

StyleSinger: Style Transfer for Out-of-Domain Singing Voice Synthesis

Yu Zhang¹, Rongjie Huang¹, Ruiqi Li¹, JinZheng He¹, Yan Xia¹, Feiyang Chen², Xinyu Duan²,
Baoxing Huai², Zhou Zhao^{1*}

¹Zhejiang University

²Huawei Cloud

{yuzhang34,rongjiehuang,ruiqili,jinzhenghe,zhaozhou}@zju.edu.cn, xiayan.zju@gmail.com,

{chenfeiyang2,duanxinyu,huaibaoxing}@huawei.com

Abstract

Style transfer for out-of-domain (OOD) singing voice synthesis (SVS) focuses on generating high-quality singing voices with unseen styles (such as timbre, emotion, pronunciation, and articulation skills) derived from reference singing voice samples. However, the endeavor to model the intricate nuances of singing voice styles is an arduous task, as singing voices possess a remarkable degree of expressiveness. Moreover, existing SVS methods encounter a decline in the quality of synthesized singing voices in OOD scenarios, as they rest upon the assumption that the target vocal attributes are discernible during the training phase. To overcome these challenges, we propose StyleSinger, the first singing voice synthesis model for zero-shot style transfer of out-of-domain reference singing voice samples. StyleSinger incorporates two critical approaches for enhanced effectiveness: 1) the Residual Style Adaptor (RSA) which employs a residual quantization module to capture diverse style characteristics in singing voices, and 2) the Uncertainty Modeling Layer Normalization (UMLN) to perturb the style attributes within the content representation during the training phase and thus improve the model generalization. Our extensive evaluations in zero-shot style transfer undeniably establish that StyleSinger outperforms baseline models in both audio quality and similarity to the reference singing voice samples. Access to singing voice samples can be found at <https://stylesinger.github.io/>.

Introduction

Singing voice synthesis (SVS) is dedicated to generating high-quality singing voices through the utilization of lyrics and musical notations. This domain has witnessed significant advancements, finding crucial applications in both the realm of professional music composition and entertainment short videos. Currently, numerous outstanding SVS techniques demonstrate remarkable efficacy in synthesizing exceptional results (Zhang et al. 2022b; Choi and Nam 2022; Kim et al. 2023; Huang et al. 2022a, 2021; He et al. 2023).

With the rapid development of SVS methods, there is a growing demand for out-of-domain (OOD) style transfer in singing voices, which seeks to generate high-quality singing voices with unseen styles derived from reference singing voice samples. To be more specific, styles of singing voices

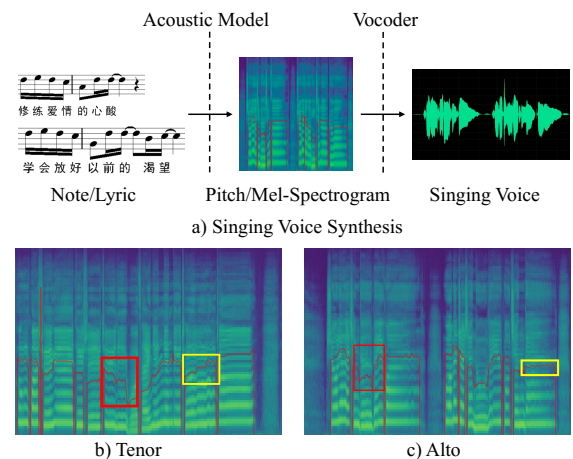


Figure 1: Figure (a) shows the singing voice synthesis overall pipeline. SVS systems use an acoustic model to transform musical notations and lyrics into intermediate features (like pitch and mel-spectrograms), and then a vocoder synthesizes the target singing voices. In this paper, our method mainly focuses on the acoustic model. Figures (b) and (c) depict the constituent elements of singing voice styles, namely pronunciation and articulation skills. Red boxed showcases pitch transitions and yellow boxes highlight the vibrato skill.

primarily include **timbre, emotion, pronunciation, and articulation skills**. Timbre represents the fundamental and distinctive quality of a singer’s voice, while emotion captures the expressive and emotional delivery conveyed during a performance. As shown in Figure 1(b) and (c), pronunciation and articulation skills involve various techniques such as vibrato, pitch transitions, and enunciation skills. However, current SVS systems lack the necessary techniques to effectively model the intricate styles of singing voices. Consequently, existing SVS methods encounter a decline in the quality of synthesized samples in OOD scenarios, as they rest upon the assumption that the target vocal attributes are discernible during the training phase.

In essence, the challenges of the style transfer for OOD SVS can be summarized as follows: 1) Modeling the intricate nuances of singing voice styles is an arduous task, as

*Corresponding author

singing voices possess a remarkable degree of expressiveness. Some methods for style modeling utilize a speaker encoder (Kumar et al. 2021). Moreover, other approaches model styles from multiple perspectives (Casanova et al. 2022). However, these methods only consider limited aspects of speech styles and do not model the detailed styles of singing voices, such as pronunciation and articulation skills. 2) Disparities between the styles of OOD reference samples and the training data often lead to a deterioration in the quality of the synthesized singing voices. Many methods for model generalization rely on extensive data (Jia et al. 2018), which will be costly for singing voices. Alternatively, some methods employ a style adaptor for unseen styles (Min et al. 2021), but they often require direct access to the target voice for model adaptation, which is not always feasible.

