



# WeSinger: Data-augmented Singing Voice Synthesis with Auxiliary Losses

Zewang Zhang, Yibin Zheng, Xinhui Li, Li Lu

Tencent Inc, Guangzhou, China

{zewangzhang, hiccupli}@tencent.com

## Abstract

In this paper, we develop a new multi-singer Chinese neural singing voice synthesis (SVS) system named WeSinger. To improve the accuracy and naturalness of synthesized singing voice, we design several specific modules and techniques: 1) A deep bi-directional LSTM-based duration model with multi-scale rhythm loss and post-processing step; 2) A Transformer-like acoustic model with progressive pitch-weighted decoder loss; 3) a 24 kHz pitch-aware LPCNet neural vocoder to produce high-quality singing waveforms; 4) A novel data augmentation method with multi-singer pre-training for stronger robustness and naturalness. To our knowledge, WeSinger is the first SVS system to adopt 24 kHz LPCNet and multi-singer pre-training simultaneously. Both quantitative and qualitative evaluation results demonstrate the effectiveness of WeSinger in terms of accuracy and naturalness, and WeSinger achieves state-of-the-art performance on the recent public Chinese singing corpus Opencpop<sup>1</sup>. Some synthesized singing samples are available online<sup>2</sup>.

**Index Terms:** singing voice synthesis, multi-scale rhythm loss, pitch-weighted progressive loss, data augmentation, 24 kHz lpcnet vocoder

## 1. Introduction

Recently, the SVS task has drawn increasing attention for its various potential applications in entertainment and multi-model technologies. Unlike the traditional text-to-speech (TTS) pipeline, a robust SVS (singing voice synthesis) system aims to generate accurate and natural singing voices from linguistic and musical features such as the phoneme, tempo, pitch, slur, and duration. A typical SVS system consists of three essential components: an acoustic model to convert lyrics and musical information into acoustic features, a duration model to predict the duration of each phoneme, and a neural vocoder model to generate singing voices from acoustic features. In the past, some SVS systems have been designed on the unit selection [1, 2] or statistical parametric speech synthesis (SPSS) methods like context-dependent hidden Markov model (HMM) [3]. Nevertheless, the quality of synthesized singing voices cannot reach the naturalness of ground-truth songs owing to the over-smoothing effects and lack of generalization.

Thanks to the emergence of deep neural network (DNN) [4] and some significant gains obtained in TTS, including popular acoustic models [5, 6] and waveform generation algorithms [7, 8, 9], a great deal of improved DNN-based models have also been suggested subsequently for acoustic modeling in SVS systems [10, 11, 12, 13]. XiaoIceSing [14] employs acoustic model inspired by FastSpeech [6] with a tradition vocoder World [15] to outperform the conventional DNN-based systems. HiFiSinger [16], N-Singer [17] and CpopSing [18]

<sup>1</sup><https://wenet.org.cn/opencpop/>

<sup>2</sup><https://zzw922cn.github.io/wesinger/>

are equipped with Transformer-like acoustic model and GAN-based neural vocoder to achieve accurate and high-fidelity singing synthesis. In recent literature, another fast GAN-based multi-singer singing voice vocoder Multi-Singer [19] is proposed to improve the performance of singing voice synthesis for the unseen singers. Apart from the feed-forward paradigm, an auto-regressive Chinese SVS system ByteSing [20] is designed based on the Tacotron-based [5] acoustic models with attention mechanism and WaveRNN [7] neural vocoder. Moreover, the pipeline of building data directly mined from the web without any high-quality singing data is also explored in DeepSinger [21].

