



Interpretable Style Transfer for Text-to-Speech with ControlVAE and Diffusion Bridge

Wenhao Guan¹, Tao Li¹, Yishuang Li², Hukai Huang¹, Qingyang Hong^{1*}, Lin Li^{2,3*}

¹School of Informatics, Xiamen University, China

²Institute of Artificial Intelligence, Xiamen University, China

³School of Electronic Science and Engineering, Xiamen University, China

{qyhong, lilin}@xmu.edu.cn

Abstract

With the demand for autonomous control and personalized speech generation, the style control and transfer in Text-to-Speech (TTS) is becoming more and more important. In this paper, we propose a new TTS system that can perform style transfer with interpretability and high fidelity. Firstly, we design a TTS system that combines variational autoencoder (VAE) and diffusion refiner to get refined mel-spectrograms. Specifically, a two-stage and a one-stage system are designed respectively, to improve the audio quality and the performance of style transfer. Secondly, a diffusion bridge of quantized VAE is designed to efficiently learn complex discrete style representations and improve the performance of style transfer. To have a better ability of style transfer, we introduce ControlVAE to improve the reconstruction quality and have good interpretability simultaneously. Experiments on LibriTTS dataset demonstrate that our method is more effective than baseline models¹.

Index Terms: speech synthesis, style transfer, variational autoencoder, diffusion probabilistic model

1. Introduction

With the rapid development of deep learning, the common speech synthesis task that only takes text as input has recently almost reached the human level in the public dataset [1]. More researchers focused on the controllable and expressive TTS [2, 3] because the controllability provides users with more flexibility, and the improvement of expressiveness makes people have better auditory experience.

At present, most TTS systems [4, 5, 6, 7] can model controllable style attributes in different ways. FastSpeech [4] which is a non-autoregressive model uses the trained autoregressive model as the teacher model to train the duration predictor, and utilizes the length regulator to control the duration information, thus indirectly controls the overall style. FastSpeech2 [5] employs the external duration aligner Montreal Forced Aligner (MFA) [8] to train the duration predictor on the basis of FastSpeech, furthermore, it trains the pitch predictor, energy predictor and predictors of other styles in a supervised way to achieve more precise style control. Most TTS system pipelines have two modules. The first module is to produce intermediate representations and the second module is to generate raw waveforms conditioned on the intermediate representations. VITS [7] connects the two modules, generated by VAE [9], through latent variables to enable efficient end-to-end training, and it proposes a stochastic duration predictor to synthesize speech with different variations. However, they don't do style transfer tasks so

they do not need reference speech as input. Subsequently, there were many works focused on designing powerful style encoders for style transfer tasks. Some works based on autoregressive models obtain fine-grained style representations by designing multi-level or multi-scale style modeling methods [10, 11, 12]. Due to the problem of low decoding speed in the autoregressive models, the following style modeling tasks mostly adopt a non-autoregressive architecture [4, 5, 13]. Meta-StyleSpeech [14] adopts the base architecture upon FastSpeech2, applying style adaptive layer norm and meta-learning algorithm to effectively synthesize style-transferred speech. Styler [15] models style factors by speech decomposition via information bottleneck. GenerSpeech [16] proposes a multi-level style adaptor and a generalizable content adaptor to efficiently model the style information. NoreSpeech [17] proposes a knowledge distillation based conditional diffusion model to generate style representation from noisy reference audios. But these methods often cannot perform style interpretability. Some works [2, 18] are proposed to demonstrate that they have the ability of style interpretability and transfer simultaneously. Global style token (GST) [18] designs a style token layer and a reference encoder to explore the expressiveness of TTS systems unsupervisedly. VAE-Tacotron [2] learns the style representation through VAE [9]. A recent work Style-Label-Free [19] extends VAE-Tacotron by proposing Quantized VAE and speaker-wise normalization in cross-speaker style transfer.

In this paper, we focus on high-fidelity style transfer and interpretability in speech synthesis which requires a better style representation and an interpretable disentangled style latent space. We utilize VAE based style encoder to have access to the interpretable latent space and integrate VAE within diffusion probabilistic models (DPM) [20] to overcome the over-smoothness problem [21]. We also propose a diffusion bridge of Quantized VAE to improve the diversity of generated style representations. We further introduce ControlVAE [22] to our system instead of original VAE to have better reconstruction quality and good interpretability.