To tackle these challenges, we propose StyleSinger, the first singing voice synthesis (SVS) model for zero-shot style transfer of out-of-domain (OOD) reference samples. To capture the diverse style information in singing voices, we introduce the Residual Style Adaptor (RSA). The RSA employs a residual quantization module to capture detailed style characteristics (e.g., pronunciation and articulation skills) in reference samples. To improve the model generalization, we propose the Uncertainty Modeling Layer Normalization (UMLN). The UMLN perturbs the style attributes within the content representation during the training phase, so the model performs better when faced with unseen reference styles during testing. Our comprehensive evaluations in zero-shot style transfer establish that StyleSinger surpasses the baseline models in singing quality and similarity to the reference style. The main contributions of this work are summarized as follows:

- We present StyleSinger, the first singing voice synthesis model for zero-shot style transfer of out-of-domain reference samples. StyleSinger excels in generating exceptional singing voices with unseen styles derived from reference singing voice samples.
- We propose the Residual Style Adaptor (RSA), which uses a residual quantization model to meticulously capture diverse style characteristics in reference samples.
- We introduce the Uncertainty Modeling Layer Normalization (UMLN) to perturb the style information in the content representation during the training phase, and thus enhance the model generalization of StyleSinger.
- Extensive experiments in zero-shot style transfer show that StyleSinger exhibits superior audio quality and similarity compared with baseline models.

Related Works

Singing Voice Synthesis

Singing voice synthesis (SVS) aims to generate singing voices of exceptional quality based on provided musical scores and lyrics. DiffSinger (Liu et al. 2022) introduces a diffusion decoder (Ho, Jain, and Abbeel 2020) to generate high-fidelity mel-spectrograms. In the multi-singer scenarios, MuSE-SVS (Kim et al. 2023) presents a multi-singer emotional singing voice synthesizer. M4Singer (Zhang et al.

2022a) releases a multi-style, multi-singer Chinese song corpus with meticulously annotated fine-grained music scores. Wesinger (Zhang et al. 2022c) proposes a Transformer-like acoustic model. Recently, RMSSinger (He et al. 2023) proposes a method based on realistic music scores, utilizing a diffusion pitch prediction model to forecast F0 and UV. However, these SVS methods encounter challenges in maintaining synthesis quality when dealing with out-of-domain singers and styles, as well as in accurately modeling the intricate nuances of singing voice styles. In this paper, our approach successfully tackles these difficulties.

Style Modeling

The field of audio research has dedicated significant efforts to the exploration of style modeling. Attention (Choi et al. 2020) introduces an attention mechanism to extract styles from reference samples. Cooper et al. (2020) proposes a speaker embedding method to model the reference samples. ZSM-SS (Kumar et al. 2021) proposes a Transformer-based architecture with an external speaker encoder using wav2vec 2.0 (Baeovski et al. 2020). Moreover, numerous methods focus on modeling multi-level audio styles apart from speaker embedding. (Li et al. 2021) incorporates global utterance-level and local phoneme-level style features in target speech. SC-GlowTTS (Casanova et al. 2021) presents a speaker-conditional architecture utilizing flow-based models. Meta-StyleSpeech (Min et al. 2021) employs a speech encoding network for synthesizing multi-speaker TTS. Styler (Lee, Park, and Kim 2021) disentangles style factors with equal supervision levels. Generspeech (Huang et al. 2022b) incorporates both global and local style adaptors to capture styles. However, these approaches focus on limited aspects of speech styles and fail to capture the pronunciation and articulation skills of singing voice styles.

Model Generalization

Enabling the model to effectively capture the essence of unfamiliar out-of-domain test data presents a formidable challenge that SVS models must confront. Prominent methodologies (Jia et al. 2018; Paul, Pantazis, and Stylianou 2020) leverage extensive data to achieve generalization. When it comes to singing voice data, acquiring a substantial amount of annotated data proves to be both costly and arduous. Min et al. (2021); Huang et al. (2022d) employ meta-learning as the style adaptor for unseen speakers not encountered during the training phase. Such style adaptation methods require accessibility to the target voice, which is not always feasible. In contrast, Casanova et al. (2022) have devised an architecture that builds upon VITS, yielding exceptional zero-shot results. In the image domain, certain approaches focus on manipulating feature statistics to improve model generalization. MixStyle (Zhou et al. 2021) utilizes linear interpolation on feature statistics and shuffles the input samples to generate synthesized samples. Similarly, pAdaIn (Nurriel, Benaim, and Wolf 2021) applies a random permutation to swap sample statistics. Nevertheless, all of these approaches primarily concentrate on the domains of speech or image, whereas our focus is on the realm of singing voices.

