VAW-GAN for Singing Voice Conversion with Non-parallel Training Data

Junchen Lu *, Kun Zhou *, Berrak Sisman ^{†*} and Haizhou Li *

* Dept. of Electrical and Computer Engineering, National University of Singapore, Singapore

[†] Information Systems Technology and Design (ISTD) Pillar, Singapore University of Technology and Design, Singapore

Abstract—Singing voice conversion aims to convert singer's voice from source to target without changing singing content. Parallel training data is typically required for the training of singing voice conversion system, that is however not practical in real-life applications. Recent encoder-decoder structures, such as variational autoencoding Wasserstein generative adversarial network (VAW-GAN), provide an effective way to learn a mapping through non-parallel training data. In this paper, we propose a singing voice conversion framework that is based on VAW-GAN. We train an encoder to disentangle singer identity and singing prosody (F0 contour) from phonetic content. By conditioning on singer identity and F0, the decoder generates output spectral features with unseen target singer identity, and improves F0 rendering. Experimental results show that the proposed framework achieves better performance than the baseline frameworks.

Index Terms: singing voice conversion, VAW-GAN, F0 conditioning

I. INTRODUCTION

Singing voice conversion (SVC) is a voice conversion (VC) technique that converts source singer's voice to sound like target singer's voice, while preserving the singing content [1]. With singing voice conversion, we can make everyone sing like a professional, overcoming the limitation of physical constraints, controlling the voice timbre freely, and expressing the emotions in variable ways [2], [3].

Singing voice conversion shares many similarities with speech voice conversion [4]–[6]. They both aim to change the vocal identity. However, they are also different in many ways. For example, in speech voice conversion, speech prosody is considered to contain the information of speaker characteristics [7]–[10]. In contrast, in singing voice conversion, we assume that source singers are always singing on the key, which means the singing style is only determined by the sheet music, thus we consider singing style as a singer-independent feature. Therefore, only singer-dependent traits, such as vocal timbre, need to be converted [11]–[13].

Early studies attempted to convert singing voice through spectral modeling. Many statistical methods, such as Gaussian mixture model (GMM)-based many-to-many eigenvoice conversion (EV-GMM) [14], direct waveform modification based on spectrum differential (DIFFSVC) [15], and DIFFSVC with global variance [16] have been proposed for SVC. With the advent of deep learning, deep neural network (DNN) [17] and generative adversarial network (GAN) [13], [18] among others have shown improved quality and naturalness.



Fig. 1. A singing voice conversion system is trained on singing voice data recorded by the source and target singers. At run-time, the system takes singing voice of the source singer as input, and converts it to that of the target singer.

However, previous studies on singing voice conversion mostly require of parallel training data between the source and target singer. In practice, collecting such parallel data is expensive and time-consuming, that motivated non-parallel SVC methods, such as deep bidirectional long short term memory (DBLSTM) based recurrent neural network (RNN) [19], [20], Wasserstein generative adversarial network (WD-GAN) [21], and StarGAN [22] for SVC. Recently, encoderdecoder based networks [23], such as variational autoencoder (VAE) [24], variational autoencoding Wasserstein generative adversarial network (VAW-GAN) [25], auxiliary classifier variational autoencoder (ACVAE) [26], [27] and cycle-consistent autoencoder (CycleVAE) [28], [29] have been successfully applied into various tasks, such as cross-lingual voice conversion [30] and emotional voice conversion [31].

Autoencoder is effective for disentanglement of mixed information [24], [25], [32]. If we are able to disentangle vocal timbre of the singer, that we call singer identity in this paper, from phonetic content and singing prosody (F0 contour), we can simply replace the singer identity and F0 contour during signal reconstruction for singing voice conversion. We adopt VAW-GAN in this paper for two reasons. First, autoencoder doesn't require parallel training data during encoding and decoding; Second, the encoder-decoder architecture allows for effective control of singer identity and singing prosody, that makes many-to-many conversion easier than other nonparallel generative models, such as cycle-consistent generative adversarial network (CycleGAN) [6], [30], [33], [34].

