

CONTUNER: Singing Voice Beautifying with Pitch and Expressiveness Condition

Jianzong Wang¹, Pengcheng Li^{1,2}, Xulong Zhang^{1✉}, Ning Cheng¹, Jing Xiao¹

¹Ping An Technology (Shenzhen) Co., Ltd.

²University of Science and Technology of China

jzwang@188.com, lipengcheng@ustc.edu, zhangxulong@ieee.org,
{chengning211, xiaojing661}@pingan.com.cn

Abstract—Singing voice beautifying is a novel task that has application value in people’s daily life, aiming to correct the pitch of the singing voice and improve the expressiveness without changing the original timbre and content. Existing methods rely on paired data or only concentrate on the correction of pitch. However, professional songs and amateur songs from the same person are hard to obtain, and singing voice beautifying doesn’t only contain pitch correction but other aspects like emotion and rhythm. Since we propose a fast and high-fidelity singing voice beautifying system called CONTUNER, a diffusion model combined with the modified condition to generate the beautified Mel-spectrogram, where the modified condition is composed of optimized pitch and expressiveness. For pitch correction, we establish a mapping relationship from MIDI, spectrum envelope to pitch. To make amateur singing more expressive, we propose the expressiveness enhancer in the latent space to convert amateur vocal tone to professional. CONTUNER achieves a satisfactory beautification effect on both Mandarin and English songs. Ablation study demonstrates that the expressiveness enhancer and generator-based accelerate method in CONTUNER are effective.

Index Terms—singing voice beautifying, speech representation learning

I. INTRODUCTION

Singing voice beautifying (SVB) is a novel task in the field of speech and music. SVB models aim to calibrate the pitch of the amateur singing voice while retaining the timbre and content of the original singing voice and then improve the singing skills and expressiveness of the singing voice. In the entertainment industry, SVB is often completed by vocal tuners using professional tools like Auto-Tune [1], which requires expensive labour costs. Since many people enjoy singing but face difficulties to obtain satisfactory songs, automatic singing beautification has great application value in our daily lives.

There are currently two types of tasks associated with SVB, singing voice conversion (SVC) and automatic pitch correction (APC). SVC changes the singer of the source singing voice and maintains the content. But in SVB, we need to keep the content and the vocal timbre. APC directly corrects the pitch of amateur singing voices, but this is insufficient to perform beautification. Current SVB work [2] relies on paired data, but it is difficult to obtain amateur and professional songs from the same person. The data we can easily obtain are amateur singing voices and original professional recordings of

the same song. Lots of improvements in previous works are made only on pitch correction, but few researchers focus on the expressiveness of a song like singing skill, emotion, rhythm, *etc.* How to generate a new song with a newly generated pitch and beautified features is also a critical problem. One way is to use a vocoder such as WORLD [3], but this kind of generation work is difficult to establish the control of expressiveness and other features of the song.

To address the mentioned problems, we propose a diffusion-based SVB model, CONTUNER, which integrates the predicted pitch information as well as the modified expressiveness into the original amateur singing voice to achieve singing voice beautification. As far as we know, we are the first to model the expressiveness in singing voices in the SVB area, and we hope to bring inspiration to future works. Furthermore, as an outstanding SVB model is expected to be fast, we choose the high-performance generative model *Diffusion* [4] as the backbone of CONTUNER. And inspired by some recent works, we speed up the diffusion model through using *generator-based* methods [5], [6]. During the whole training and reasoning process of the model, CONTUNER only needs to extract target information from amateur singing voice constraints, thereby avoiding the problem of paired data. In addition, CONTUNER beautifies the amateur singing voice from the perspective of pitch and expressiveness. Moreover, CONTUNER chooses to establish control of the condition during the generation of the beautified Mel-spectrogram instead of establishing control of the condition in the vocoder stage. To sum up, our contributions can be summarized as follows:

- We propose a novel model called CONTUNER to solve the task of beautification *without* professional-amateur paired data from the same singer. We provide a new perspective that obtains conditions from pitch and expressiveness and establishes the control of the condition in the process of Mel-spectrogram generation.
- We *predict* pitch instead of fitting pitch in terms of pitch correction. We establish a mapping relationship from MIDI, spectral envelope to pitch curve to get the corrected pitch curve. We derive an expressiveness representation from Mel-spectrogram and obtain the beautified expressiveness via the designed expressiveness enhancer.