Although previous attempts bring different degrees of performance improvements in acoustic modeling or parallel neural vocoder for SVS tasks, they pay little or no attention to better rhythm modeling, such as the dependency between syllables and phonemes in Chinese. Meanwhile, previous works either adopt Mel-spectrogram without introducing pitch information explicitly as the target for acoustic modeling or consider pitch information as acoustic feature input for GAN-based neural vocoder. As a result, the acoustic model ignores the importance of pitch for SVS and may cause the predicted F0 contour out of tune, while the GAN-based neural vocoders may produce some unsatisfactory artifacts in the quality of synthesized singing voices especially when the human ear is sensitive to the coherence of sound. In addition, how to utilize multi-singer pre-training with data augmentation method to improve naturalness and robustness for SVS tasks has also not been explored yet.

To address the above challenges and meet the demands of naturalness in rhythm and accuracy in pitch, we propose a new end-to-end SVS system named WeSinger to imitate the voice and style of any singer. Our contributions are summarized below. First, we design two auxiliary loss functions for optimizing the duration model and acoustic model effectively, including multi-scale rhythm loss and progressive pitch-weighted loss. Second, to achieve the high-quality and pitch-aware synthesis, we adopt a 24 kHz pitch-aware LPCNet vocoder to convert acoustic features into singing waveforms. Last, we propose a data augmentation method and combine the augmented multi-singer corpus with pre-training and then do fine-tuning to further enhance the performance of WeSinger. Experimental results indicate that the data-augmented WeSinger with auxiliary losses can synthesize singing voices with much higher quality.

## 2. Methodology

### 2.1. Architecture

The overall architecture of WeSinger is described in Figure 1, each module of which will be introduced in detail.

#### 2.1.1. Feature Representation

Given MusicXML-format [22] files and human-annotated phoneme-level interval files, we first convert the Chinese syllable

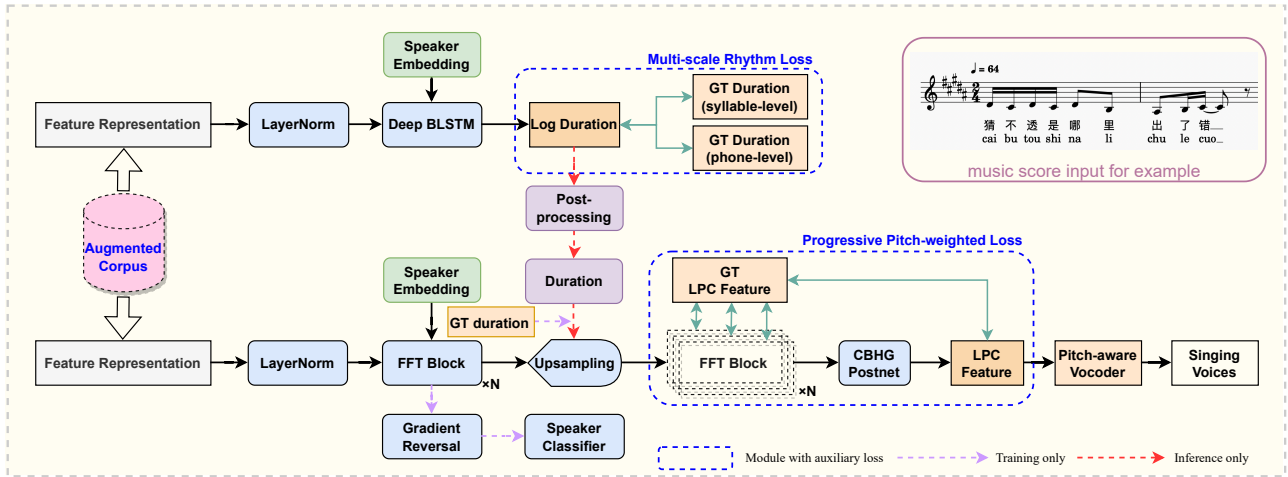


Figure 1: The system architecture of our proposed SVS system WeSinger. First, each song recording is collected with MusicXML-format [22] transcription and human-annotated phoneme-level interval file. Then, a data augmentation method of variable-length segmentation is applied to attain a larger multi-singer corpus. The training stage contains multi-speaker pre-training and fine-tuning with data of a specific singer. At the synthesis stage, WeSinger first predicts the duration of each phoneme by duration model with a post-processing procedure and then produces pitch-aware LPC features as input for the pitch-aware 24 kHz LPCNet vocoder.