It is known that F0 and spectral features are inherently correlated [35], [36]. In autoencoder-based VC, the latent code from encoder contains F0 information from the source, that adversely affects the output. Recent studies have also shown that disentangling F0 from latent code improves the performance of speech voice conversion [37], [38] and emotion

Codes & Singing Samples: https://kunzhou9646.github.io/singvaw-gan/

conversion [31], that motivates the studies in this paper. We are motivated to study disentanglement of both singer identity, and F0 information from singing content. This can be achieved by providing both singer identity and F0, in addition to latent code, to the decoder during training. In this way, we aim to obtain a latent code that is singer and F0 independent. At run-time, we add both singer identity and F0 as inputs to the decoder to control the signal reconstruction.

The main contributions of this paper include: 1) we propose a framework for singing voice conversion with non-parallel training data; 2) we achieve high quality converted singing voice; 3) we eliminate the need for parallel training data, time alignment procedures and other external modules, such as automatic speech recognition (ASR); 4) we show the effectiveness of the F0 conditioning mechanism for SVC. To our best knowledge, this is the first attempt to use F0 conditioning for non-parallel singing voice conversion.

This paper is organized as follows: In Section II, we recap the study of VAW-GAN for speech synthesis. In Section III, we introduce our proposed singing voice conversion framework. In Section IV, the experiments and results are reported. Conclusions are given in Section V.

II. RELATED WORK: VAW-GAN IN SPEECH SYNTHESIS

Recently, encoder-decoder networks such as variational Wasserstein generative adversarial network (VAW-GAN) [25] have drawn much attention because of their generating ability and controllability. VAW-GAN makes it possible to train a model without parallel data or any other time alignment procedures through an encoder-decoder structure.

The main idea of VAW-GAN is based on the probabilistic graphical model (PGM). Given spectral features x_s from source speaker and x_t from target speaker, the PGM tries to explain the observation x using two latent variables: the speaker representation vector y and the phonetic content vector z. It is noted that y is determined solely by the speaker identity and z is a speaker-independent variable. According to the PGM, the voice conversion function $f(\cdot)$ can be divided into two stages: 1) a speaker-independent encoder \mathcal{E}_{ϕ} with parameter set ϕ infers a latent vector from the source spectral features $z = \mathcal{E}_{\phi}(x_s)$, and 2) a speaker-dependent decoder \mathcal{G}_{θ} with parameter set θ reconstructs the input with the latent code z and a target speaker representation vector y_t . Therefore, the task of voice conversion is then reformulated as:

$$\boldsymbol{x}_t \approx f(\boldsymbol{x}_s) = \mathcal{G}_{\theta}(\mathcal{E}_{\phi}(\boldsymbol{x}_s), \boldsymbol{y}_t)$$
 (1)

During training, the frames that belong to the same phoneme class hinge on a similar z. With the latent content vector z, the decoder can generate voice of a specific speaker by varying the speaker representation vector y.

Different from variational encoding networks, generative adversarial network (GAN) produces sharper spectra since it optimizes a loss function between two distributions in a more direct fashion [25]. In order to achieve better conversion performance, VAW-GAN incorporates the discriminator from GAN models and assigns VAE's decoder as GAN's generator. In the case of voice conversion, the Jensen-Shannon divergence [39] in the GAN objective is renovated with a Wasserstein objective:

$$\mathcal{J}_{wgan} = \mathbb{E}_{\boldsymbol{x} \sim p_t^*} [\mathcal{D}_{\psi}(\boldsymbol{x})] - \mathbb{E}_{\boldsymbol{z} \sim q_{\phi}(\boldsymbol{z}|\boldsymbol{x})} [\mathcal{D}_{\psi}(\mathcal{G}_{\theta}(\boldsymbol{z}), \boldsymbol{y}_t)] \quad (2)$$

where p_t^* is the distribution of x_t , $q_{\phi}(z|x)$ is the inference model, and \mathcal{D}_{ψ} is the discriminator with parameter set ψ .