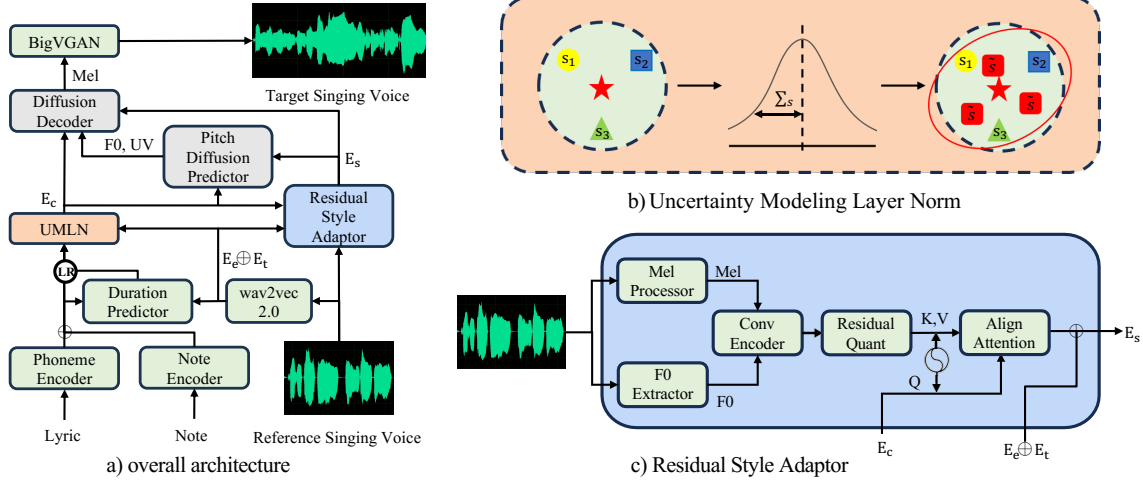


Figure 2: The architecture of StyleSinger. In Figure (a), UMLN is the Uncertainty Modeling Layer Normalization. LR means length regulator. E_t and E_e represent the embedding of timbre and emotion respectively, while E_c and E_s denote the style-agnostic representation and style-specific representation. In Figure (b), s and \tilde{s} are the input and output style information. In Figure (c), mel-spectrograms and f0 are extracted from the reference singing voice .

StyleSinger

In this section, we first define the task of style transfer for out-of-domain singing voice synthesis. Then we overview the proposed StyleSinger. After that, we introduce several critical components including the Uncertainty Modeling Layer Normalization (UMLN), the Residual Style Adaptor (RSA), and architectural details. Finally, we elaborate on the pre-training, training, and inference pipeline of StyleSinger.

Problem Formulation

Given target lyrics and notes, the objective of style transfer for out-of-domain (OOD) singing voice synthesis (SVS) is to generate high-quality target singing voices with unseen styles (such as timbre, emotion, pronunciation, and articulation skills) extracted from reference singing voice samples.

Overview

The architecture of StyleSinger is illustrated in Figure 2(a). Lyrics are encoded through the phoneme encoder, while the note encoder captures musical notes. To extract timbre and emotion embedding from the reference singing voice, we utilize a pre-trained wave2vec 2.0 (Baevski et al. 2020). Then we split our model into style-agnostic and style-specific parts to achieve better generalization (Li et al. 2017, 2019). After predicting the duration, we utilize the Uncertainty Modeling Layer Normalization (UMLN) to perturb the style information in the content representation. This approach enhances the model generalization of StyleSinger and acquires the style-agnostic representation. The reference singing voice is then processed by the Residual Style Adaptor (RSA), which employs a residual quantization module to capture detailed style information (such as pronunciation

and articulation skills) and thus gets the style-specific representation. Subsequently, the pitch diffusion predictor gets both style-agnostic and style-specific representations as inputs to generate F0 and UV. The diffusion decoder then generates mel-spectrograms. Finally, the target singing voice is generated by BigVGan (Lee et al. 2022b).

Uncertainty Modeling Layer Normalization

In general, the style vector is commonly incorporated into the generator by concatenating it with the encoder output. However, this approach can lead to a decline in model performance when encountering OOD scenarios. To address this issue, Chen et al. (2021) introduces conditional layer normalization for style adaptation, allowing for scaling and shifting of the normalized input features based on the style embedding. In this work, we propose the Uncertainty Modeling Layer Normalization (UMLN), which enhances the generalization performance of StyleSinger by incorporating regularization techniques that introduce perturbations to the style information in training samples.

To be more detailed, we can compute the mean μ and variance δ with a hidden vector x . Additionally, given the style vector s , we utilize two simple linear layers to convert the vector into the bias vector $\beta(s)$ and scale vector $\gamma(s)$. To perturb style information, we utilize a Gaussian distribution to model the uncertainty scope of style embedding. By sampling from the uncertainty scope, we can simulate a wide range of diverse unseen style information and effectively prevent the model from generating style-consistent representations. Notably, several studies (Shen and Zhou 2021; Wang et al. 2019) have showcased that the variances observed within features bear implicit semantic connotations. To capture the uncertainties inherent in style embedding, we

calculate the variances of the scale and bias vectors:

$$\begin{aligned}\Sigma_\gamma^2(s) &= \frac{1}{B} \sum_{b=1}^B (\gamma(s) - \mathbb{E}_b[\gamma(s)])^2, \\ \Sigma_\beta^2(s) &= \frac{1}{B} \sum_{b=1}^B (\beta(s) - \mathbb{E}_b[\beta(s)])^2,\end{aligned}\quad (1)$$

where Σ_γ and Σ_β represent the uncertainty estimation of style embedding s . The magnitudes of uncertainty estimation provide the potential transformations that may transpire within the style embedding.

As shown in Figure 2(b), we employ random sampling to perturb the style information in training samples and foster the cultivation of a style-agnostic representation. Drawing inspiration from the previous work (Li et al. 2022), we update the scale and bias vectors:

$$\begin{aligned}\gamma_{um}(s) &= \gamma(s) + \epsilon_\gamma \Sigma_\gamma^2(s), \\ \beta_{um}(s) &= \beta(s) + \epsilon_\beta \Sigma_\beta^2(s),\end{aligned}\quad (2)$$

where ϵ_γ and ϵ_β are drawn from the standard Gaussian distribution $\mathcal{N}(0, 1)$. Upon updating the scale and bias vectors, the style-agnostic hidden representation becomes:

$$UMLN(x, s) = \gamma_{um}(s) \frac{x - \mu(x)}{\delta(x)} + \beta_{um}(s). \quad (3)$$

Ultimately, the model assiduously refines the input features, so attains style-agnostic representation. To strike a delicate balance within this module, we introduce a hyperparameter p , which denotes the probability of using UMLN during the training phase. For the pseudo-code of the algorithm, please refer to Algorithm 1 provided in Appendix B.

