



Adversarially Trained Multi-Singer Sequence-To-Sequence Singing Synthesizer

Jie Wu, Jian Luan

Xiaoice, Software Technology Center Asia, Microsoft

{jiewu, jianluan}@microsoft.com

Abstract

This paper presents a high quality singing synthesizer that is able to model a voice with limited available recordings. Based on the sequence-to-sequence singing model, we design a multi-singer framework to leverage all the existing singing data of different singers. To attenuate the issue of musical score unbalance among singers, we incorporate an adversarial task of singer classification to make encoder output less singer dependent. Furthermore, we apply multiple random window discriminators (MRWDs) on the generated acoustic features to make the network be a GAN. Both objective and subjective evaluations indicate that the proposed synthesizer can generate higher quality singing voice than baseline (4.12 vs 3.53 in MOS). Especially, the articulation of high-pitched vowels is significantly enhanced.

Index Terms: singing voice synthesis, multi-singer, adversarial training, GAN, sequence-to-sequence modeling

1. Introduction

Singing voice synthesis (SVS) is a task to synthesize specific singer's singing voice from musical score (e.g. lyric, melody and rhythm). Recent years, many deep learning based methods have been introduced into SVS to generate high quality singing voices, such DNN [1] and LSTM [2]. In addition, auto-regressive models, like Tacotron2 [3] have been successfully applied to SVS task [4, 5, 6].

Although auto-regressive model can achieve high quality, it suffers from exposure bias and time-consuming inference due to the forward dependency. To avoid these issues, Blaauw [7] proposed a sequence-to-sequence (Seq2Seq) singing synthesizer based on feed-forward transformer architecture, which can generate acoustic features in parallel. Feed-forward transformer has also shown its superior performance in text-to-speech (TTS) task [8]. However, to achieve a high performance, Seq2Seq singing synthesizer requires a large amount of training data from one singer. It is hard and expensive to collect them in customization application scenario. In order to reduce the amount of training data for target singer, we expect to construct a multi-singer Seq2Seq model by leveraging many existing singing data of other singers.

To build a multi-singer singing model with limited training data, one challenge is data unbalance issue. It refers to the unbalance distribution of training data among singers, such as lyric and melody. The deviation of training data distribution among singers might be considered as the singer's identification during training. To attenuate this issue, an adversarial loss [9] is incorporated by employing a singer classifier to encourage the encoder to learn singer-independent representation from musical scores. Adversarial training has demonstrated its ability in many fields, like cross-language voice cloning [10] and singing voice conversion [11, 12].

With recent development of generative adversarial net (GAN) [13] in many tasks, such as image generation [14, 15], text-to-speech [16, 17] and others [18, 19], it has also been successfully applied to SVS field as a powerful generative model. Hono [20] proposed a DNN-based SVS system with conditional GAN to optimize the distribution of acoustic features. Lee [21] adopted a conditional adversarial training network in end-to-end SVS system. WGANSing [22] presented a block-wise generative singing model with Wasserstein-GAN [23] framework. In [24], an auto-regressive singing model based on boundary equilibrium GAN was proposed. However, all the aforementioned GAN-based SVS systems only adopted a single discriminator directly operating on the whole sample sequences. They have a limitation, that is the lack of diversity in the sample distribution evaluation. Different from them, we introduce multiple random window discriminators (MRWDs) [16] into the multi-singer singing model to make the network be a GAN. It is an ensemble of discriminators operating on random sub-sampled fragments of samples. MRWDs allows for the evaluation of samples in different complementary ways, where it analyses the samples in the general realism, as well as in the correspondence between the generated samples and input conditions. Moreover, using different sizes of random windows, rather than the whole sample sequences, has a data augmentation effect. It is quite helpful since the training data are limited for each singer.

In this paper, we propose a multi-singer sequence-to-sequence singing model based on adversarial training strategy. Our contributions includes: (1) Scale a Seq2Seq network to support multi-singer training. It improves the performance when only limited recordings are available for one singer. (2) Incorporate an adversarial task of singer classification to make encoder output less singer dependent to handle the data unbalance issue. (3) Apply multiple random window discriminators on the generated acoustic features to make the network be a GAN.