Therefore, the final objective loss function of VAW-GAN is given as follow:

$$\mathcal{J}_{vawgan} = -\mathcal{D}_{KL}(q_{\phi}(\boldsymbol{z}|\boldsymbol{x}) || p_{\theta}(\boldsymbol{z})) + \mathbb{E}_{\boldsymbol{z} \sim q_{\phi}(\boldsymbol{z}|\boldsymbol{x})}[\log p_{\theta}(\boldsymbol{x}|\boldsymbol{z}, \boldsymbol{y})] + \alpha \mathcal{J}_{wgan}$$
(3)

where α is a coefficient which emphasizes \mathcal{J}_{wgan} , \mathcal{D}_{KL} is the Kullback-Leibler divergence, $p_{\theta}(z)$ is the prior distribution model of z, and $p_{\theta}(x|z, y)$ is the synthesis model. Through the adversarial learning, the decoder minimizes the loss, while the discriminator maximizes it, until an optimal pseudo pair is found through this min-max game. This objective is shared across all three main components in VAW-GAN: the encoder, the decoder and the discriminator.

VAW-GAN has been successfully applied in the field of speech synthesis, such as voice conversion [25], [30] and emotion conversion [31]. We expect that the way it characterizes speaker identity also applies to singer identity. In this paper, we propose a VAW-GAN framework for singing voice conversion, which will be the focus of Section III.

III. VAW-GAN FOR SINGING VOICE CONVERSION

In this section, we propose the use of VAW-GAN for disentanglement of singer identity and F0 information from the phonetic content. The proposed VAW-GAN includes a singer-independent encoder, that generates latent code z, and a decoder that takes a triplet input, namely latent code, singer identity and F0.

A. Training Phase

The training phase is illustrated in Fig. 2. We first use WORLD vocoder [36] to extract spectral features (SP) and F0 from the singing waveform. The encoder takes input frames from multiple singers, and generate a singer-independent latent code z. We assume that the latent code z only contains the information of phonetic content.

We use a one-hot vector singer ID and source F0 as the input to the decoder, in addition to latent code. In this way, the encoder learns to disentangle singer ID and F0 from the latent code z after being exposed to singing data of multiple singers. By conditioning on singer ID and F0, the decoder, as formulated in Eq. (1), can be re-written as follows,

$$\boldsymbol{x} \approx f(\boldsymbol{x}, F0) = \mathcal{G}_{\theta}(\mathcal{E}_{\phi}(\boldsymbol{x}), \boldsymbol{y}, F0)$$
 (4)

The decoder learns to reconstruct the spectral features and the discriminator tries to distinguish whether the spectral features are from real singing voice or not. Through this min-max game, the encoder, decoder and discriminator are encouraged to find an optimal pseudo pairs during the training.



Fig. 2. The training phase of the proposed VAW-GAN (SID+F0) singing voice conversion framework. The encoder learns to disentangle singer identity and fundamental frequency (F0) from the phonetic content. Blue boxes are involved in the training.



Fig. 3. The run-time conversion phase of the proposed VAW-GAN (SID+F0) singing voice conversion framework. The decoder is conditioned on singer identity and fundamental frequency (F0) to generate spectral features for unseen target singer, and improve F0 rendering. Red boxes have been trained during the training phase.

B. Run-time Conversion

The conversion phase is illustrated in Fig. 3. We first extract spectral features and F0 from the source singing waveform using WORLD vocoder. The spectral features are then encoded into a latent code through the encoder and F0 is converted by logarithm Gaussian (LG)-based linear transformation [24].

At run-time conversion, the decoder is conditioned on the converted F0 features $\hat{F}0$. The converted spectral features \hat{x}_t are given as:

$$\hat{\boldsymbol{x}}_t = f(\boldsymbol{x}_s, \hat{F}0) = \mathcal{G}_{\theta}(\mathcal{E}_{\phi}(\boldsymbol{x}_s), \boldsymbol{y}_t, \hat{F}0)$$
(5)

where y_t is the designated singer ID.

The converted spectral features are then reconstructed by the decoder together with the converted F0 and the designated singer ID. Finally, we use WORLD vocoder to synthesis the converted singing waveform.