✉ Corresponding author.

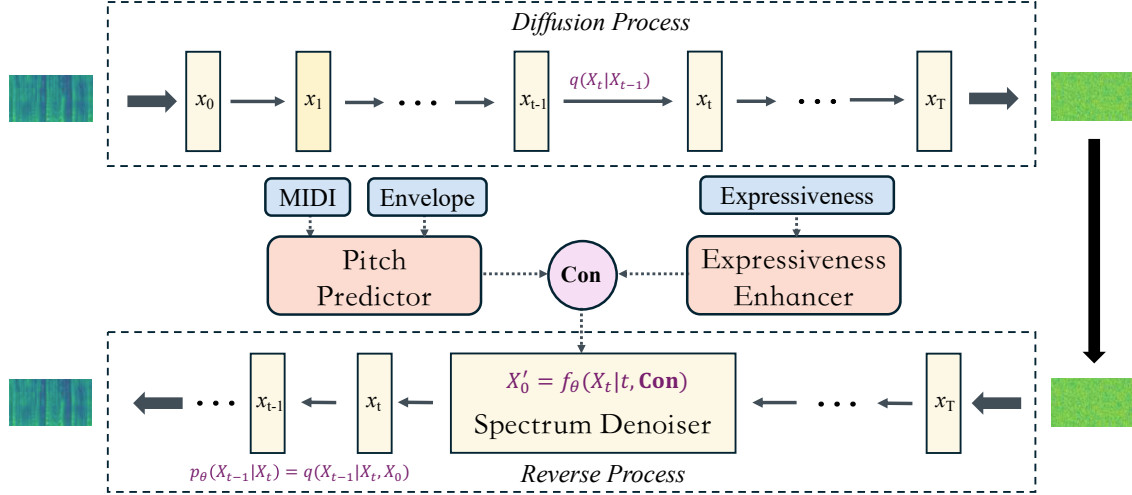


Fig. 1. Architecture of CONTUNER. The pitch predictor conducts mapping from MIDI and envelope to pitch, while the expressiveness enhancer disentangles the expressiveness representation from the singing voice. The outputs from them are combined as the condition that takes part in the denoising process.

II. RELATED WORKS

A. Neural Singing Voice Processing

Utilizing neural networks for singing voice processing is an area that has begun to attract attention in recent years [7], [8]. Different from speech voice processing (like speech voice conversion [9], [10], text-to-speech [11], [12], *etc.*), singing voice contains richer and more varied pitch and rhythmic information, which requires more accurate control and optimization of these information. Existing works on neural singing voice processing contain automatic pitch correction [13]–[15], singing voice conversion [16], [17], and singing voice beautifying [2]. APC attempts to modify the unsatisfactory pitch of singing voices from amateur singers but APC doesn't further consider the beautification of the processed singing voice, such as singing skills, emotion, expressiveness, *etc.* SVC is a downstream task of voice conversion, which aims to change the singer of a given singing voice. SVC methods usually first disentangle the timbre and the content from the target song and source song respectively, and then combine the timbre information from the target singer with the content information from the source song [18]–[20]. SVC is quite different from SVB, as SVC will erase all the uniqueness of the singing voice while SVB tries to keep the content and the vocal timbre unchanged. Simply leveraging the SVC method for the SVB task will greatly destroy the original singers' unique characteristics.