Similar to [7], we only focus on spectrum modeling and assume F_0 are given. Meanwhile, to avoid the impact of different duration models, ground-truth phoneme duration are used in both training and inference phases. WORLD vocoder [25] is used to extract acoustic features, which allows the explicit control of F_0 . In this way, we evaluate the contribution of each component and demonstrate our proposed singing model.

2. Single-singer Seq2Seq SVS system

Following [7], the system is shown in the solid lines of Figure 1, including encoder, length regulator and decoder.

Firstly, the encoder takes phoneme embedding, pitch embedding and position encoding from musical score as input and then obtain its score encoding through a series of gated linear units (GLU) blocks [26]. GLU block is a convolutional block including a 1-D convolutional layer with gated linear

units and a residual connection. The length regulator expands the score encoding from phoneme-level to frame-level according to phoneme duration. Finally, the feed-forward decoder transforms the expanded score encoding sequence to its corresponding acoustic feature sequence through several GLU blocks and self-attention layers [27].

3. Proposed Architecture

As illustrated in the dotted lines of Figure 1, the proposed architecture consists of three modules: multi-singer SVS module, singer classifier and multiple random window discriminators.

To focus on the spectrum and exclude the influence of duration model in prosody, in length regulator, we use ground-truth phoneme alignment in both training and inference phases.

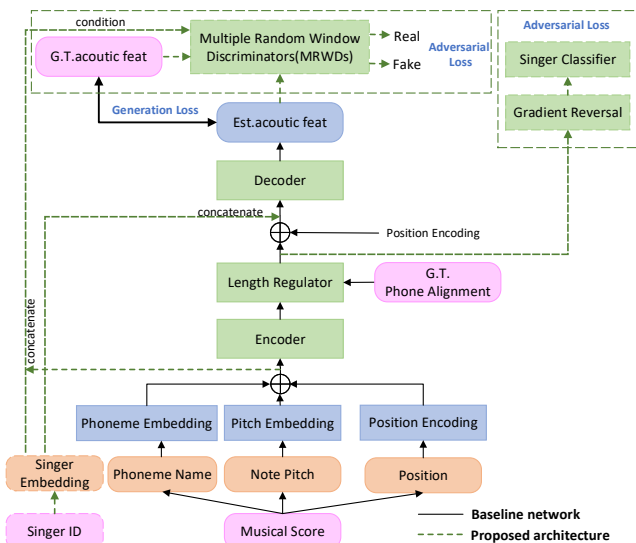


Figure 1: Overview of the proposed adversarially trained multi-singer sequence-to-sequence singing synthesizer. Green dotted lines represents additional components in the proposed architecture while black solid lines is the baseline network.

3.1. Multi-singer Seq2Seq SVS System

For the purpose of reducing the training data size of target singer, one solution is to leverage singing data from other singers. For example, in [29], a learned singer embedding from acoustic features was used to represent singer identity. An singer identity encoder was designed to produce singer’s identity vector in SVS system [30].

Following [31], we construct our multi-singer singing model with trainable singer embedding. These singer embedding vectors are based on singer ID, which are initially randomized and then updated during training. It is concatenated with encoder output and frame position encoding as the input of decoder.

We use the cross-entropy loss function for voiced/unvoiced (VUV) flag prediction while L1 loss for mel-generalized coefficient (MGC) and band aperiodicity (BAP) prediction. Therefore the generation loss for acoustic features is:

$$L_G = L_1^{mgc} + L_1^{bap} + L_{ce}^{vuv} \quad (1)$$

3.2. Singer Classifier

One challenge in multi-singer SVS system is the data unbalance issue among singers, which is also mentioned in cross-language