IV. EXPERIMENTS

We conduct both objective and subjective experiments to assess the performance of the proposed VAW-GAN for singing voice conversion. We use NUS Sung and Spoken Lyrics Corpus (NUS-48E corpus) [40], which consists of the sung and spoken lyrics of 48 English songs by 12 professional singers. We choose 2 male singers and 1 female singer for all the experiments. For each singer, 6 songs are used for training and evaluation.

We construct two systems: a) VAW-GAN with the decoder conditioning on singer ID (SID) and F0 (as illustrated in Figure 3) to convert the spectrum between different singers, denoted as VAW-GAN (SID+F0); b) VAW-GAN with the decoder conditioning only on singer ID, denoted as VAW-GAN (SID), that is similar to the VAW-GAN in [25] for speech voice conversion.

We use VAW-GAN (SID) as the reference baseline to show the effect of the proposed VAW-GAN (SID+F0), and report the performance in both objective and subjective evaluations. It is noted that both frameworks are trained with non-parallel singing voice data.

A. Experimental Setup

The singing voice data is down-sampled at 16kHz. We first use WORLD vocoder [36] to extract 513-dimensional spectral features (SP), F0, and aperiodicity (AP). The frame length is 25 ms with a frame shift of 5 ms. F0 is re-scaled to the range of [-1, 1]. The input SP of each frame is normalized to unit-sum. The normalization factor, known as the energy, is taken out as an independent feature. We use the log energy-normalized SP for VAW-GAN training.

The model architecture of our proposed VAW-GAN (SID+F0) framework is given in Table I. The encoder, the decoder, and the discriminator of both frameworks are all 1D convolutional neural networks (CNN), of which each layer is followed by a fully connected layer. The latent space is 128-dimensional and is assumed to have a standard normal distribution. The dimension of the speaker representation is set to be 10. For both frameworks, during the training phase, we first train the encoder-decoder pair for 15 epochs, and then train the whole VAW-GAN for 45 epochs. The framework is trained using RMSProp with a learning rate of 0.0001. During the conversion phase, SP and F0 are converted on a frame-by-frame basis, while AP and the energy remain unmodified.

B. Objective Evaluation

We use Mel-cepstral distortion (MCD) [2], [41] to measure the distortion between the converted and target Mel-cepstra, that is given as follows:

$$MCD[dB] = \frac{10}{T\ln 10} \sum_{i=1}^{T} \sqrt{2\sum_{d=1}^{D} (mc_{i,d}^{t} - mc_{i,d}^{c})^{2}}$$
(6)

THE MODEL ARCHITECTURE OF THE ENCODER, DECODER AND DISCRIMINATOR OF OUR PROPOSED FRAMEWORK VAW-GAN (SID+F0).

	# of Layers	Kernel Size	Stride	Output Channel
Encoder	5	$\{7, 7, 7, 7, 7\}$	$\{3, 3, 3, 3, 3, 3\}$	$\{16, 32, 64, 128, 256\}$
Decoder	4	{9, 7, 7, 1025}	$\{3, 3, 3, 1\}$	$\{32, 16, 8, 1\}$
Discriminator	3	{7, 7, 115}	{3, 3, 3}	{16, 32, 64}

TABLE II A COMPARISON OF THE MCD RESULTS BETWEEN VAW-GAN (SID+F0), VAW-GAN (SID) FOR MALE-TO-MALE AND MALE-TO-FEMALE SINGING VOICE CONVERSION.

Framework	MCD [dB]		
Flamework	$male \rightarrow male$	male \rightarrow female	
Zero effort	10.05	13.43	
VAW-GAN (SID)	7.20	7.39	
VAW-GAN (SID+F0)	5.51	6.57	

where $mc_{i,d}^c$ and $mc_{i,d}^t$ represent the d^{th} coefficient of the converted and target MCEPs sequences at the i^{th} frame, respectively. D is the dimension of MCEP features and T represents the total number of frames. In this paper, we extract 24-dimensional MCEPs at each frame, thus D is 24. We note that a lower value of MCD indicates a smaller spectrum distortion and a better conversion performance.