Residual Style Adaptor

To intricately model the singing voice styles, we firstly use a wav2vec 2.0 (Baeviski et al. 2020) to capture the timbre and emotion attributes. However, the complexity of styles in singing voices is remarkably high. So we propose the Residual Style Adaptor (RSA) to capture additional style information, like pronunciation and articulation skills.

As illustrated in Figure 2(c), we extract and encode mel-spectrograms and f0 from the reference singing voice sample. In this process, we utilize parselmouth (Jadoul, Thompson, and De Boer 2018) to extract f0 information. Subsequently, we employ a Residual Quantization (RQ) module (Lee et al. 2022a) to extract the detailed style features, which establishes an information bottleneck and effectively eliminates non-style information. RQ has typically been used in the image field. Due to the ability of RQ to extract multiple layers of information, it enables more comprehensive and detailed modeling of style information across various hierarchical levels. In more concrete terms, pronunciation and articulation skills encompass pitch transitions between musical notes and vibrato within a musical note, where the multi-level modeling capability of RQ is highly suitable.

To be specific, the conv encoder generates an output E . With a quantization depth of N , the RQ module represents E as a sequence of N ordered codes. Let $RQ_e(E)$

denote the process of representing E as a stacked map of codes and extracting code embedding. The representation of E in the RQ module at depth $n \in [N]$ is denoted as $\hat{E}^n = \sum_{i=1}^n RQ_e(E)$. To ensure that the input representation adheres to a discrete embedding, a commitment loss (Lee et al. 2022a) is employed:

$$\mathcal{L}_c = \sum_{d=1}^D \left\| E - sg[\hat{E}^n] \right\|_2^2, \quad (4)$$

where the notation sg represents the stop gradient operator. It is important to note that \mathcal{L}_c is the cumulative sum of quantization errors across all n iterations, rather than a single term. The objective is to ensure that \hat{E}^n progressively reduces the quantization error of E as the value of n increases.

After generating the detailed style embedding in the RQ module, it becomes necessary to align the embedding with the content representation E_c . To achieve this, we introduce the Align Attention module, which incorporates the Scaled Dot-Product Attention mechanism (Vaswani et al. 2017). Before feeding the detailed style embedding into the attention module, we include positional encoding embedding. In the attention module, E_c serves as the query, while the detailed style embedding E_d serves as both the key and the value, and d represents the dimensionality of the key and query:

$$\begin{aligned}Attention(Q, K, V) &= Attention(E_c, E_d, E_d) \\ &= Softmax\left(\frac{E_c E_d^T}{\sqrt{d}}\right) E_d.\end{aligned}\quad (5)$$

In the end, we acquire the detailed style representation, which we integrate with the content representation, as well as the timbre and emotion embedding generated from wav2vec 2.0. This integration results in the attainment of the style-specific representation.

Architectural Details

The overall architecture is depicted in Figure 2(a), and we shall briefly introduce a few other crucial components apart from UMLN and RSA.

Encoder Our encoder consists of a note encoder and a phoneme encoder. To be more specific, the phoneme encoder adopts the architecture in FastSpeech2 (Ren et al. 2020), which accepts phonemes as input and yields phoneme features. Meanwhile, the note encoder handles musical scores. It takes note pitches, note types (including rest, slur, grace, etc.), and note duration as inputs, and results in note features. We combine the note and phoneme features to form the content representation. For more detailed information on the encoder, please refer to Appendix A.2.

Pitch Diffusion Predictor When confronted with ever-evolving and dynamic singing voices, simple pitch predictor approaches demonstrate limited effectiveness. To capture the diverse styles in singing voices, we introduce the pitch diffusion predictor. The pitch diffusion predictor consists of the style-specific pitch diffusion predictor and the style-agnostic pitch diffusion predictor, both of which adhere to the same architectural principles as the previous pitch

diffusion model (He et al. 2023). By combining the outputs of them, we obtain the final predictions for F0 and UV. The optimization of this module is achieved through the utilization of Gaussian diffusion loss and multinomial diffusion loss (He et al. 2023). For more details about the pitch diffusion predictor, please refer to Appendix A.4.

Diffusion Decoder The dynamic nature of singing voice poses a challenge for traditional mel decoders, as they can not effectively capture the nuances of mel-spectrograms in singing voices. To tackle this challenge, we employ the diffusion decoder to generate mel-spectrograms. In our approach, we adopt the structure of the teacher model from ProDiff (Huang et al. 2022c), a 4-step generator-based diffusion model. To train the diffusion decoder, we use both the Mean Absolute Error (MAE) loss and Structural Similarity Index (SSIM) loss (Wang et al. 2004). For more details about the diffusion decoder, please refer to Appendix A.5.

Pre-training, Training and Inference Procedures

The final loss terms of StyleSinger consist of the following parts: 1) Duration prediction loss \mathcal{L}_{dur} : MSE between the predicted and the GT phoneme-level duration in log scale; 2) Pitch reconstruction loss $\mathcal{L}_{gdiff}, \mathcal{L}_{mdiff}$: Gaussian diffusion loss and multinomial diffusion loss between the GT and the pitch spectrogram predicted by the pitch diffusion predictor; 3) RQ loss \mathcal{L}_c : the commitment loss for residual quantization layer; 4) Mel reconstruction loss $\mathcal{L}_{mae}, \mathcal{L}_{ssim}$: MAE loss and SSIM loss of the diffusion decoder.

During the pre-training phase, we train the wav2vec 2.0 model to classify timbres and emotions by the AM softmax loss. When training StyleSinger, the reference and target singing voices remain unchanged. During the inference stage, we input lyrics and notes of the target singing voice, and with unseen reference samples, we synthesize target singing voices with OOD reference styles.