Table 1: Abbreviations and descriptions of some input symbols in WeSinger.

Abbreviation	Description
<b>Ph</b>	Chinese phoneme (ch, ang)
<b>Pt</b>	Chinese phoneme type (initial, final, single final)
<b>Pi</b>	Note pitch (C4, D4)
<b>Sr</b>	Slur flag (start, continue, stop, null)
<b>Bt</b>	Duration of note beat (frames)

ble sequences in lyrics to phoneme sequences using the Chinese grapheme-to-phoneme algorithm and map the combination of *step*, *octave* and *alter* into a single note pitch according to the MIDI standard [23]. Chinese syllables are usually composed of initials and finals. In this case, we build phoneme-type sequences from phoneme sequences for richer feature representation. The duration of each syllable is calculated by its associated notes and beats per minute (BPM), and we observed that it works best to divide the duration evenly between the syllable’s initial and final. Notably, we don’t make use of Chinese tone information due to the existence of musical notes. Furthermore, we assign each phoneme a slur type from *start*, *continue*, *stop* and *null* to indicate whether to be played smoothly or not. The ground-truth duration of each phoneme is extracted from the annotated interval file. All durations are converted to the number of frames with a unit of 10 milliseconds. Finally, we rearrange all of the pitch sequence, slur sequence, beat sequence, phoneme-type sequence, and duration sequence to manually match the phoneme sequence. For convenience, we abbreviate some input symbols and illustrate them in Table 1.

### 2.1.2. Duration Model

One discrepancy between TTS and SVS is that the starting and ending time of each phoneme in SVS systems is largely con-

strained by music scores, while traditional TTS systems often do duration prediction with higher degrees of freedom. We evaluate the different forms of feature representation and finally take the most informative and effective phoneme-level  $\mathbf{X}_D = [\mathbf{Ph}, \mathbf{Pt}, \mathbf{Pi}, \mathbf{Sr}, \mathbf{Bt}]$  as the input of BLSTM-based duration model and log-scale duration as the target. We also find that it can make training more stably and efficiently by adopting the layer normalization [24] technique for the concatenated input embedding sequences. To make the rhythm of synthesized singing voice more accurate and natural, we design a multi-scale rhythm loss including the commonly used phoneme-level minimum mean squared error (MMSE) loss and an auxiliary syllable-level MMSE loss to optimize the duration model. Normally, for a syllable sequence  $(s_1, s_2, \dots, s_M)$  and its phoneme-level sequence  $(p_1, p_2, \dots, p_N)$ , if we denote the ground-truth phoneme-level duration as  $d_{p_i}$  and the predicted phoneme-level duration as  $\hat{d}_{p_i}$ , thus, the loss function of the duration model can be written as:

$$\mathcal{L}_{dur} = \frac{1}{M} \sum_{i=1}^M \left| \sum_{p_k \in s_i} (d_{p_k} - \hat{d}_{p_k}) \right|^2 + \frac{1}{N} \sum_{j=1}^N |d_{p_j} - \hat{d}_{p_j}|^2 \quad (1)$$

Indeed, there are time-lags [11] between the starting time of musical notes and the starting time of ground-truth audios. To enforce the constraint that the duration of a syllable should be consistent with musical notes, we adopt a post-processing procedure to refine the predicted duration by scaling the predicted duration of phonemes in each syllable in proportion to the duration of corresponding musical notes, limiting the duration of its consonant to a maximum of 100 milliseconds according to the duration distribution of the initials in the training set, and allocating the remaining duration to the final.