We report MCD results of our proposed framework VAW-GAN (SID+F0) and the baseline framework VAW-GAN (SID). Zero effort represents the cases where we directly compare the singing voice of source and target singers without any conversion. As reported in Table II, we observe that our proposed framework VAW-GAN (SID+F0) outperforms the baseline framework VAW-GAN (SID) for both male-to-male and male-to-female (5.51 vs. 6.20 and 6.57 vs. 7.39), which we believe is remarkable. The results indicate that our proposed VAW-GAN (SID+F0) with F0 conditioning achieves a better performance of spectrum conversion than the baseline framework VAW-GAN (SID) without any condition in both intergender and intra-gender SVC.



Fig. 4. MOS results with 95 % confidence interval between the proposed VAW-GAN (SID+F0) and VAW-GAN (SID) baseline.



Fig. 5. XAB preference test results with 95 % confidence interval between the proposed VAW-GAN (SID+F0) and VAW-GAN (SID) baseline.

C. Subjective Evaluation

We further conduct subjective evaluation to assess the performance of the proposed *VAW-GAN* for singing voice conversion in terms of voice quality and singer similarity. 20 subjects participate in all the listening tests, and each of them listens to 120 converted singing voice samples in total.

We conduct mean opinion score (MOS) [42], [43] to assess the voice quality of the converted singing voices. Listeners are asked to score the quality of the converted singing voice on a five-point scale (5: excellent, 4: good, 3: fair, 2: poor, 1: bad). As shown in Fig. 4, our proposed framework *VAW-GAN* (*SID+F0*) outperforms the baseline framework *VAW-GAN* (*SID*) in terms of voice quality by achieving higher MOS values of 3.03 ± 0.28 for male-to-male singing voice conversion and 2.90 ± 0.31 for male-to-female singing voice on F0 improves voice quality remarkably, which is consistent with the observation in objective evaluation.

We also conduct XAB preference test [44], [45] in terms of the singer similarity. The subjects are asked to listen to the reference target singing samples and the converted singing samples of the VAW-GAN (SID) baseline, and the proposed VAW-GAN (SID+F0), and choose the one which sounds closer to the target in terms of singer similarity. As shown in Fig. 5, our proposed framework VAW-GAN (SID+F0) outperforms the VAW-GAN (SID) baseline in terms of singer similarity (84.7 % vs. 13 % for male-to-male SVC and 56.7 % vs. 31 % for male-to-female SVC). The results prove the effectiveness of our proposed framework in terms of singer identity conversion.

V. CONCLUSION

In this paper, we propose a parallel-data-free singing voice conversion framework with VAW-GAN. We first propose to conduct singer-independent training with a encoding-decoding process. We also propose to condition the decoder on F0 to improve the singing voice conversion performance. We eliminate the need for parallel training data or other timealignment procedures, and achieve high performance of converted singing voices. Experimental results show the efficiency of our proposed SVC framework on both intra-gender and inter-gender singing voice conversion.

VI. ACKNOWLEDGEMENT

This work is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-100E-2018-006, AISG-GC-2019-002), the National Robotics Programme (Programmatic Grant Number: 192 25 00054), Human Robot Collaborative AI for AME Programmatic Grant (Programmatic Grant No. A18A2b0046) and Neuromorphic Computing Programmatic Grant (Programmatic Grant No. A1687b0033) from the Singapore Government's Research, Innovation and Enterprise 2020 plan in the Advanced Manufacturing and Engineering domain. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

This work is also supported by SUTD Start-up Grant Artificial Intelligence for Human Voice Conversion (SRG ISTD 2020 158) and SUTD AI Grant titled 'The Understanding and Synthesis of Expressive Speech by AI' (PIE-SGP-AI-2020-02).