Our main contributions can be summarized as follows:

- We propose a new TTS system that incorporates VAE and DPM to get refined mel-spectrograms. Specifically, a two-stage and a one-stage training pipeline are designed respectively to improve the performance of style transfer.
- A diffusion bridge of Quantized VAE is proposed to model the diversity of style representations in latent space so that the TTS system achieve better performance of style transfer.
- We introduce ControlVAE in our system to replace original VAE to a better reconstruction ability so that improve the style transfer quality and have good style interpretability.

*Corresponding authors

¹Audio samples are publicly available at <https://gwh22.github.io/>.

2. Related Work

2.1. Diffusion probabilistic models

The concept of diffusion was first defined in [23] and then researchers proposed denoising diffusion probabilistic model (DDPM) [20] which greatly promotes the development of generative models. DDPM has recently succeeded to advance the state-of-the-art results in speech synthesis [13, 24].

Diffusion process and reverse process are given by diffusion probabilistic models, which could be used for the denoising neural networks θ to learn data distribution. Similar as previous work [23, 20], we define the data distribution as $q(x_0)$. Let x_1, \dots, x_T be a sequence of variables with the same dimension. The diffusion process is defined by a fixed Markov chain from data x_0 to the latent variable x_T :

$$q(x_t|x_{t-1}) = N(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I) \quad (1)$$

$$q(x_1, \dots, x_T|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}) \quad (2)$$

The reverse process aims to recover samples from Gaussian noise, which is Markov chain from x_T to x_0 parameterized by shared θ :

$$p_\theta(x_{t-1}|x_t) = N(x_{t-1}; \mu_\theta(x_t, t), \sigma_t^2 I) \quad (3)$$

$$p_\theta(x_0, \dots, x_{T-1}|x_T) = \prod_{t=1}^T p_\theta(x_{t-1}|x_t) \quad (4)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, μ_θ and σ_t^2 represent the mean and standard derivation respectively. Finally, we can get the training objective as follows:

$$L_{DDPM} = E_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t)\|_2^2] \quad (5)$$

where ϵ is the Gaussian noise and $\epsilon_\theta(\cdot)$ is the output of model. For sampling phase, the sampling formulation is computed as follows:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) + \sigma_t z, \quad (6)$$

where $\epsilon \sim N(0, I)$, $p_z = N(z; 0, I)$ and $\sigma_t = \sqrt{\frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t}$. As a result, the final data distribution p_{x_0} is obtained through iterative sampling over all of the time steps.

2.2. Variational Autoencoder

VAE has been applied for latent representation learning of natural speech for years [2, 19]. It is assumed that the observed data distribution $p(x)$ is generated by some random process from a random latent variable z . The true posterior distribution $p_\theta(z|x)$ is intractable because of the undifferentiable marginal likelihood $p_\theta(x)$. To address this problem, $q_\phi(z|x)$ is introduced as an approximation to the true posterior distribution $p_\theta(z|x)$. Finally, we can get the formulation of $\log p_\theta(x)$:

$$\begin{aligned} \log p_\theta(x) &\geq E_{q_\phi(z|x)} [\log \frac{p_\theta(x, z)}{q_\phi(z|x)}] \\ &= E_{q_\phi(z|x)} [\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x) \| p_\theta(z)) \end{aligned} \quad (7)$$

The encoder of the VAE is chosen to model a multivariate Gaussian with diagonal covariance, and the prior is often assumed to be a standard multivariate Gaussian:

$$q_\phi(z|x) = N(z; \mu_\phi(x), \sigma_\phi^2(x)I) \quad (8)$$

$$p_z = N(z; 0, I) \quad (9)$$

where $\mu(x)$ and $\sigma^2(x)$ in $q_\phi(z|x)$ are learned via neural network, and the reparameterization trick is introduced to VAE framework to avoid non-derivable problem. Thus, each z is computed as a deterministic function of input x and auxiliary noise variable ϵ , where \odot represents an element-wise product.

$$z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon \quad (10)$$

2.3. Quantized VAE

Quantized VAE is proposed in [19] to improve the representation ability of original VAE encoder. Quantized VAE simply extends VAE by adding a discrete codebook component to the VAE output z . z is compared with all the vectors in the codebook, and the closest codebook vector is fed into the VAE decoder. The vector quantization (VQ) loss which consists of commitment loss and codebook loss is as follows:

$$L_Q = \|sg[z] - q\|_2^2 + \gamma \|z - sg[q]\|_2^2, \quad (11)$$

where z is referred as VAE output, q is referred as the codebook vector, γ represents the weight of commitment loss, $sg[\cdot]$ refers to the operation of stop gradient.