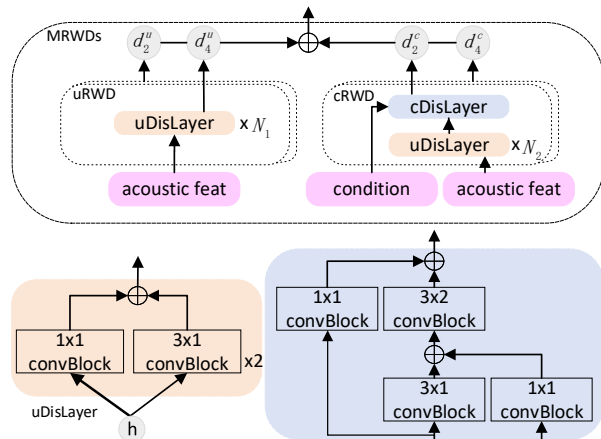


Figure 2: The architecture of multiple random window discriminators (MRWDs). It contains 2 unconditional (uRWD, left) and 2 conditional (cRWD, right) discriminators with [32,64] random window size. The convBlock is a 1-D convolutional network with ReLU activation and spectral normalization [28].

voice cloning task [10]. To address this problem, we incorporate an adversarial loss by employing a singer classifier based on the musical score encoding. The encoder is expected to learn a latent representation which is strong on acoustic feature generation task while weak on singer classification. In other words, singer identity information is removed from encoder output.

As usual, a gradient reversal layer (GRL) [9] is applied prior to the singer classifier. It multiplies the gradient by a certain negative constant (λ) during back propagation of training while only acts an identity transformation during forward propagation.

For singer classifier, cross entropy loss function is applied as below:

$$L_{adv_singer} = \sum_i L_{ce}(S(E(x_i)), s_i) \quad (2)$$

where S is singer classifier and E represents encoder, i means the i -th singer and s_i is the singer identity label.

3.3. Multiple Random Window Discriminators

Expecting to generate more realistic singing voice, we adopt multiple random window discriminators (MWRDs) [16] to make the network be a GAN.

As shown in Figure 2, MRWDs contains unconditional and conditional discriminators with random window sizes (uRWD and cRWD). All these discriminators are independent of each other. It can measure the similarity of distribution between generated and real acoustic features as well as consider whether the generated samples match the input conditions (e.g., phoneme, pitch and singer identity). Furthermore, random window sizes in MRWDs allow different sub-segments feeding into these discriminators. On the one hand, using different sizes of random window has a data augmentation effect. It is very helpful since the training data are limited for each singer. On the other hand, feeding sub-segments instead of the whole sample sequence can also reduce the computational complexity in training. Finally, outputs of all discriminators are added together as the final MRWDs output to calculate the loss. The final MRWDs output is:

$$MRWD_{S(32,64)} = \sum_{k \in (32,64)} (uRWD_k + cRWD_k)$$

Following the vanilla GAN loss in [13], the adversarial loss terms (L_{adv_D} and L_{adv_G}) for MRWDs D and generator G are shown below:

$$L_{adv_D} = -\mathbb{E}_{y \sim P_r} [\log(D(y))] - \mathbb{E}_{x \sim P_x} [\log(1 - D(G(x)))] \quad (3)$$

$$L_{adv_G} = \mathbb{E}_{x \sim P_x} [\log(1 - D(G(x)))] \quad (4)$$

where y is a sample from real data distribution and x is the input of the generator. Different from the vanilla GAN [13], the input x is musical score feature instead of random noise.

4. Experiments

4.1. Dataset

- **Single-singer SVS system:** 770 Chinese pop songs (total ~ 10 hours) of clean recordings are collected from a female singer ‘‘F1’’.
- **Multi-singer SVS system:** 200 songs (~ 3.5 hours) of ‘‘F1’’ are randomly selected from 770 songs. Meanwhile, about 200 Chinese pop songs are also collected from six singers, named ‘‘F2’’, ‘‘F3’’, ‘‘F4’’, ‘‘M1’’, ‘‘M2’’ and ‘‘M3’’ respectively¹. The specific number of songs from each singer are listed in Table 1.

Table 1: The number of songs from seven singers.

	F1	F2	F3	F4	M1	M2	M3
Number of songs	200	200	200	210	205	200	153

All the recordings are collected in a professional recording studio while singers listen to the accompaniment through headphones. The musical scores are hereafter manually revised according to the audios. The phoneme duration is obtained by HMM-based force alignment tool [32].

Recordings are sampled at 48kHz with 16bits per sample in mono. They are segmented into pieces at silence boundaries. Each piece is forced to be not longer than 10 seconds. Acoustic features are then extracted from audio pieces at every 15ms by WORLD vocoder, including 60-dimensional MGC, 5-dimensional BAP and 1-dimensional VUV flag.