REFERENCES

- F. Villavicencio and J. Bonada, "Applying voice conversion to concatenative singing-voice synthesis," in *Eleventh annual conference of the international speech communication association*, 2010.
- [2] K. Kobayashi, T. Toda, and S. Nakamura, "Intra-gender statistical singing voice conversion with direct waveform modification using logspectral differential," *Speech communication*, vol. 99, pp. 211–220, 2018.
- [3] Y. Luo, C. Hsu, K. Agres, and D. Herremans, "Singing voice conversion with disentangled representations of singer and vocal technique using variational autoencoders," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 3277–3281.
- [4] B. Sisman, J. Yamagishi, S. King, and H. Li, "An overview of voice conversion and its challenges: From statistical modeling to deep learning," *arXiv preprint arXiv:2008.03648*, 2020.
- [5] B. Sisman, M. Zhang, and H. Li, "A voice conversion framework with tandem feature sparse representation and speaker-adapted wavenet vocoder." in *Interspeech*, 2018, pp. 1978–1982.
- [6] 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.
- [7] B. Şişman, H. Li, and K. C. Tan, "Transformation of prosody in voice conversion," in 2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC). IEEE, 2017, pp. 1537–1546.
- [8] B. Sisman and H. Li, "Wavelet analysis of speaker dependent and independent prosody for voice conversion." in *Interspeech*, 2018, pp. 52–56.
- [9] B. Sisman, M. Zhang, and H. Li, "Group sparse representation with wavenet vocoder adaptation for spectrum and prosody conversion," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 6, pp. 1085–1097, 2019.
- [10] B. Sisman, G. Lee, H. Li, and K. C. Tan, "On the analysis and evaluation of prosody conversion techniques," in 2017 International Conference on Asian Language Processing (IALP). IEEE, 2017, pp. 44–47.
- [11] Y. Kawakami, H. Banno, and F. Itakura, "Gmm voice conversion of singing voice using vocal tract area function," *IEICE technical report. Speech (Japanese edition)*, vol. 110, no. 297, pp. 71–76, 2010.
- [12] H. Doi, T. Toda, T. Nakano, M. Goto, and S. Nakamura, "Singing voice conversion method based on many-to-many eigenvoice conversion and training data generation using a singing-to-singing synthesis system," in *Proceedings of The 2012 Asia Pacific Signal and Information Processing Association Annual Summit and Conference*, 2012, pp. 1–6.
- [13] B. Sisman, K. Vijayan, M. Dong, and H. Li, "SINGAN: Singing voice conversion with generative adversarial networks," 2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference, APSIPA ASC 2019, no. December, 2019.
- [14] T. Toda, Y. Ohtani, and K. Shikano, "One-to-many and many-to-one voice conversion based on eigenvoices," in 2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07, vol. 4. IEEE, 2007, pp. IV–1249.
- [15] K. Kobayashi, T. Toda, G. Neubig, S. Sakti, and S. Nakamura, "Statistical singing voice conversion with direct waveform modification based on the spectrum differential," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.