B. Diffusion Model

Denoising diffusion probabilistic model (DDPM) [4], [21] gradually transforms simple distributions into complex data distributions with Markov chains. It has gained satisfactory performance in diverse tasks like image generation [22], speech synthesis [23], and even object detection [24]. The

forward diffusion process and reverse generation process are built in DDPM to learn the transformation. In the forward process, Gaussian noise is gradually incorporated into the data, this helps the model to explore the various data distributions, while the model tries to denoise to restore the origin data through the reverse process. The reverse process makes the model gain the capability to generate real samples. Compared to the classic generated model, generative adversarial network (GAN), DDPM is more stable but relatively slow. How to improve the inference speed of DDPM also becomes an issue worth exploring [25]–[27].

III. METHODOLOGY

A. Generator- and Condition-based Diffusion

The overall architecture of CONTUNER is shown in Fig. 1. Diffusion model is adapted as the backbone of CONTUNER, combining with modified condition (*i.e.* the output of the pitch predictor and expressiveness enhancer). Through the diffusion process and reserve process, we can get the Mel-spectrogram of the beautified singing voice. The diffusion process is non-parametric. We input the Mel-spectrogram x_0 of amateur song, and after t steps of sampling, we get the Mel-spectrogram x_t with noise. With the pre-defined noise schedule β and diffusion step t , we compute the corresponding constants in Eq. 1 respectively.

$$\alpha_t = \prod_{i=1}^t \sqrt{1 - \beta_i} \quad \sigma_t = \sqrt{1 - \alpha_t^2} \quad (1)$$

In conventional gradient-based training, the noise ϵ loss is calculated through optimizing a random term of t with stochastic gradient descent as shown in Eq. 2.

$$\mathcal{L}_\theta^G = \|\epsilon_\theta(\alpha_t x_0 + \sqrt{1 - \alpha_t^2} \epsilon) - \epsilon\|_2^2, \epsilon \sim \mathcal{N}(0, 1) \quad (2)$$

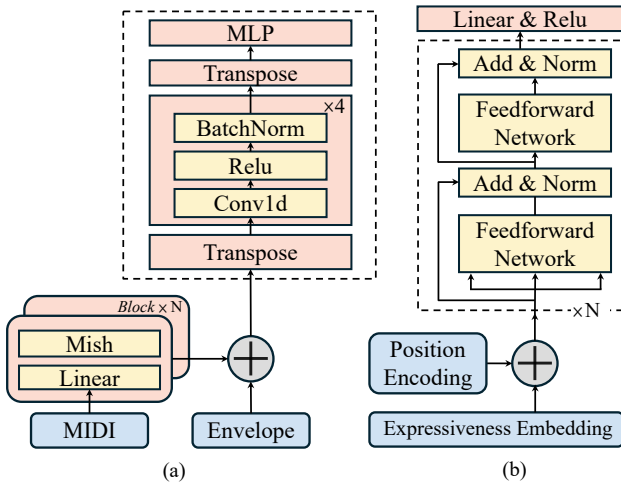


Fig. 2. Details of the pitch predictor and expressiveness enhancer.

It is well known that in DDPMs, x_t has different degrees of perturbation, thus using a single gradient-based parameterization network directly in different t predictions for x_{t-1} is difficult. Inspired by the recent works [5], [28], [29] in the field of text-to-speech, which point out that the generator-based diffusion model does not need to estimate the gradient of the data density. So we predict clean professional Mel-spectrogram x_0^p by undisturbed x_0' , and then add the anti-perturbation through the posterior distribution $q(x_{t-1}|x_t, x_0')$.

Therefore, in the reverse sampling process of the diffusion model, x_{t-1} is sampled with the posterior distribution given x_t and the predicted x_0' .

$$x_{t-1} \sim p_\theta(x_{t-1}|x_t) = q(x_{t-1}|x_t, x_0') \quad (3)$$

In this way, CONTUNER can greatly reduce the quantity of sampling steps t . We put the beautifying condition Con into the denoiser to guide the direction of generation as shown in Eq. 4.

$$x_0' = f_\theta(x_t|t, Con) \quad (4)$$

Finally, we constrain the loss of the denoiser which is defined as a mean squared error (MSE) in the data space x . Efficient training is to optimize a random term of t with stochastic gradient descent.