### 2.1.3. Acoustic Model

The acoustic model of WeSinger consists of an encoder-decoder module modified from FastSpeech [6]. A carefully designed phoneme-level  $\mathbf{X}_D = [\mathbf{Ph}, \mathbf{Pt}, \mathbf{Pi}, \mathbf{Sr}]$  is represented as the input for encoder. Apart from adopting the layer normalization tech-

nique for the concatenated embedded input sequences and substituting the intermediate fully-connected layers of each Transformer FFT block with Conv1D layers for stronger context-dependent performance[25], we made two key improvements for better acoustic modeling. First, to discourage the encoder from also memorizing speaker characteristics in multi-singer pre-training and improve the generalization, we employ the domain adversarial training (DAT) strategy [26] by appending a speaker classifier with a gradient reversal layer (GRL) [27, 28] to the encoder outputs. We find it effective to set the weight for GRL loss as 0.02. To our knowledge, this is the first time that GRL is successfully applied to the training of a multi-singer SVS system. Second, in contrast to the commonly-used Mel-spectrogram as the target [16, 20], we choose to adopt the pitch-aware features including Bark-scale Frequency Cepstral Coefficients (BFCCs) and pitch information as the intermediate target for WeSinger. We design a progressive pitch-weighted loss function for optimizing WeSinger’s acoustic model. Specifically, for output states of all the FFT blocks and post-net in the decoder, we insert a fully-connected layer to project the intermediate states to the target feature space and calculate the L1 loss between the projected feature and the target acoustic feature. Besides, we observe that re-weighting the L1 loss by multiplying 1.2 in the dimension of log-scale F0 can achieve better pitch prediction without harming pronunciation. Normally, given the sequence length of the decoder’s output as  $T$ , the number of all decoder blocks as  $B$ , the target LPC feature sequence as  $Y$ , the 26-dim weighting vector as  $w = [1, 1, \dots, 1.2, 1]$ , and the transformation functions as  $\mathcal{F}$ , then the formula for the decoder’s loss function can be written as follows:

$$\mathcal{L}_{decoder} = \frac{1}{T} \sum_{i=1}^T \frac{1}{B} \sum_{j=1}^B w \circ |\mathcal{F}_j(state_{i,j}) - Y_i| \quad (2)$$

#### 2.1.4. Neural Vocoder

Despite the recent popularity of parallel GAN-based vocoders [9], we still find some unsatisfactory artifacts in the quality of synthesized singing voices [16, 17, 19] especially when the human ear is sensitive to the coherence of sound. Different from ByteSinger [20] which adopts WaveRNN [7] as a neural vocoder without modeling the pitch information explicitly, we design a pitch-aware neural vocoder to synthesize singing waveforms with a sampling rate of 24 kHz. Our pitch-aware neural vocoder is modified from the LPCNet [8], which was originally proposed as a neural vocoder for TTS tasks to synthesize 16 kHz speech in an auto-regressive manner. Considering the widely held view that a lower sampling rate as 16 kHz is far away from representing singing voices more accurately, we expand the 18-dimensional BFCCs in the original LPCNet to 24-dimensional BFCCs and combine 24-dimensional BFCCs with two pitch parameters (log-scale F0 and pitch correlation) as input features for our proposed 24 kHz pitch-aware neural vocoder. In the meantime, we increase the dimension in the hidden state of  $GRU_A$  from the original 384 to 512 and keep that of  $GRU_B$  unchanged. To the best of our knowledge, this is the first attempt to adopt an auto-regressive pitch-aware neural vocoder for the Chinese SVS tasks.