3. Proposed Method

Figure 1 shows our proposed TTS method named IST-TTS (The abbreviation of Interpretable Style Transfer for Text-to-Speech). Our proposed method consists of three parts: the (1) diffusion refiner, (2) diffusion bridge, and (3) ControlVAE.

3.1. Model architecture

The overall architecture of our method is illustrated in Figure 1 (a). A reference Mel-spectrogram is fed into a reference encoder to extract style information, and the style information is then pass through ControlVAE to obtain the interpretable latent space Z . The quantized embedding is obtained by proposed diffusion bridge, and it is fed into the acoustic model. We use the same architecture for diffusion bridge as [25] to learn diverse style representations.

The acoustic model is shown as Figure 1 (b), which is based on the architecture of FastSpeech [4]. And a diffusion refiner [13] is designed for two-stage and one-stage training pipeline to explore the effectiveness of combining VAE and DPM. Additionally, the speaker embedding is extracted by x-vector [26]. Note that we use MFA [8] to replace the previous distillation method to effectively train the duration predictor.

3.2. Diffusion Refiner

We propose to incorporate VAE framework and DPM in two-stage and one-stage training pipeline respectively to explore the effectiveness of our method.

3.2.1. Two-stage and one-stage training pipeline

As shown in the left of Figure 1 (b), in two-stage training pipeline, the model firstly generates an intermediate mel-spectrogram and it can be fed into the vocoder to get the intermediate waveform, the model is named as VAEFS in this paper. And then the intermediate mel-spectrogram which is processed by a linear layer is fed into diffusion refiner as the condition of diffusion model, the model is named as VAEFS+2s in this paper.

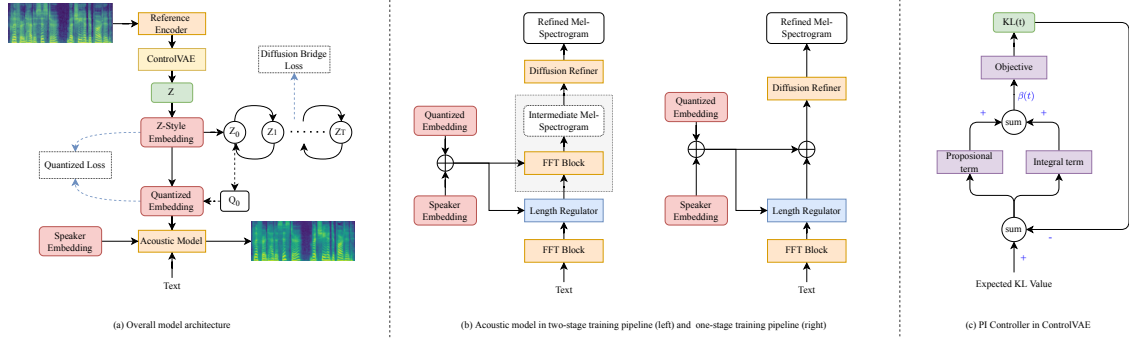


Figure 1: (a) the overall architecture of IST-TTS. (b) the acoustic model in two-stage training pipeline (left) and one-stage training pipeline (right). (c) the PI controller algorithm in ControlVAE.

As shown in the right of Figure 1 (b), in one-stage training pipeline, it will degrade into the method of Diff-TTS [13] in acoustic model, which is named as VAEFS+1s in this paper. Note that the designs in the rest of this paper are based on the acoustic model in one-stage training pipeline.

