medicine Sciences: M/S*, vol. 21, no. 5, pp. 546–550, 2005.
- [22] L. Hantrakul, J. H. Engel, A. Roberts, and C. Gu, "Fast and flexible neural audio synthesis," in *ISMIR*, 2019.
- [23] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4401–4410.
- [24] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, "Least squares generative adversarial networks," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2794–2802.
- [25] 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 *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6199–6203.
- [26] K. Kumar, R. Kumar, T. de Boissiere, L. Geste, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "Melgan: Generative adversarial networks for conditional waveform synthesis," *Advances in neural information processing systems*, vol. 32, 2019.
- [27] Z. Duan, H. Fang, B. Li, K. C. Sim, and Y. Wang, "The nus sung and spoken lyrics corpus: A quantitative comparison of singing and speech," in *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*. IEEE, 2013, pp. 1–9.
- [28] B. Sharma, X. Gao, K. Vijayan, X. Tian, and H. Li, "Nhss: A speech and singing parallel database," *Speech Communication*, vol. 133, pp. 9–22, 2021.
- [29] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," *INTERSPEECH*, 2020.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.