## 2.2. Data Augmentation

It is challenging to obtain an exhaustive singing dataset that sufficiently covers the diversity of rhythms, lyrics, and melodies. Furthermore, it is a costly and laborious task to gather correct

annotations. In WeSinger, we try to alleviate such a dilemma by an effective data augmentation method of variable-duration segmentation (VS). Since the acoustic model of WeSinger is based on the Transformer architecture with MHSA having a global receptive field, in which the modeling ability is inevitably sensitive to the length of the input sequence. Under this perspective, we segment each song recording into smaller fragments of three different time intervals: 0 ~ 5 seconds, 5 ~ 8 seconds, and 8 ~ 12 seconds. That is to say, each song recording would be split into short audio clips three times. By the way, we also tried the other data augmentation method such as pitch shifting. Specifically, the overall pitch of each song is raised by one semitone and lowered by one semitone in turn. A similar approach has been proposed in [11]. This allows us to obtain two additional distinct variants based on each training audio clip. However, we find that doing so could lead to somewhat perceptible changes in timbre. Based on the existing VS method, we found no significant benefit from adopting the additional pitch-shifting method. Therefore, in this paper, we only present the performance of adopting the VS method.

## 3. Experiments

### 3.1. Experimental Setup

#### 3.1.1. Dataset

We collect three datasets for SVS experiments totally:

1. a 30-hour singing dataset collected from nearly 160 amateur singers in different noisy environments.
2. Opencpop [18], a recent public Mandarin singing corpus including a pre-defined training set and test set.
3. a 5-hour internal high-quality Mandarin singing dataset collected from a professional female singer. We pick 5 songs as the test set for our experiments and it does not overlap with all other training sets.

All songs are down-sampled to 24 kHz with 16-bit quantization. All 26-dim acoustic features are re-scaled by min-max normalization for acoustic modeling and then de-normalized for training the LPCNet vocoder. Normally, each song is split into singing segments by the pause, these segments last from 1 second to 5 seconds (4 seconds on average). As for the VS-augmented data, each song is split into different intervals as described in Section 2.2.

#### 3.1.2. Experiments

Our experiments are conducted in two groups. One group is to explore the optimal performance following different training recipes with only the data of the professional female singer. Based on the conclusion drawn from the first group, we further perform the other group to enhance performance with the help of multi-singer pre-training and data augmentation. In detail, we design several experiments under different conditions as below:

- **WeSinger** Training the proposed WeSinger system with multi-scale duration loss and progressive pitch-weighted decoder loss.
- **WS- $w/o$ -syllable** Training WeSinger without auxiliary syllable loss for duration model.
- **WS- $w/o$ -weighted** Training WeSinger without weighted loss on pitch dimension.

Table 2: Duration Accuracy of with and without auxiliary syllable loss.

System	Dur Acc	Dur CORR
WS- <i>w/o</i> -syllable	87%	0.96
WeSinger	88%	0.97

Table 3: Quantitative performance of different SVS systems on an internal professional female singer corpus.

System	F0 RMSE	F0 CORR	V/UV error	BFCCD
WS- <i>w/o</i> -weighted	14.5	0.98	0.06	64.2
WS- <i>w/o</i> -progressive	15.3	0.97	0.06	66.7
WeSinger	13.9	0.98	0.06	63.8
WS- <i>w/-VS</i>	14.1	0.98	0.05	63.5
WS-Finetune	13.6	0.98	0.05	63.3
WS- <i>w/-VS</i> -Finetune	12.9	0.98	0.05	63.7

- **WS-*w/o*-progressive** Training WeSinger without progressive loss.
- **WS-*w/-VS*** Training WeSinger with additional VS-augmented data.
- **WS-Finetune** Pre-training WeSinger on multi-singer data and then do fine-tuning.
- **WS-*w/-VS*-Finetune** Pre-training WeSinger on multi-singer data with additional VS-augmented data and then do fine-tuning.

To compare the performance of WeSinger with CpopSing in [18], we further conduct two experiments following the above training recipes of WeSinger and WS-*w/-VS*-Finetune based on the public corpus Openpop.

### 3.2. Evaluation

In both quantitative and qualitative tests, we keep the rhythm and lyrics of the test set consistent among different models to examine the audio quality. Notably, the generated short singing voices are combined into each submitted singing segment for evaluation.