3.2.2. Conditional diffusion model

Diffusion refiner in our method is a conditional diffusion model because the input needs external intermediate mel-spectrogram or decoder input to be as the condition of diffusion model. c represents the condition. The training objective is as follows:

$$L_R = E_{t, x_0, \epsilon, c} [\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, t, c)\|_2^2] \quad (12)$$

3.3. Diffusion Bridge

Quantized VAE discretizes the latent features using vector quantization to generate more expressive samples. Inspired by recent work [27], we propose a new Diffusion Bridge to improve the ability of expressiveness in Quantized VAE. Specifically, we consider using diffusion models in a continuous space, which is the z latent space of VAE output in our work, to efficiently learn complex discrete distributions. Note that the sampling process of diffusion bridge is only used in inference stage. The process is shown in Figure 1 (a) and the training loss is:

$$L_B = E_{t, z_0, \epsilon} [\|\epsilon - \epsilon_\theta(\sqrt{\alpha_t}z_0 + \sqrt{1 - \alpha_t}\epsilon, t)\|_2^2] \quad (13)$$

3.4. ControlVAE

The original VAE models may suffer from KL vanishing and low reconstruction quality problem. To address the issues, we propose ControlVAE [22] which combines a controller with the basic VAE as an alternative of original VAE. Specifically, a new non-linear proportional-integral (PI) controller is designed to automatically tune the weight added in the VAE objective using the output KL-divergence as feedback during model training. The PI controller algorithm can be illustrated as Figure 1 (c), the weight $\beta(t)$ and the loss of ControlVAE are as follows:

$$\beta(t) = \frac{K_p}{1 + \exp(e(t))} - K_i \sum_{j=0}^t e(j) + \beta_{min} \quad (14)$$

where K_p and K_i represent the coefficient of propositional term and integral term. $e(t)$ is the error between current KL value and expected KL value. β_{min} is a constant.

$$L_C = E_{q_\phi(z|x)} [\log p_\theta(x|z)] - \beta(t) D_{KL}(q_\phi(z|x) || p(z)) \quad (15)$$

Table 1: The objective results of style transfer. The last three rows show the results of ablation study.

Method	FD ↓	MCD ↓
VAEFS	6.36	5.82
VAEFS+2s	3.69	5.79
VAEFS+1s	4.47	5.84
GenerSpeech	4.83	5.66
IST-TTS (ours)	3.37	5.63
w/o ControlVAE	3.60	5.90
w/o VQ	4.27	6.21
w/o Diffusion Bridge	3.81	5.67

Finally, the total training loss of our method is:

$$L_{AU} = L_C + L_R + L_Q + L_B \quad (16)$$

4. Experimental Setup

4.1. Dataset

As for dataset, we use LibriTTS [28], which is a multispeaker English corpus containing 586 hours of speech clips with 2456 speakers, to evaluate our method. And we convert the sampling rate to 22050Hz and extract the spectrogram with the FFT size of 1024, hop size of 256, and window size of 1024.

4.2. Training setting

The proposed model was trained for 320K iterations using Adam optimizer [29] on a single NVIDIA TESLA V100 GPU. Additionally, we utilize a pretrained HiFi-GAN [30] as the neural vocoder to convert mel-spectrogram to waveform.

The hyperparameters of reference encoder is the same as [2], the dimension of z , the codebook size and γ of vector quantization are set to 32, 1024 and 0.25, the hyperparameters β_{min} , K_p , K_i and expected KL value in ControlVAE are set to 0, 0.01, 0.0001 and 3 respectively. Exponential moving averages (EMA) [31] is used instead of codebook loss in Quantized VAE to get faster convergence speed. We use cost annealing method in original VAE. T in diffusion refiner (bridge) is set to 1000.

4.3. Baseline models

We compare the quality and similarity of generated audios by our IST-TTS with public baseline model GenerSpeech [16] and our proposed baseline models, VAEFS(+1s/2s).

Table 2: The subjective results of style transfer. 1s and 2s represent one stage and two stage training with diffusion refiner respectively. B,S,O denote baseline, same and ours respectively. The last three rows show the results of ablation study.