$$\mathcal{L}_\theta^{Denoiser} = \|x_\theta(\alpha_t x_0 + \sqrt{1 - \alpha_t^2} \epsilon) - x_0^p\|_2^2 \quad (5)$$

B. Pitch Predictor and Expressiveness Enhancer

1) *Pitch Predictor*: The structure of the *pitch predictor* (PP) is shown in Fig. 2(a). As we hope to build the beautified pitch curve and maintain the pitch characteristics of the amateur singer at the same time, so we combine the spectral envelope and MIDI from professional singers and amateur singers respectively. We establish a mapping $pitch \leftarrow (envelope, MIDI)$. The spectral envelope feature implicitly contains the pitch curve [30], as previous work predicts the pitch curve from the spectral envelope with high accuracy.

We feed spectral envelope into the designed PP. Besides, MIDI can be seen as the standard pitch representation of a song. We perform HMM smoothing to extract standard MIDI note sequences from vocals via *pyin* [31]. With the help of the WORLD vocoder [3], we extract the spectral envelope and pitch curve (set as the pitch label p). The MIDI vector goes through 4 linear layers, each layer is followed by a mish activation function, then spliced and flipped with the spectral envelope dimension. The main body of PP contains four Conv1d layers with ReLU activation function and batch normalization followed, finally the predicted pitch curve p' can be obtained from the MLP layer.

During the training process of the model, we input the MIDI and the spectral envelope of amateur singers, while the labels are also extracted from amateur singing voices. It is worth noting that PP does not require any information of professional songs during the training stage. In the inference stage, we input the spectral envelope from professional singers and the MIDI from amateur singers to obtain the beautified pitch curves.

2) *Expressiveness Enhancer*: The structure of *expressiveness enhancer* (EE) is shown in Fig. 2(b). We define the singing skills, rhythm, and emotion of a singing voice as the expressiveness characteristics. We first obtain expressiveness representation through disentanglement similar to voice conversion. Defining a pair of expressiveness representations (W_a, W_p), which represent the amateur expressiveness W_a and professional expressiveness W_p of the same song respectively. We disentangle expressiveness representation with a pre-trained encoder. A feed-forward Transformer block is adapted as the main body of the expressiveness enhancer, which is a stack of self-attention layers [32]. The latent expressiveness representation W_a' represents the expressiveness after being modified by EE. What we need to do is to strengthen the similarity between W_a' and W_p .

3) *Condition Loss*: Condition loss consists of pitch loss and expressiveness loss as shown in Eq. 6, λ_1 and λ_2 are weight coefficients of the two losses. The pitch loss is the L_2 distance between the predicted pitch p' and the pitch label p , while the expressiveness loss is retrieved from the professional and modified amateur expressiveness representation.

$$\mathcal{L}_{con} = \lambda_1 \|p' - p\|_2 + \lambda_2 \|W_a' - W_p\|_2 \quad (6)$$

C. Denoiser

Following the previous work [29], we adopt a non-causal WaveNet [33] architecture as our spectrogram denoiser as shown in Fig. 3. The denoiser comprises a 1×1 convolution layer and N convolution blocks with residual connections to project the input hidden sequence with 256 channels. For any step t , we use a cosine schedule with $\beta_t = \cos(0.5\pi t)$. The condition first passes through the length regulator to get the same dimension to x_t through padding. Finally, the output part of the denoiser consists of a 2-layer 1D-convolutional network with ReLU activation.

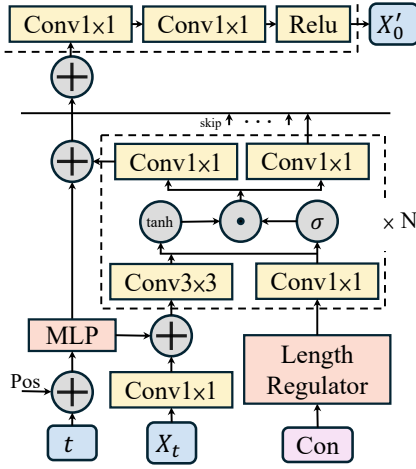


Fig. 3. Structure of the spectrogram denoiser.