Experiments

Experimental Setup

In this section, we first provide an overview of the datasets used in our study. Next, we present the implementation details of our StyleSinger. We then discuss the training and evaluation details for the task. Finally, we introduce the baseline models that we employed for comparison purposes.

Dataset Currently, there are no publicly available SVS datasets with style information. In this endeavor, we collect and annotate a Chinese song corpus (including 12 singers and 20 hours) by recruiting professional singers in a professional recording studio. Additionally, to include more acoustic variation, we incorporate the M4Singer dataset (Zhang et al. 2022a) (including 20 singers and 30 hours), which is used under license CC BY-NC-SA 4.0. Under the guidance of music experts, we manually annotate these datasets with vocal range and emotion labels, categorizing them into 8 classes: tenor happy, tenor sad, soprano happy, soprano sad, bass happy, bass sad, alto happy, and alto sad. Finally, we randomly designate 2 of these classes (tenor happy and alto sad) and 8 singers as unseen styles to evaluate StyleSinger

in the OOD scenario, and then randomly select 20 sentences with unseen styles to construct the OOD testing set.

Implementation Details We utilize pypinyin to convert Chinese lyrics into phonemes. We extract mel-spectrograms from raw waveforms and set the sample rate to 48000Hz, the window size to 1024, the hop size to 256, and the number of mel bins to 80. The default size of the codebook in the RQ is set to 128, and the depth of the RQ is 4. For more information, please refer to Appendix A.1.

Training Details We train our model for 20000 steps using 1 NVIDIA 2080Ti GPU. Adam optimizer is used with $\beta_1 = 0.9, \beta_2 = 0.98$. It takes about 24 hours for training on 1 NVIDIA 2080Ti GPU.

Evaluation Details In our experimental analysis, we employ both objective and subjective evaluation metrics to assess the synthesis quality and style similarity of the test set. For objective evaluation, we utilize the Speaker Cosine Similarity (Cos) to quantify the timbre resemblance between the synthesized and reference singing voices, and F0 Frame Error (FFE) to quantify the synthesis quality. Regarding subjective evaluation, we rely on the Mean Opinion Score (MOS) to gauge naturalness and employ the Similarity Mean Opinion Score (SMOS) (Min et al. 2021) to assess style similarity. Additionally, in the ablation study, we conduct Comparative Mean Opinion Score (CMOS) and Comparative Similarity Mean Opinion Score (CSMOS) evaluations. All these metrics are rated from 1 to 5 and reported with 95% confidence intervals. Moreover, we employ an AXY test (Skerry-Ryan et al. 2018) to evaluate the style transfer performance. We employ the BigVGAN (Lee et al. 2022b) for all experiments. For more detailed information on the evaluation process, please refer to Appendix C.

Baseline Models We conduct a comparative analysis of the quality and similarity of the singing voice samples generated by our esteemed StyleSinger system with other systems, encompassing the following: 1) Reference: The original reference singing voice sample; 2) Reference (vocoder): We transform the reference singing voice sample into mel-spectrograms and subsequently regenerate it into a singing voice using BigVGAN; 3) Styler (Lee, Park, and Kim 2021): We incorporate a module for handling note embedding into Styler, enabling it to generate singing voice performances. 4) GenerSpeech (Huang et al. 2022b): within GenerSpeech, we add a note encoder enabling GenerSpeech to accomplish style transfer of singing voice performances; 5) YourTTS (Casanova et al. 2022): Incorporating the architecture of YourTTS, we likewise integrate a module for note embedding to process the singing voice data. 6) Multi-Style RMSSinger (He et al. 2023) (MS RMSSinger): we enrich the architecture of RMSSinger by integrating the timbre and emotion vectors extracted by wav2vec 2.0 into its backbone, allowing it to handle style transfer tasks.

Main Results

We randomly select singing voice samples from the OOD testing sets as references to assess the style transfer capabilities of StyleSinger and baseline models. Based on the

Method	MOS \uparrow	SMOS \uparrow	Cos \uparrow	FFE \downarrow
Refernece	4.53 \pm 0.05	-	-	-
Reference (vocoder)	4.13 \pm 0.07	4.26 \pm 0.09	0.97	0.05
Styler (Lee, Park, and Kim 2021)	3.52 \pm 0.08	3.79 \pm 0.07	0.77	0.39
GenerSpeech (Huang et al. 2022b)	3.59 \pm 0.07	3.83 \pm 0.08	0.83	0.36
YourTTS (Casanova et al. 2022)	3.65 \pm 0.09	3.85 \pm 0.10	0.85	0.31
MS RMSSinger (He et al. 2023)	3.84 \pm 0.08	3.90 \pm 0.11	0.88	0.29
StyleSinger (ours)	3.90 \pm 0.05	4.03 \pm 0.07	0.90	0.27

Table 1: The quality and style similarity of parallel style transfer when extended to out-of-domain test sets. For subjective measurement, we employ MOS and SMOS. In objective evaluation, we utilize Cos and FFE.