#### 3.2.1. Quantitative evaluation

We conduct several quantitative tests to compare different systems. First, we try to verify the effectiveness of using multi-scale duration loss. As shown in Table 2, when the auxiliary syllable duration loss is adopted, the duration accuracy has been improved from 87% to 88%. Then, to measure the quality of synthesized singing voices, we adopt F0 RMSE (F0 root mean square error), F0 CORR (F0 correlation), and V/UV error (voice/unvoiced error rate), and BFCCD (BFCC distortion) as objective metrics to calculate. For brevity, the ground-truth durations are used as the target duration for systems to synthesize. The calculated results are listed in Table 3 and we can come to the following conclusions: 1) With the progressive pitch-weighted loss, the F0 RMSE can be reduced from 15.3 to 13.9; 2) With VS-augmented data, the quality of predicted BFCC can be slightly improved; 3) Combining VS-augmented data with multi-singer pre-training, the optimal F0 RMSE of 12.9 can be obtained. Overall, data-augmented WeSinger with

Table 4: MOS ratings for the internal professional female singer corpus with the confidence interval 95%.

System	MOS	
	predicted duration	ground-truth duration
WeSinger	2.91 ± 0.11	3.02 ± 0.08
WS- <i>w/-VS</i>	3.18 ± 0.08	3.42 ± 0.08
WS- <i>w/-Finetune</i>	3.40 ± 0.09	3.46 ± 0.06
WS- <i>w/-VS</i> -Finetune	3.48 ± 0.06	3.62 ± 0.05
Ground Truth	4.12 ± 0.04	

Table 5: MOS ratings for the public corpus Openpop with the confidence interval 95%.

System	MOS	
	predicted duration	ground-truth duration
CpopSing	3.20 ± 0.10	3.35 ± 0.15
WeSinger	3.25 ± 0.08	3.45 ± 0.12
WS- <i>w/-VS</i> -Finetune	3.35 ± 0.06	3.60 ± 0.10
Ground Truth	4.01 ± 0.15	

auxiliary losses can achieve better performance in quantitative metrics.

#### 3.2.2. Qualitative evaluation

To compare the performances of different systems qualitatively, we conduct the Mean Opinion Score (MOS) evaluations for naturalness and audio quality on the test set. We collected synthesized samples from different systems based on both ground-truth duration and predicted duration. Twenty listeners were asked to rate the quality of each singing voice segment on a scale from 1 to 5, in which 1 means very bad and 5 means excellent. As shown in Table 4, when training from scratch, a significant gain of 0.3 in MOS is obtained with the VS-augmented data. With multi-singer pre-training, we also benefit from the VS-augmented data and achieve the 3.48 MOS finally. As for the public Openpop corpus, WeSinger also outperforms CpopSing with a MOS margin of 0.1, and WS-*w/-VS*-Finetune achieves a final 3.60 MOS. The detailed evaluation result between WeSinger and CpopSing is listed in Table 5, which indicates that WeSinger has achieved state-of-the-art performance on the public corpus Openpop.

## 4. Conclusion

In this work, we introduced an end-to-end SVS system named WeSinger to bridge the gap between accuracy and naturalness. We described the specifically designed auxiliary loss functions including multi-scale rhythm loss and progressive pitch-weighted decoder loss. Besides, we adopted a data augmentation method of variable-duration segmentation to improve the performance. Experiment results show that the data-augmented WeSinger system can synthesize singing voices with higher quality and naturalness and achieve state-of-the-art performance on the public corpus Openpop. For future work, we will investigate how to improve the robustness efficiently and explore how to adapt to any unseen singer with lower cost.