D. Training and Inference

1) *Training*: The final loss term in training CONTUNER consists of the following parts:

- The reconstruction loss $\mathcal{L}_{\theta}^{Denoiser}$ in Eq. 5, which is the MSE between the generated and target Mel-spectrograms.
- The condition loss \mathcal{L}_{con} , which is the distance between the predicted pitch and the target pitch, as well as the modified expressiveness and the professional expressiveness, as shown in Eq. 6.

It is worth noting that the pitch predictor, expressiveness enhancer and denoiser perform gradient back-propagation at the same time during training.

2) *Inference*: CONTUNER iteratively predicts unperturbed x_0^p and then adds back perturbation via the posterior distribution. Specifically, the denoising model $f_{\theta}(x_t|t, con)$ predicts x_0' firstly, and then x_{t-1} is sampled using the posterior distribution $q(x_{t-1}|x_t, x_0')$ that given x_t and the predicted x_0' . In addition to the amateur singing voice, we **only** need a professional MIDI in the inference stage.

IV. EXPERIMENT

A. Experimental Setup

Since SVB is a novel task, with few public unaccompanied datasets so far. NSVB [2] comes up with PopButFy, an SVB dataset composed of paired data (*i.e.* amateur and professional vocals from the same person). However, we aim to reduce the SVB model's dependency on paired data, so we collect professionally recorded songs and songs from amateur singers to produce an SVB dataset called Professional and Amateur Singing Voice dataset **PASV**. PASV consists of about 400 Mandarin and English pop songs (≈ 42 hours) in total. In order to get closer to the scene of SVB, the amateur songs in PASV contain out-of-tune samples and extremely amateur samples. For each amateur sample, there is a one-to-one correspondence

with the original professional song. For the recorded song, we use Spleeter [34] to separate the singing voice from the accompaniment to extract the pure human singing voice.

We utilize the Griffin-Lim algorithm [35] as the vocoder to obtain waveform from the generated Mel-spectrogram in all our experiments. CONTUNER is trained on a 12G NVIDIA 3080Ti GPU with 400k steps. The warm-up learning rate is set to 10^{-4} and an Adam optimizer [36] with $\beta_1 = 0.9, \beta_2 = 0.99, \epsilon = 10^{-9}$ is built for training. Audio samples are available at <https://largeaudiomodel.com/contuner/>.

B. Metrics and Measurement

We use subjective metric mean opinion score (MOS, and comparison MOS, **CMOS** for ablation study) and the objective metric Mel-cestral distortion (MCD) to evaluate the performance of our proposed model on the test set. In addition, following NSVB [2], we also leverage pitch alignment accuracy (PAA) as an objective metric to measure pitch correction. For the beautified singing voices, we analyzed the MOS from two aspects: audio quality (such as naturalness, singing voice quality, *etc.*, denoted as **MOS-Q**) and expressiveness (such as singing skills, emotion, rhythm, *etc.*, denoted as **MOS-E**).

C. Experimental Results

1) *Pitch Correction*: Fig. 4 shows the results of the comparison between our proposed method and other methods on PAA. Dynamic Time Warping (DTW) [37] and Canonical Time Warping (CTW) [13] are two classic algorithms for pitch correction [14], while KaraTuner [15] is a Transformer-based method that performs pitch correction. CONTUNER(P) represents the pitch predictor in our proposed model. Results show that PP in CONTUNER outperforms other time-aligned methods. This is mainly because previous time-warping algorithms only focus on the forced alignment in time but ignore the direction of the pitch curve.




Fig. 4. The pitch alignment accuracy of different algorithms on Mandarin and English songs.

2) *Quality and Expressiveness*: In order to analyze the beautification capability of our proposed model, we further compare the quality MOS-Q, expressive quality MOS-E and objective metric MCD with other baseline methods. The following models and ground-truths are compared,

TABLE I
COMPARISON OF DIFFERENT METHODS
(WITH 95% CONFIDENCE INTERVALS).