Baseline	Parallel				Non-Parallel			
	7-point score	Perference (%)			7-point score	Perference (%)		
		X	Neutral	Y		X	Neutral	Y
Styler	1.31 \pm 0.14	18%	26%	56%	1.40 \pm 0.12	10%	22%	68%
GenerSpeech	1.29 \pm 0.11	28%	20%	52%	1.33 \pm 0.09	12%	24%	64%
YourTTS	1.28 \pm 0.08	26%	24%	50%	1.30 \pm 0.10	16%	20%	64%
MS RMSSinger	1.14 \pm 0.10	28%	30%	42%	1.26 \pm 0.12	24%	16%	60%

Table 2: The AXY preference test results for parallel and non-parallel style transfer are presented. From the testing sets, we have selected 20 samples for evaluation. Raters were requested to assign a 7-point score (ranging from -3 to 3) and select the samples that sounded closer to the target style. In this context, X represents a baseline model, while our StyleSinger is denoted as Y. A higher score indicates that Y is closer to the target style compared to X.

content consistency between the reference and generated singing voices, we categorize the experiments into parallel and non-parallel style transfer (Skerry-Ryan et al. 2018).

Parallel Style Transfer In the context of out-of-domain (OOD) scenarios, where the content of the reference voice remains unchanged, the primary outcomes are presented in Table 1. Based on both objective and subjective evaluations, the following observations can be made: 1) StyleSinger demonstrates exceptional audio quality, as evidenced by the highest Mean Opinion Score (MOS) among all models. This signifies the model’s remarkable universality in handling out-of-domain (OOD) scenarios. 2) StyleSinger also excels in style similarity, as indicated by the highest Style Mean Opinion Score (SMOS). This showcases the model’s exceptional ability to accurately model and capture the nuances of different singing styles. 3) As measured by the objective indicators Cos and FFE, StyleSinger consistently delivers the best results. These findings collectively demonstrate the remarkable effectiveness of StyleSinger in OOD scenarios for singing voice synthesis and style transfer. This can be attributed to the exceptional generalization capability of UMLN, the adeptness of the RSA in modeling style representations, and the integration of the pitch diffusion predictor and the diffusion decoder, which imbue the generated OOD singing voices with enhanced details and vividness. For more details, please refer to Appendix D.1.

Non-Parallel Style Transfer In out-of-domain (OOD) scenarios, we utilize unseen reference samples with target notes and lyrics to synthesize the target singing voice. To

evaluate the performance, we conducted an AXY preference test by randomly selecting 20 unseen reference singing voice samples with target notes and lyrics. Then we compared the synthesis results of StyleSinger with baseline models. As shown in Table 2, the results demonstrate a clear preference for the synthesis generated by StyleSinger over the baselines. This affirms the effectiveness of our Residual Style Adaptor (RSA) and uncertainty modeling layer norm (UMLN) in achieving successful unseen style transfer.

We proceed to visualize mel-spectrograms and pitch contour in the context of non-parallel style transfer. In Figure 3, it can be observed that: 1) StyleSinger excels at capturing the intricate nuances of the reference style. The pitch curve generated by StyleSinger displays a greater range of variations and finer details, closely resembling the reference style. To be more precise, StyleSinger effectively captures the vibrato style, as well as the nuances of pronunciation and articulation skills. In contrast, the curves generated by other methods appear relatively flat, lacking distinctions in singing techniques. 2) StyleSinger excels in modeling mel-spectrograms of higher quality and intricate details. The mel-spectrograms generated by StyleSinger exhibit superior quality, showcasing rich details in frequency bins between adjacent harmonics and high-frequency components. In contrast, the mel-spectrograms generated by other methods for out-of-domain (OOD) samples demonstrate lower quality and a lack of intricate details.

When listening to the demo, it is evident that our model effectively captures the timbre, emotion, pitch transitions, vibrato, pronunciation, and articulation skills present in the

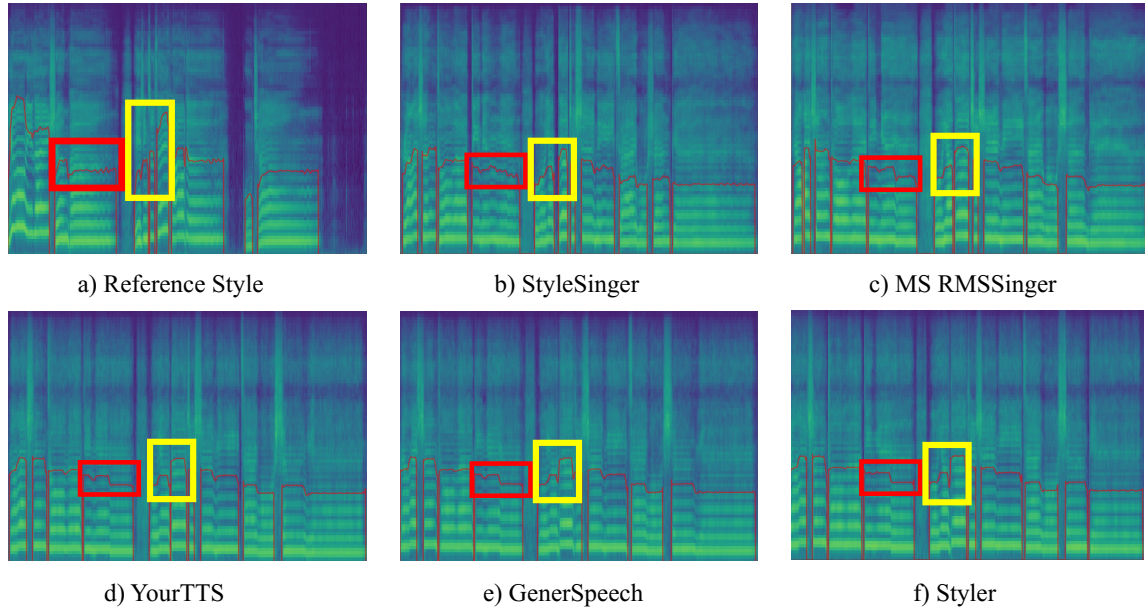


Figure 3: The mel-spectrograms depicting the results of non-parallel style transfer. StyleSinger effectively captures the vibrato style indicated by red boxes, along with the pronunciation and articulation skills highlighted in yellow boxes.