- [16] —, "Statistical singing voice conversion based on direct waveform modification with global variance," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [17] Y. Hono, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Singing voice synthesis based on generative adversarial networks," in *ICASSP* 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2019, pp. 6955–6959.
- [18] B. Sisman and H. Li, "Generative adversarial networks for singing voice conversion with and without parallel data," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 238–244.
- [19] 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.
- [20] 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), 2019, pp. 292–296.
- [21] W. Zhao, W. Wang, Y. Sun, and T. Tang, "Singing voice conversion based on wd-gan algorithm," in 2019 IEEE 4th Advanced Information Technology, Electronic and Automation Control Conference (IAEAC), vol. 1, 2019, pp. 950–954.
- [22] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "Stargan-vc: Nonparallel many-to-many voice conversion using star generative adversarial networks," in 2018 IEEE Spoken Language Technology Workshop (SLT). IEEE, 2018, pp. 266–273.
- [23] G. E. Henter, J. Lorenzo-Trueba, X. Wang, and J. Yamagishi, "Deep encoder-decoder models for unsupervised learning of controllable speech synthesis," arXiv preprint arXiv:1807.11470, 2018.
- [24] 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.
- [25] —, "Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks," in *Proc. Interspeech*, 2017, pp. 3364–3368.
- [26] H. Kameoka, T. Kaneko, K. Tanaka, and N. Hojo, "Acvae-vc: Nonparallel many-to-many voice conversion with auxiliary classifier variational autoencoder," arXiv preprint arXiv:1808.05092, 2018.
- [27] —, "Acvae-vc: Non-parallel voice conversion with auxiliary classifier variational autoencoder," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 9, pp. 1432–1443, 2019.
- [28] P. L. Tobing, Y.-C. Wu, T. Hayashi, K. Kobayashi, and T. Toda, "Nonparallel voice conversion with cyclic variational autoencoder," arXiv preprint arXiv:1907.10185, 2019.
- [29] D. Yook, S.-G. Leem, K. Lee, and I.-C. Yoo, "Many-to-many voice conversion using cycle-consistent variational autoencoder with multiple decoders," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 215–221.
- [30] B. Sisman, M. Zhang, M. Dong, and H. Li, "On the study of generative adversarial networks for cross-lingual voice conversion," *IEEE ASRU*, 2019.
- [31] K. Zhou, B. Sisman, M. Zhang, and H. Li, "Converting anyone's emotion: Towards speaker-independent emotional voice conversion," arXiv preprint arXiv:2005.07025, 2020.
- [32] K. Zhou, B. Sisman, R. Liu, and H. Li, "Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset," 2020.
- [33] K. Zhou, B. Sisman, and H. Li, "Transforming Spectrum and Prosody for Emotional Voice Conversion with Non-Parallel Training Data," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 230–237.
- [34] J. Lorenzo-Trueba, F. Fang, X. Wang, I. Echizen, J. Yamagishi, and T. Kinnunen, "Can we steal your vocal identity from the internet?: Initial investigation of cloning obama's voice using gan, wavenet and lowquality found data," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 240–247.
- [35] H. Kawahara, I. Masuda-Katsuse, and A. De Cheveigne, "Restructuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based f0 extraction: Possible role of a repetitive structure in sounds," *Speech communication*, vol. 27, no. 3-4, pp. 187–207, 1999.
- [36] M. Morise, F. Yokomori, and K. Ozawa, "World: a vocoder-based high-quality speech synthesis system for real-time applications," *IEICE*

TRANSACTIONS on Information and Systems, vol. 99, no. 7, pp. 1877–1884, 2016.

- [37] K. Qian, Z. Jin, M. Hasegawa-Johnson, and G. J. Mysore, "F0-consistent many-to-many non-parallel voice conversion via conditional autoencoder," 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), May 2020.
- Signal Processing (ICASSP), May 2020.
 [38] W.-C. Huang, Y.-C. Wu, C.-C. Lo, P. L. Tobing, T. Hayashi, K. Kobayashi, T. Toda, Y. Tsao, and H.-M. Wang, "Investigation of F0 Conditioning and Fully Convolutional Networks in Variational Autoencoder Based Voice Conversion," in *Proc. Interspeech 2019*, 2019, pp. 709–713.
- [39] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," arXiv preprint arXiv:1701.07875, 2017.
- [40] Z. Duan, H. Fang, B. Li, K. Chai Sim, and Y. Wang, "The nus sung and spoken lyrics corpus: A quantitative comparison of singing and speech," *APSIPA*, 2013.
- [41] R. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," *Communications, Computers and Signal Processing*, pp. 125–128, 1993.
- [42] R. C. Streijl, S. Winkler, and D. S. Hands, "Mean opinion score (mos) revisited: methods and applications, limitations and alternatives," *Multimedia Systems*, vol. 22, no. 2, pp. 213–227, 2016.
- [43] R. Liu, B. Sisman, J. Li, F. Bao, G. Gao, and H. Li, "Teacher-student training for robust tacotron-based tts," in 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2020.
- [44] E. Helander, T. Virtanen, J. Nurminen, and M. Gabbouj, "Voice conversion using partial least squares regression," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 5, pp. 912–921, 2010.
 [45] R. Liu, B. Sisman, F. Bao, G. Gao, and H. Li, "Wavetts: Tacotron-
- [45] R. Liu, B. Sisman, F. Bao, G. Gao, and H. Li, "Wavetts: Tacotronbased tts with joint time-frequency domain loss," *arXiv preprint arXiv*:2002.00417, 2020.