4.2. Experimental setup

In the experiments, lyrics and note pitch in MIDI standard format [33] are obtained from musical score files. The sizes of phoneme and note pitch vocabularies are 71 and 84 respectively. As shown in Figure 1, phoneme and note pitch are both embedded into 384-dimensional vectors and added together with position encoding as encoder input. Besides, additional 64-dimensional trainable singer embedding vector is used to represent singer identity.

Following the design in [7], the networks of encoder and decoder are shown in Figure 3. The encoder contains two linear layers and a single GLU block. The input/output size of each layer is 384/256, 256/64 and 64/384 respectively. Decoder is a 6-layers network. Each layer consists of a single-head self-attention sub-layer and a GLU sub-layer, all with 448 channels. The output linear layer maps the 448-dimensional hidden vector to 66-dimensional acoustic feature vector. The kernel sizes of convolutional layers in encoder and decoder are both 3.

The singer classifier is a 2-layers 1-D convolutional network with ReLU activation and spectral normalization [28]. The kernel sizes are all set to be 3, with input/output size of

¹Here, ‘‘F’’ means female singer while ‘‘M’’ means male singer.




Figure 3: A diagram of the encoder and decoder networks, where GLU block, GLU sub-layer and attention sub-layer follows the design of [7].

384/128 for the first layer and 128/128 for the second layer. The output linear layer finally converts the 128-dimensional hidden vector into a 7-dimensional singer identity vector with softmax activation.

As shown in the top of Figure 2, MRWDs consists of 2 uRWDs and 2 cRWDs with input window size of [32,64] in frames respectively. uRWD is a stack of 4 uDisLayers while cRWD contains 3 cDisLayers and 1 cDisLayer. The networks of uDisLayer and cDisLayer are shown in the bottom of Figure 2, which are both a stack of convolutional networks with different kernel sizes. uRWD and cRWD have the same output channels of [64, 128, 256, 1] for 4 layers. In cRWDs, the input condition is a 448-dimensional vector which is the concatenation of 384-dimensional score encoding and 64-dimensional singer embedding. Finally, all output of these 4 discriminators are added together as the final MRWDs output to calculate loss.

In the training of GAN, discriminator and generator are updated alternately sharing the initial learning rate of 1×10^{-4} . The mini batch size is 32 and the Adam optimizer is with $\beta_1 = 0.9, \beta_2 = 0.98$. Dropout probability is 0.1 through the entire model.

We implement five Seq2Seq SVS systems to evaluate the contributions of three proposed modules. The differences among the five systems are described in Table 2. The baseline and three modules are:

- **baseline:** Single-singer singing model.
- **module1:** Multi-singer singing model.
- **module2:** Singer classifier.
- **module3:** Multiple random window discriminators.

In the training phase, the total generative loss is represented as follows:

$$L_{total} = \lambda_G L_G + \lambda_S L_{adv_{singer}} + \lambda_D L_{adv_G} \quad (5)$$

where λ_G, λ_S and λ_D are the weights for the losses of the three modules.

Five systems have different weights for each module, which are specifically listed in Table 2.

Table 2: Five systems description of modules and loss weights.

Systems	System1	System2	System3	System4	System5
Modules	baseline	+(module1)	+(module1,2)	+(module1,3)	+(module1,2,3)
Loss weights	[1,0,0]	[1,0,0]	[1,1,0]	[10,0,1]	[10,2,1]

4.3. Evaluation

To evaluate the performance of different systems, we conduct objective and subjective evaluations on the target singer ‘‘F1’’.

4.3.1. Objective Evaluation

In objective test, averaged global variances (GVs) [34] of MGC is calculated for different systems. As shown in Figure 4, system2 and system3 have similar averaged GV and both slightly better than System1. It indicates that multi-singer module can even improve the averaged GV with less training data of target singer. Furthermore, the GV trajectories of system4 and system5 are on par but much closer to ground truth than others. It shows that MRWDs can significantly alleviate the over-smoothing problem of generated features. Also, it can be noted that singer classifier has no obvious impact to GV.