## 5. References

- [1] M. W. Macon, L. Jensen-Link, J. Oliverio, M. A. Clements, and E. B. George, "A singing voice synthesis system based on sinusoidal modeling," in *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 1. IEEE, 1997, pp. 435–438.
- [2] J. Bonada, M. Umberto Morist, and M. Blaauw, "Expressive singing synthesis based on unit selection for the singing synthesis challenge 2016," *Morgan N, editor. Interspeech 2016; 2016 Sep 8-12; San Francisco, CA.[place unknown]: ISCA; 2016. p. 1230-4., 2016.*
- [3] K. Oura, A. Mase, T. Yamada, S. Muto, Y. Nankaku, and K. Tokuda, "Recent development of the hmm-based singing voice synthesis system—sinsy," in *Seventh ISCA Workshop on Speech Synthesis*, 2010.
- [4] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [5] 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," *arXiv preprint arXiv:1703.10135*, 2017.
- [6] 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 *NeurIPS 2019*, November 2019.
- [7] 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*. PMLR, 2018, pp. 2410–2419.
- [8] J.-M. Valin and J. Skoglund, "Lpcnet: Improving neural speech synthesis through linear prediction," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5891–5895.
- [9] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in Neural Information Processing Systems*, vol. 33, pp. 17 022–17 033, 2020.
- [10] 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.
- [11] 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.
- [12] K. Nakamura, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Singing voice synthesis based on convolutional neural networks," *arXiv preprint arXiv:1904.06868*, 2019.
- [13] Y. Hono, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Sinsy: A deep neural network-based singing voice synthesis system," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2803–2815, 2021.
- [14] P. Lu, J. Wu, J. Luan, X. Tan, and L. Zhou, "Xiaoicesing: A high-quality and integrated singing voice synthesis system," *arXiv preprint arXiv:2006.06261*, 2020.
- [15] 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.
- [16] J. Chen, X. Tan, J. Luan, T. Qin, and T.-Y. Liu, "Hifisinger: Towards high-fidelity neural singing voice synthesis," *arXiv preprint arXiv:2009.01776*, 2020.
- [17] G.-H. Lee, T.-W. Kim, H. Bae, M.-J. Lee, Y.-I. Kim, and H.-Y. Cho, "N-singer: A non-autoregressive korean singing voice synthesis system for pronunciation enhancement," *arXiv preprint arXiv:2106.15205*, 2021.
- [18] Y. Wang, X. Wang, P. Zhu, J. Wu, H. Li, H. Xue, Y. Zhang, L. Xie, and M. Bi, "Opencpop: A high-quality open source chinese popular song corpus for singing voice synthesis," *arXiv preprint arXiv:2201.07429*, 2022.
- [19] R. Huang, F. Chen, Y. Ren, J. Liu, C. Cui, and Z. Zhao, "Multi-singer: Fast multi-singer singing voice vocoder with a large-scale corpus," in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 3945–3954.
- [20] Y. Gu, X. Yin, Y. Rao, Y. Wan, B. Tang, Y. Zhang, J. Chen, Y. Wang, and Z. Ma, "Bytesing: A chinese singing voice synthesis system using duration allocated encoder-decoder acoustic models and wavernn vocoders," in *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2021, pp. 1–5.
- [21] Y. Ren, X. Tan, T. Qin, J. Luan, Z. Zhao, and T.-Y. Liu, "Deepsinger: Singing voice synthesis with data mined from the web," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1979–1989.
- [22] M. Good *et al.*, "Musicxml: An internet-friendly format for sheet music," in *Xml conference and expo*. Citeseer, 2001, pp. 03–04.
- [23] M. M. Association *et al.*, "The complete midi 1.0 detailed specification," *Los Angeles, CA, The MIDI Manufacturers Association*, 1996.
- [24] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [25] 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," *arXiv preprint arXiv:2005.08100*, 2020.
- [26] 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.
- [27] 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.
- [28] Y. Zheng, Z. Zhang, X. Li, W. Su, and L. Lu, "Zero-shot cross-lingual transfer using multi-stream encoder and efficient speaker representation," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 8027–8031.