Setting	CMOS	CSMOS
StyleSinger	0.00	0.00
w/o UMLN	-0.28	-0.25
w/o RSA	-0.12	-0.29
w/o Pitch	-0.40	-0.23
w/o Decoder	-0.52	-0.19

Table 3: Audio quality and similarity comparisons for ablation study with CMOS and CSMOS. UMLN and RSA are the Uncertainty Modeling Layer Normalization and the Residual Style Adaptor, while Pitch and Decoder mean the pitch diffusion predictor and the diffusion decoder.

reference singing voice samples. Furthermore, it can be discerned that StyleSinger surpasses baseline models in synthesis quality and similarity to reference singing voice samples.

Ablation Study

As depicted in Table 3, we undertake ablation studies to showcase the efficacy of various designs incorporated within StyleSinger. We conduct CMOS (comparative mean opinion score) and CSMOS (comparative similarity mean opinion score) evaluations. 1) When we eliminate the uncertainty modeling layer norm (UMLN), the quality and similarity decline, indicating the enhancement our method brings to model generalization performance. 2) As the Residual Style Adaptor (RSA) is removed, the similarity significantly decreases, demonstrating the effectiveness of our method in modeling the intricate styles in singing voices. 3) Exclud-

ing the pitch diffusion predictor, we utilize the simple pitch predictor in FastSpeech2 (Ren et al. 2020), and the quality deteriorates further, highlighting the improvement our pitch predictor brings to f0 modeling. 4) Without the diffusion decoder, we employ a transformer decoder (Ren et al. 2020) instead. The significant decline in audio quality highlights the crucial role of the diffusion decoder in generating high-quality mel-spectrograms. For more detailed results of the ablation study, please refer to Appendix D.2.

Conclusion

In this paper, we present a pioneering approach StyleSinger, the first singing voice synthesis model capable of achieving high-quality zero-shot style transfer for out-of-domain voices. We primarily enhance the model’s performance through two key components: 1) the Residual Style Adaptor (RSA) which employs a residual quantization module to capture diverse style characteristics in singing voices, and 2) the Uncertainty Modeling Layer Normalization (UMLN) to perturb the style attributes within the content representation during the training phase and thus improve the model generalization. Our extensive evaluations in zero-shot style transfer undeniably establish that StyleSinger outperforms baseline models in both audio quality and similarity to the reference samples. For future work, we aim to expand the capabilities of StyleSinger to encompass a broader range of scenarios, such as multilingual tasks. Additionally, it is a good idea to explore training models that integrate both speech and singing styles, which can generate singing voices with styles extracted from OOD reference speech.

Ethics Statement

StyleSinger’s ability to perform out-of-domain style transfer for singing voices raises concerns regarding potential unfair competition and the potential displacement of professionals within the music industry. Additionally, its application in the entertainment sector, including short videos, may give rise to copyright issues. Therefore, we will impose restrictions on our code and models to prevent unauthorized usage.

Acknowledgements

This work was supported in part by the National Key R&D Program of China under Grant No.2022ZD0162000, National Natural Science Foundation of China under Grant No.62222211, Grant No.61836002 and Grant No.62072397, and Yiwise.