Method	5-scale Score		Parallel Style Transfer			Non-Parallel Style Transfer				
	MOS	SMOS	7-point score	Preference (%)		7-point score	Preference (%)			
				B	S	O	B	S	O	
Reference	4.52±0.062	—								
Reference(Mel+Voc)	4.49±0.076	4.41±0.065								
VAEFS	3.31±0.081	2.85±0.083	1.75	12	14	74	1.60	22	28	50
VAEFS+2s	3.75±0.092	3.21±0.076	1.30	15	18	67	1.52	26	33	41
VAEFS+1s	3.79±0.079	3.26±0.105	1.21	20	25	55	1.32	15	38	47
GenerSpeech	3.96±0.086	3.99±0.089	0.21	32	34	34	0.51	32	32	36
IST-TTS (ours)	4.17±0.055	4.02±0.091								
w/o ControlVAE	4.09±0.056	3.86±0.083	0.76	25	35	40	0.60	26	40	34
w/o VQ	3.85±0.061	3.81±0.076	0.91	19	37	44	0.69	29	36	35
w/o Diffusion Bridge	4.01±0.081	3.79±0.073	0.71	30	27	43	0.55	34	28	38




Figure 2: The comparison of (a) VAEFS, (b) VAEFS+2s and (c) VAEFS+1s.




Figure 3: The generated mel-spectrograms for disentangling.

5. Experiment Results and Analysis

5.1. Style Transfer Quality and Similarity

We evaluate the style transfer quality and similarity by objective metrics and subjective evaluations. The objective metrics include frechet distance (FD) [32] metric and mel cepstral distortion (MCD) metric. The FD in audio indicates the similarity between generated samples and target samples and MCD evaluates the compatibility between the spectra of two audio clips. As for subjective evaluations, we conduct 5-scale Mean Opinion Score (MOS) and Similarity Mean Opinion Score (SMOS) test between IST-TTS and the baselines. To further evaluate the style transfer performance, an AXY test used in [16] is conducted, the range of 7-point score is from -3 to 3, 0 represents “Both are about the same distance”.

5.1.1. Parallel Style Transfer

As for parallel style transfer, The objective results of FD and MCD shown in Table 1 show that our proposed IST-TTS outperform baseline models, indicating that the proposed model generates higher quality results. The MOS, SMOS and parallel style transfer results in Table 2 show our IST-TTS achieve higher quality and similarity.

Furthermore, we show the results of comparison between VAEFS, VAEFS+2s and VAEFS+1s in Figure 2, inferring that our diffusion refiner can overcome the oversmoothness problem

in VAEFS so that obtain better performance of quality.

5.1.2. Non-parallel Style Transfer

Non-parallel style transfer represents that the text is changed from the reference utterance. The non-parallel style transfer results are shown in Table 2, which denote that our method has a better performance in non-parallel style transfer than baseline models by using ControlVAE and diffusion bridge to get more expressive results.

5.1.3. Ablation Study

As shown in the last three rows of Table 1 and Table 2, we conduct ablation studies to demonstrate the effectiveness of several designs in IST-TTS, including ControlVAE, vector quantization (VQ) and diffusion bridge. Both the objective metrics and subjective scores drop when removing VQ or diffusion bridge, and replacing ControlVAE with the original VAE results in decreased quality and similarity. These demonstrate the efficiency of the proposed method in modelling expressive style representations.

5.2. Style Interpretability

To evaluate the performance of style interpretability, we experiment on single dimension of ControlVAE latent space z to explore the ability of disentangling as shown in Figure 3 which represents some different speaking styles of energy, pitch variation and pitch level.

The limitation in the analysis of style interpretability is that the dataset we chose does not have detailed style labels, so there is no way to conduct quantitative interpretable comparisons.

6. Conclusions

In this paper, we propose a TTS method for interpretable style transfer named IST-TTS, which incorporates VAE with diffusion refiner to improve the audio quality and the performance of style transfer. Furtherly we propose diffusion bridge to perform better style transfer. And finally we introduce ControlVAE to achieve better style transfer quality and good disentanglement. In the future, we will explore better method for interpretability.

7. Acknowledgements

Thanks to the National Natural Science Foundation of China (Grant No.62276220, No.62001405 and No.61876160) for funding.