Figure 4: Averaged global variances (GVs) of mel-generalized coefficient (MGC) from evaluation set for different systems.

We visualize the mel-spectrograms of a testing sample generated by different systems in Figure 5. These green rectangles refer to a character “qi a_h nn_h” whose note pitch is 80 in MIDI standard. As shown in the green rectangles, the mel-spectrogram of the character in System1 is almost invisible while in System2 it has normal energy although quite fuzzy harmonics. This phenomenon can be attributed to the benefit of leveraging data from other singers. Based on it, both System3 and System4 show correct and clear harmonics against System2. It indicates that either singer classifier or MRWDs can obviously enhance the articulation of high-pitched vowels. Moreover, when combining singer classifier and MRWDs together, System5 achieves further improvement and defeat all other systems. Overall, from System1 to System5, it can be easily found that the proposed synthesizer enhances the articulation of high-pitched vowels step by step.




Figure 5: Mel-spectrograms of a testing sample generated by different systems. Green rectangles show their differences of the same character.

4.3.2. Subjective Evaluation

For subjective evaluation, the female singer “F1” is regarded as target singer. WORLD vocoder is used to synthesize the singing voice with generated MGC, BAP and ground-truth $F0$.

Firstly, Mean Opinion Score (MOS) test is carried out to evaluate the quality of generated voices from different systems. 10 listeners are invited to judge 30 samples in each system. Each sample is not longer than 10 seconds.

As shown in Figure 6, System2 achieves a higher MOS score than System1, proving the benefit of leveraging data of other singers. Multi-singer singing model can improve the quality even with less singing data of target singer. Meanwhile, the MOS scores of System3 and System4 are both higher than that of System2, confirming the effectiveness of integrating singer classifier or MRWDs individually. Our proposed System5 achieves the highest score of 4.12. It demonstrates that further advancement can be achieved if combining singer classifier and MRWDs together.

Particularly, two A/B preference tests are conducted to validate the effect of singer classifier: System2 vs. System3 and System4 vs. System5. For each system, 30 samples are generated and each one contains 1-5 vowels whose pitch ranges from 78 to 81 in MIDI standard. 10 listeners are required to listen to all the pairs and choose which one is better on the articulation.

The results are shown in Figure 7. Both A/B tests prove the overwhelming advantage of systems with adversarial task of singer classification. They are mainly contributed by the model quality improvement on sparse input (such as high pitch). It is likely because the model of one singer can better leverage others’ data when the distribution of encoder output is less singer dependent.




Figure 6: Mean Opinion Score (MOS) test results of singing voice quality with 95% confidence intervals.




Figure 7: A/B preference tests on the articulation of high-pitched vowels between different systems. The p -values for the two bars are 1.14×10^{-9} and 2.17×10^{-9} respectively.

5. Conclusions

In this paper, we extend conventional sequence-to-sequence singing synthesizer to a multi-singer one with trainable singer embedding which can model each singer well with limited recording data. Moreover, adversarial training strategies are introduced into the multi-singer synthesizer. One is the singer classifier to exclude singer identity information from encoder output. The other is MRWDs for acoustic feature generation loss to help the synthesizer to generate more realistic singing voice. Experimental results prove that our methods can achieve higher quality singing voice and especially significant improvement on the articulation of high-pitched vowels. By this method, only limited recordings are required from target singer for a high quality singing model. In other words, the singing voice customization becomes easy and economic. Some samples for subjective evaluation are available via this link².