References

- Baevski, A.; Zhou, Y.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33: 12449–12460.
- Casanova, E.; Shulby, C.; Gölge, E.; Müller, N. M.; de Oliveira, F. S.; Junior, A. C.; Soares, A. d. S.; Aluisio, S. M.; and Ponti, M. A. 2021. Sc-glowtts: an efficient zero-shot multi-speaker text-to-speech model. *arXiv preprint arXiv:2104.05557*.
- Casanova, E.; Weber, J.; Shulby, C. D.; Junior, A. C.; Gölge, E.; and Ponti, M. A. 2022. Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone. In *International Conference on Machine Learning*, 2709–2720. PMLR.
- Chen, M.; Tan, X.; Li, B.; Liu, Y.; Qin, T.; Zhao, S.; and Liu, T.-Y. 2021. Adaspeech: Adaptive text to speech for custom voice. *arXiv preprint arXiv:2103.00993*.
- Choi, S.; Han, S.; Kim, D.; and Ha, S. 2020. Attention: Few-shot text-to-speech utilizing attention-based variable-length embedding. *arXiv preprint arXiv:2005.08484*.
- Choi, S.; and Nam, J. 2022. A melody-unsupervision model for singing voice synthesis. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7242–7246. IEEE.
- Cooper, E.; Lai, C.-I.; Yasuda, Y.; Fang, F.; Wang, X.; Chen, N.; and Yamagishi, J. 2020. Zero-shot multi-speaker text-to-speech with state-of-the-art neural speaker embeddings. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6184–6188. IEEE.
- He, J.; Liu, J.; Ye, Z.; Huang, R.; Cui, C.; Liu, H.; and Zhao, Z. 2023. RMSSinger: Realistic-Music-Score based Singing Voice Synthesis. *arXiv preprint arXiv:2305.10686*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, R.; Chen, F.; Ren, Y.; Liu, J.; Cui, C.; and Zhao, Z. 2021. Multi-singer: Fast multi-singer singing voice vocoder with a large-scale corpus. In *Proceedings of the 29th ACM International Conference on Multimedia*, 3945–3954.
- Huang, R.; Cui, C.; Chen, F.; Ren, Y.; Liu, J.; Zhao, Z.; Huai, B.; and Wang, Z. 2022a. Singgan: Generative adversarial network for high-fidelity singing voice generation. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2525–2535.
- Huang, R.; Ren, Y.; Liu, J.; Cui, C.; and Zhao, Z. 2022b. Generspeech: Towards style transfer for generalizable out-of-domain text-to-speech synthesis. *arXiv preprint arXiv:2205.07211*.
- Huang, R.; Zhao, Z.; Liu, H.; Liu, J.; Cui, C.; and Ren, Y. 2022c. Prodiff: Progressive fast diffusion model for high-quality text-to-speech. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2595–2605.
- Huang, S.-F.; Lin, C.-J.; Liu, D.-R.; Chen, Y.-C.; and Lee, H.-y. 2022d. Meta-tts: Meta-learning for few-shot speaker adaptive text-to-speech. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30: 1558–1571.
- Jadoul, Y.; Thompson, B.; and De Boer, B. 2018. Introducing parselmouth: A python interface to praat. *Journal of Phonetics*, 71: 1–15.
- Jia, Y.; Zhang, Y.; Weiss, R.; Wang, Q.; Shen, J.; Ren, F.; Nguyen, P.; Pang, R.; Lopez Moreno, I.; Wu, Y.; et al. 2018. Transfer learning from speaker verification to multispeaker text-to-speech synthesis. *Advances in neural information processing systems*, 31.
- Kim, S.; Kim, Y.; Jun, J.; and Kim, I. 2023. MuSE-SVS: Multi-Singer Emotional Singing Voice Synthesizer that Controls Emotional Intensity. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Kumar, N.; Goel, S.; Narang, A.; and Lall, B. 2021. Normalization Driven Zero-Shot Multi-Speaker Speech Synthesis. In *Interspeech*, 1354–1358.
- Lee, D.; Kim, C.; Kim, S.; Cho, M.; and Han, W.-S. 2022a. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11523–11532.
- Lee, K.; Park, K.; and Kim, D. 2021. Styler: Style factor modeling with rapidity and robustness via speech decomposition for expressive and controllable neural text to speech. *arXiv preprint arXiv:2103.09474*.
- Lee, S.-g.; Ping, W.; Ginsburg, B.; Catanzaro, B.; and Yoon, S. 2022b. Bigvgan: A universal neural vocoder with large-scale training. *arXiv preprint arXiv:2206.04658*.
- Li, D.; Yang, Y.; Song, Y.-Z.; and Hospedales, T. M. 2017. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, 5542–5550.
- Li, X.; Dai, Y.; Ge, Y.; Liu, J.; Shan, Y.; and Duan, L.-Y. 2022. Uncertainty modeling for out-of-distribution generalization. *arXiv preprint arXiv:2202.03958*.
- Li, X.; Song, C.; Li, J.; Wu, Z.; Jia, J.; and Meng, H. 2021. Towards multi-scale style control for expressive speech synthesis. *arXiv preprint arXiv:2104.03521*.
- Li, Y.; Yang, Y.; Zhou, W.; and Hospedales, T. 2019. Feature-critic networks for heterogeneous domain generalization. In *International Conference on Machine Learning*, 3915–3924. PMLR.

- Liu, J.; Li, C.; Ren, Y.; Chen, F.; and Zhao, Z. 2022. Diff-singer: Singing voice synthesis via shallow diffusion mechanism. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 11020–11028.
- Min, D.; Lee, D. B.; Yang, E.; and Hwang, S. J. 2021. Meta-style speech: Multi-speaker adaptive text-to-speech generation. In *International Conference on Machine Learning*, 7748–7759. PMLR.
- Nuriel, O.; Benaim, S.; and Wolf, L. 2021. Permuted adain: Reducing the bias towards global statistics in image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9482–9491.
- Paul, D.; Pantazis, Y.; and Stylianou, Y. 2020. Speaker conditional WaveRNN: Towards universal neural vocoder for unseen speaker and recording conditions. *arXiv preprint arXiv:2008.05289*.
- Ren, Y.; Hu, C.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T.-Y. 2020. FastSpeech 2: Fast and high-quality end-to-end text to speech. *arXiv preprint arXiv:2006.04558*.
- Shen, Y.; and Zhou, B. 2021. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1532–1540.
- Skerry-Ryan, R.; Battenberg, E.; Xiao, Y.; Wang, Y.; Stanton, D.; Shor, J.; Weiss, R.; Clark, R.; and Saurous, R. A. 2018. Towards end-to-end prosody transfer for expressive speech synthesis with tacotron. In *international conference on machine learning*, 4693–4702. PMLR.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.
- Wang, Y.; Pan, X.; Song, S.; Zhang, H.; Huang, G.; and Wu, C. 2019. Implicit semantic data augmentation for deep networks. *Advances in Neural Information Processing Systems*, 32.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Zhang, L.; Li, R.; Wang, S.; Deng, L.; Liu, J.; Ren, Y.; He, J.; Huang, R.; Zhu, J.; Chen, X.; et al. 2022a. M4singer: A multi-style, multi-singer and musical score provided mandarin singing corpus. *Advances in Neural Information Processing Systems*, 35: 6914–6926.
- Zhang, Y.; Cong, J.; Xue, H.; Xie, L.; Zhu, P.; and Bi, M. 2022b. Visinger: Variational inference with adversarial learning for end-to-end singing voice synthesis. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7237–7241. IEEE.
- Zhang, Z.; Zheng, Y.; Li, X.; and Lu, L. 2022c. Wesinger: Data-augmented singing voice synthesis with auxiliary losses. *arXiv preprint arXiv:2203.10750*.
- Zhou, K.; Yang, Y.; Qiao, Y.; and Xiang, T. 2021. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*.