Lang.	Method	MOS-Q \uparrow	MOS-E \uparrow	MCD \downarrow
Mandarin	GT MelA	4.31 \pm 0.21	3.16 \pm 0.28	-
	GT MelP	4.42 \pm 0.18	4.39 \pm 0.14	-
	KaraTuner [15]	4.15 \pm 0.15	4.02 \pm 0.13	7.21 \pm 0.09
	CONTUNER	4.21 \pm 0.13	4.24 \pm 0.12	6.97 \pm 0.14
English	GT MelA	4.18 \pm 0.25	3.03 \pm 0.13	-
	GT MelP	4.24 \pm 0.11	4.25 \pm 0.11	-
	KaraTuner [15]	4.01 \pm 0.07	3.86 \pm 0.13	8.81 \pm 0.11
	CONTUNER	4.06 \pm 0.15	4.03 \pm 0.12	7.20 \pm 0.10

- Ground-truth Mel includes amateur (GT MelA) and professional (GT MelP) versions. We first convert ground-truth audio into Mel-spectrogram, and then convert the Mel-spectrogram back to waveform via the vocoder to eliminate the effect of vocoder on evaluation.
- KaraTuner [15], a neural pitch correlation model.
- CONTUNER, the proposed SVB model.

Table I shows the results of CONTUNER and other models on both Mandarin and English singing voice data. Results show that

- CONTUNER significantly achieves promising results in both MOS-Q and MOS-E, with audio quality degradation with 0.12, as well as MOS-E being more than those for ground-truth amateur recordings by 1.08 and 1.00 on Mandarin and English data respectively. This result proves the strong performance of the CONTUNER in singing voice beautification.
- As for MOS-E, CONTUNER is less than those for ground-truth professional recordings by only 0.15 and 0.21 in Mandarin and English singing voice data respectively, which proves that the CONTUNER has strong language generalization performance.

Comparing to other methods, CONTUNER achieves satisfactory results. It shows that under the control of the condition, the Mel-spectrogram will be optimized with a more accurate pitch and better expressiveness.

D. Ablation Study

The number of sampling steps is the key parameter of the diffusion model. As the number of sampling steps increases, the quality of the audio generated will also improve but results in more time-consuming. Therefore, how to balance sampling steps and audio quality is an important problem. In order to explore the influence of different sampling steps on the performance of CONTUNER, an ablation experiment is conducted. We also compare the gradient-based diffusion and our generator-based diffusion. The comparison of the effect of sampling steps is shown in Table II.

As Table II shows, with several sample steps and a large distribution of noise schedule, gradient-based or generator-based diffusion models could produce high-fidelity speech samples with similar results. When the amount of sampling

TABLE II
COMPARISON OF THE EFFECT OF SAMPLING STEPS ON MOS-Q \uparrow OF
GRADIENT- AND GENERATOR-BASED DIFFUSION MODEL
(WITH 95% CONFIDENCE INTERVAL).

Sampling Steps	Gradient-Based	Generator-Based
200	4.23 \pm 0.11	4.24 \pm 0.09
100	4.11 \pm 0.06	4.22 \pm 0.07
50	4.02 \pm 0.10	4.22 \pm 0.06
10	3.65 \pm 0.08	4.21 \pm 0.13
GT MelA	4.31 \pm 0.21	4.31 \pm 0.21
GT MelP	4.42 \pm 0.18	4.42 \pm 0.18

steps gradually decreases to 10, the audio quality of generator-based diffusion does not decrease significantly. This indicates that CONTUNER can weaken the inherent trade-off problem of the diffusion model to a certain extent. We hold the view that the generator-based diffusion model is free from estimating the gradient for data density, which only needs to predict unperturbed x_0 and then add back perturbation using the posterior distribution. So our generator-based diffusion model can achieve a nice beautification effect without hundreds of sample steps.