8. References

- [1] X. Tan, J. Chen, H. Liu, J. Cong, C. Zhang, Y. Liu, X. Wang, Y. Leng, Y. Yi, L. He *et al.*, “Naturalspeech: End-to-end text to speech synthesis with human-level quality,” *arXiv preprint arXiv:2205.04421*, 2022.
- [2] Y.-J. Zhang, S. Pan, L. He, and Z.-H. Ling, “Learning latent representations for style control and transfer in end-to-end speech synthesis,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6945–6949.
- [3] X. An, F. K. Soong, and L. Xie, “Disentangling style and speaker attributes for tts style transfer,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 646–658, 2022.
- [4] 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 *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, 2019, pp. 3171–3180.
- [5] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, “Fastspeech 2: Fast and high-quality end-to-end text to speech,” in *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- [6] J. Kim, S. Kim, J. Kong, and S. Yoon, “Glow-tts: A generative flow for text-to-speech via monotonic alignment search,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 8067–8077, 2020.
- [7] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 5530–5540.
- [8] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldif,” in *Interspeech*, vol. 2017, 2017, pp. 498–502.
- [9] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [10] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, A. Rosenberg, B. Ramabhadran, and Y. Wu, “Generating diverse and natural text-to-speech samples using a quantized fine-grained vae and autoregressive prosody prior,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6699–6703.
- [11] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, and Y. Wu, “Fully-hierarchical fine-grained prosody modeling for interpretable speech synthesis,” in *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2020, pp. 6264–6268.
- [12] X. Li, C. Song, J. Li, Z. Wu, J. Jia, and H. Meng, “Towards Multi-Scale Style Control for Expressive Speech Synthesis,” in *Proc. Interspeech 2021*, 2021, pp. 4673–4677.
- [13] M. Jeong, H. Kim, S. J. Cheon, B. J. Choi, and N. S. Kim, “Diff-tts: A denoising diffusion model for text-to-speech,” *proceedings of INTERSPEECH*, 2021.
- [14] D. Min, D. B. Lee, E. Yang, and S. J. Hwang, “Meta-stylespeech: Multi-speaker adaptive text-to-speech generation,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 7748–7759.
- [15] K. Lee, K. Park, and D. Kim, “Styler: Style factor modeling with rapidity and robustness via speech decomposition for expressive and controllable neural text to speech,” in *INTERSPEECH 2021*. International Speech Communication Association, 2021, pp. 3431–3435.
- [16] R. Huang, Y. Ren, J. Liu, C. Cui, and Z. Zhao, “Generspeech: Towards style transfer for generalizable out-of-domain text-to-speech synthesis,” *arXiv preprint arXiv:2205.07211*, 2022.
- [17] D. Yang, S. Liu, J. Yu, H. Wang, C. Weng, and Y. Zou, “Nore-speech: Knowledge distillation based conditional diffusion model for noise-robust expressive tts,” *arXiv preprint arXiv:2211.02448*, 2022.
- [18] Y. Wang, D. Stanton, Y. Zhang, R.-S. Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 5180–5189.
- [19] C. Qiang, P. Yang, H. Che, X. Wang, and Z. Wang, “Style-label-free: Cross-speaker style transfer by quantized vae and speaker-wise normalization in speech synthesis,” in *2022 13th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, 2022, pp. 61–65.
- [20] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [21] Y. Ren, X. Tan, T. Qin, Z. Zhao, and T.-Y. Liu, “Revisiting over-smoothness in text to speech,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 8197–8213.
- [22] H. Shao, S. Yao, D. Sun, A. Zhang, S. Liu, D. Liu, J. Wang, and T. Abdelzaher, “Controlvae: Controllable variational autoencoder,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 8655–8664.
- [23] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265.
- [24] 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 *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 2595–2605.
- [25] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro, “Diffwave: A versatile diffusion model for audio synthesis,” in *International Conference on Learning Representations*, 2021.
- [26] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-vectors: Robust dnn embeddings for speaker recognition,” in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [27] M. Cohen, G. Quispe, S. Le Corff, C. Ollion, and E. Moulines, “Diffusion bridges vector quantized variational autoencoders,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 4141–4156.
- [28] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, “LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech,” in *Proc. Interspeech 2019*, 2019, pp. 1526–1530. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2019-2441>
- [29] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [30] 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.
- [31] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, “Neural discrete representation learning,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 6309–6318.
- [32] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “Gans trained by a two time-scale update rule converge to a local nash equilibrium,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 6629–6640.