²<https://jiewu-demo.github.io/INTERSPEECH2020/>

6. References

- [1] M. Nishimura, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Singing voice synthesis based on deep neural networks," in *Interspeech*, 2016, pp. 2478–2482.
- [2] J. Kim, H. Choi, J. Park, S. Kim, J. Kim, and M. Hahn, "Korean singing voice synthesis system based on an lstm recurrent neural network," in *Proc. INTERSPEECH*, 2018.
- [3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-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.
- [4] M. Blaauw and J. Bonada, "A neural parametric singing synthesizer modeling timbre and expression from natural songs," *Applied Sciences*, vol. 7, no. 12, p. 1313, 2017.
- [5] Y.-H. Yi, Y. Ai, Z.-H. Ling, and L.-R. Dai, "Singing voice synthesis using deep autoregressive neural networks for acoustic modeling," *arXiv preprint arXiv:1906.08977*, 2019.
- [6] R. Valle, J. Li, R. Prenger, and B. Catanzaro, "Mellotron: Multispeaker expressive voice synthesis by conditioning on rhythm, pitch and global style tokens," *arXiv preprint arXiv:1910.11997*, 2019.
- [7] M. Blaauw and J. Bonada, "Sequence-to-sequence singing synthesis using the feed-forward transformer," *arXiv preprint arXiv:1910.09989*, 2019.
- [8] 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.
- [9] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [10] Y. Zhang, R. J. Weiss, H. Zen, Y. Wu, Z. Chen, R. Skerry-Ryan, Y. Jia, A. Rosenberg, and B. Ramabhadran, "Learning to speak fluently in a foreign language: Multilingual speech synthesis and cross-language voice cloning," *arXiv preprint arXiv:1907.04448*, 2019.
- [11] E. Nachmani and L. Wolf, "Unsupervised singing voice conversion," *arXiv preprint arXiv:1904.06590*, 2019.
- [12] C. Deng, C. Yu, H. Lu, C. Weng, and D. Yu, "Pitchnet: Unsupervised singing voice conversion with pitch adversarial network," *arXiv preprint arXiv:1912.01852*, 2019.
- [13] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [14] S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative adversarial text to image synthesis," *arXiv preprint arXiv:1605.05396*, 2016.
- [15] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [16] M. Bińkowski, J. Donahue, S. Dieleman, A. Clark, E. Elsen, N. Casagrande, L. C. Cobo, and K. Simonyan, "High fidelity speech synthesis with adversarial networks," *arXiv preprint arXiv:1909.11646*, 2019.
- [17] B. Bollepalli, L. Juvela, and P. Alku, "Generative adversarial network-based glottal waveform model for statistical parametric speech synthesis," *arXiv preprint arXiv:1903.05955*, 2019.
- [18] S. Pascual, A. Bonafonte, and J. Serra, "Segan: Speech enhancement generative adversarial network," *arXiv preprint arXiv:1703.09452*, 2017.
- [19] L.-C. Yang, S.-Y. Chou, and Y.-H. Yang, "Midinet: A convolutional generative adversarial network for symbolic-domain music generation," *arXiv preprint arXiv:1703.10847*, 2017.
- [20] 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.
- [21] J. Lee, H.-S. Choi, C.-B. Jeon, J. Koo, and K. Lee, "Adversarially trained end-to-end korean singing voice synthesis system," *arXiv preprint arXiv:1908.01919*, 2019.
- [22] P. Chandna, M. Blaauw, J. Bonada, and E. Gómez, "Wgansing: A multi-voice singing voice synthesizer based on the wasserstein-gan," in *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019, pp. 1–5.
- [23] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 214–223.
- [24] S. Choi, W. Kim, S. Park, S. Yong, and J. Nam, "Korean singing voice synthesis based on auto-regressive boundary equilibrium gan," in *ICASSP 2020*, 2020, pp. 7234–7238.
- [25] 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.
- [26] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 2017, pp. 933–941.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [28] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," *arXiv preprint arXiv:1802.05957*, 2018.
- [29] M. Blaauw, J. Bonada, and R. Daido, "Data efficient voice cloning for neural singing synthesis," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6840–6844.
- [30] J. Lee, H.-S. Choi, J. Koo, and K. Lee, "Disentangling timbre and singing style with multi-singer singing synthesis system," *arXiv preprint arXiv:1910.13069*, 2019.
- [31] W. Ping, K. Peng, A. Gibiansky, S. O. Arik, A. Kannan, S. Narang, J. Raiman, and J. Miller, "Deep voice 3: Scaling text-to-speech with convolutional sequence learning," *arXiv preprint arXiv:1710.07654*, 2017.
- [32] K. Sjölander, "An hmm-based system for automatic segmentation and alignment of speech," in *Proceedings of Fonetik*, vol. 2003, 2003, pp. 93–96.
- [33] M. M. Association, "Midi manufacturers association," <https://www.midi.org>.
- [34] 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.