Furthermore, in order to verify the specific effects of the designed EE in the singing beautification task, we compare the performance of CONTUNER with CONTUNER(E) which lacks the expressive enhancer. Comparison mean opinion score (CMOS-E and CMOS-Q) is obtained for the comparison. It can be seen from Table III that the EE mainly improves MOS-E and has no significant impact on audio quality, which indicates that the EE can improve the expressiveness of singing voices, achieving the purpose of singing voice beautifying.

TABLE III
COMPARISON OF CONTUNER(E) AND CONTUNER.

Language	Method	CMOS-Q	CMOS-E
Mandarin	CONTUNER(E)	-0.003	-0.140
	CONTUNER	0.000	0.000
English	CONTUNER(E)	-0.000	-0.230
	CONTUNER	0.000	0.000

E. Discussion

Sing voice beautifying is currently in its early stages. In this section, We analyze some limitations in our work as well as something that can be done in the future. We hope that there will be more works to promote the development of SVB.

In this paper, we only consider the beautification of singing voices in a single scene. In the future, new beautification scenarios will bring new problems. We decouple expressiveness features via a pre-trained encoder for expressiveness enhancement, this introduces some limitations to CONTUNER. On the other hand, The establishment of the control of the condition makes the consideration of more factors like singing skills and emotions for SVB possible. This could be useful for the quest for SVB in the future works.

V. CONCLUSION

In this paper, we propose CONTUNER, a fast diffusion-based model for high-fidelity singing voice beautifying with pitch and expressiveness conditions. The proposed model does not need paired data for inference, it can establish the control of the condition in the process of spectrogram generation. With fewer sampling steps, our model achieves a beneficial effect on singing voice beautification. Experimental results show that CONTUNER gets satisfactory performance on PAA and outperforms the baseline method with higher quality and better expressiveness in both Mandarin and English singing voice samples.

VI. ACKNOWLEDGEMENT

Supported by the Key Research and Development Program of Guangdong Province (grant No. 2021B0101400003) and Corresponding author is Xulong Zhang (zhangxulong@ieee.org).

REFERENCES

- [1] S. Yong and J. Nam, "Singing expression transfer from one voice to another for a given song," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 151–155.
- [2] J. Liu, C. Li, Y. Ren, Z. Zhu, and Z. Zhao, "Learning the beauty in songs: Neural singing voice beautifier," in *the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 7970–7983.
- [3] 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.
- [4] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 6840–6851.
- [5] S. Liu, D. Su, and D. Yu, "Diffgan-tts: High-fidelity and efficient text-to-speech with denoising diffusion gans," *arXiv preprint arXiv:2201.11972*, 2022.
- [6] T. Salimans and J. Ho, "Progressive distillation for fast sampling of diffusion models," in *the 10th International Conference on Learning Representations*, 2022.
- [7] Y. Sun, X. Zhang, X. Chen, Y. Yu, and W. Li, "Investigation of singing voice separation for singing voice detection in polyphonic music," in *the 9th Conference on Sound and Music Technology*, 2023, pp. 79–90.
- [8] X. Zhang, J. Wang, N. Cheng, and J. Xiao, "Susing: Su-net for singing voice synthesis," in *2022 International Joint Conference on Neural Networks*, 2022, pp. 1–7.
- [9] J. Chou and H. Lee, "One-shot voice conversion by separating speaker and content representations with instance normalization," in *the 20th Annual Conference of the International Speech Communication Association*, G. Kubin and Z. Kacic, Eds., 2019, pp. 664–668.
- [10] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "Autovc: Zero-shot voice style transfer with only autoencoder loss," in *the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 97, 2019, pp. 5210–5219.
- [11] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 4779–4783.
- [12] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "FastSpeech 2: Fast and high-quality end-to-end text to speech," in *the 9th International Conference on Learning Representations*, 2021.
- [13] F. Zhou and F. Torre, "Canonical time warping for alignment of human behavior," in *Advances in Neural Information Processing Systems*, vol. 22, 2009.
- [14] Y.-J. Luo, M.-T. Chen, T.-S. Chi, and L. Su, "Singing voice correction using canonical time warping," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 156–160.
- [15] X. Zhuang, H. Yu, W. Zhao, T. Jiang, P. Hu, S. Lui, and W. Zhou, "Karatuner: Towards end to end natural pitch correction for singing voice in karaoke," in *the 23rd Annual Conference of the International Speech Communication Association*, 2022.
- [16] T. Kaneko and H. Kameoka, "CycleGAN-vc: Non-parallel voice conversion using cycle-consistent adversarial networks," in *the 26th European Signal Processing Conference*, 2018, pp. 2100–2104.
- [17] T. Kaneko, H. Kameoka, K. Tanaka, and N. Hojo, "CycleGAN-vc2: Improved cycleGAN-based non-parallel voice conversion," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 6820–6824.
- [18] S. Liu, Y. Cao, D. Su, and H. Meng, "Diffsvc: A diffusion probabilistic model for singing voice conversion," in *IEEE Automatic Speech Recognition and Understanding Workshop*, 2021, pp. 741–748.
- [19] S. Liu, Y. Cao, N. Hu, D. Su, and H. Meng, "Fastsvc: Fast cross-domain singing voice conversion with feature-wise linear modulation," in *International Conference on Multimedia and Expo. IEEE*, 2021, pp. 1–6.
- [20] J. Lu, K. Zhou, B. Sisman, and H. Li, "Vaw-gan for singing voice conversion with non-parallel training data," in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, 2020, pp. 514–519.
- [21] A. Q. Nichol and P. Dhariwal, "Improved denoising diffusion probabilistic models," in *International Conference on Machine Learning*, 2021, pp. 8162–8171.
- [22] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with CLIP latents," *CoRR*, vol. abs/2204.06125, 2022.
- [23] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, "Diffwave: A versatile diffusion model for audio synthesis," in *the 9th International Conference on Learning Representations*, 2021.
- [24] S. Chen, P. Sun, Y. Song, and P. Luo, "DiffusionDet: Diffusion model for object detection," in *the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 19830–19843.
- [25] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, "Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps," in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 5775–5787.
- [26] X. Ma, G. Fang, and X. Wang, "Deepcache: Accelerating diffusion models for free," *arXiv preprint arXiv:2312.00858*, 2023.
- [27] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, "Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models," *arXiv preprint arXiv:2211.01095*, 2022.
- [28] M. W. Y. Lam, J. Wang, D. Su, and D. Yu, "Bddm: Bilateral denoising diffusion models for fast and high-quality speech synthesis," in *the 10th International Conference on Learning Representations*, 2022.
- [29] R. Huang, Z. Zhao, H. Liu, J. Liu, C. Cui, and Y. Ren, "Prodiff: Progressive fast diffusion model for high-quality text-to-speech," in *the 30th ACM International Conference on Multimedia*, 2022, pp. 2595–2605.
- [30] T. En-Najjary, O. Rosec, and T. Chonavel, "A new method for pitch prediction from spectral envelope and its application in voice conversion," in *Annual Conference of the International Speech Communication Association*, 2003.
- [31] M. Mauch and S. Dixon, "Pyin: A fundamental frequency estimator using probabilistic threshold distributions," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2014, pp. 659–663.
- [32] 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*, vol. 30, 2017.
- [33] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *the 9th ISCA Speech Synthesis Workshop*, 2016, p. 125.
- [34] R. Hennequin, A. Khelif, F. Voituret, and M. Moussallam, "Spleeter: a fast and efficient music source separation tool with pre-trained models," *Journal of Open Source Software*, vol. 5, no. 50, p. 2154, 2020.
- [35] Y. Masuyama, K. Yatabe, Y. Koizumi, Y. Oikawa, and N. Harada, "Deep griffin-lim iteration," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 61–65.
- [36] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *the 3rd International Conference on Learning Representations*, 2015.
- [37] M. Müller, "Dynamic time warping," *Information Retrieval for Music and Motion*, pp. 69–84, 2007.