



PromptStyle: Controllable Style Transfer for Text-to-Speech with Natural Language Descriptions

Guanghou Liu¹, Yongmao Zhang¹, Yi Lei¹, Yunlin Chen², Rui Wang², Zhifei Li², Lei Xie^{1*}

¹Audio, Speech and Language Processing Group (ASLP@NPU)
School of Computer Science, Northwestern Polytechnical University, Xi'an, China
²Shanghai Mobvoi Information Technology Co., Ltd

ghliu@mail.nwpu.edu.cn, yunlinchen@mobvoi.com, lxie@nwpu.edu.cn

Abstract

Style transfer TTS has shown impressive performance in recent years. However, style control is often restricted to systems built on expressive speech recordings with discrete style categories. In practical situations, users may be interested in transferring style by typing text descriptions of desired styles, without the reference speech in the target style. The text-guided content generation techniques have drawn wide attention recently. In this work, we explore the possibility of controllable style transfer with natural language descriptions. To this end, we propose *PromptStyle*, a text prompt-guided cross-speaker style transfer system. Specifically, *PromptStyle* consists of an improved VITS and a cross-modal style encoder. The cross-modal style encoder constructs a shared space of stylistic and semantic representation through a two-stage training process. Experiments show that *PromptStyle* can achieve proper style transfer with text prompts while maintaining relatively high stability and speaker similarity. Audio samples are available in our demo page¹.

Index Terms: text-to-speech, style transfer, style prompt

1. Introduction

Text-to-speech (TTS) [1, 2] aims to produce human-like speech from input text. Recent progress in deep learning approaches has greatly improved the naturalness of speech [3]. With the wide applications of TTS in real-world human-computer interaction, expressive TTS with diverse styles attracts more attention. Generating stylistic speech for a specific speaker intuitively needs the same speaker's high-quality expressive speech recordings, which incurs a high cost for data collection. To solve the problem of synthesizing expressive speech for the target speaker without diverse speaking styles, cross-speaker style transfer [4, 5, 6, 7, 8] is a feasible solution.

For the style representations in style transfer scenarios, existing works mainly include two different methods, i.e. the pre-defined style id category index [9, 10, 11, 12] and hidden variables extracted from the reference signal [13, 14, 15, 16, 17]. However, the id-based methods are limited to the styles of fixed discrete categories, which leads to less flexibility. Although the reference-based methods can produce various speech through different references, the extracted style representation is not interpretable. Moreover, in practical applications, it is difficult to accurately and conveniently select an appropriate reference for arbitrary textual content.

With the success of text and image generation from prompt descriptions [18, 19], some prompt-based TTS methods are pro-

posed to improve the expressiveness and naturalness of synthetic speech. Style-Tagging-TTS (ST-TTS) [20] produces expressive speech based on style tags, which are stylistic words or phrases labeled from audiobook datasets. But it is difficult to describe complex styles by a single word or phrase, and the style tags are hard to label in common audiobook datasets as most utterances may be recorded in a less-expressive reading style. *PromptTTS* [21] proposes to use a style prompt from five different factors (i.e. gender, pitch, speaking speed, volume, and emotion) to guide the style expression for the generated speech. The recent *InstructTTS* [22] can synthesize stylistic speech with the guidance of natural language descriptions without formal constraints as style prompts. It's a three-stage training approach to capture semantic information from natural language style prompts as conditioning to the TTS system. Intuitively, natural language description is a convenient and user-friendly way to describe the desired style since no prior acoustic knowledge is required.

In this study, we focus on style transfer in the audiobook generation, where the target speaker has little expressive data and no style description prompts. Through other expressive data and transferring diversified styles with natural language descriptions, expressive audiobook speech can be generated for the target speaker. Specifically, this paper proposes to leverage natural language description prompts to transfer style from the source speaker to the target speaker who has no expressive speech data. The proposed approach - *PromptStyle* - has a two-stage procedure utilizing text prompts to model the style appearance. In the first stage, based on an improved VITS [23], we use a style encoder to extract a style hidden representation from reference speech as the condition of the TTS system. Multi-speaker multi-style expressive data without style annotation is involved in this stage to achieve cross-speaker style transfer through diverse reference speech. In the second stage, we design a prompt encoder to model the style embedding from the style prompt. Expressive speech data with style annotations in natural language descriptions is involved to fine-tune the pre-trained language model and TTS acoustic model, capturing the relationship between prompt embedding and style embedding space. Due to the generalization capability of the language model, style transfer from unseen prompts is also feasible.

We summarize the contributions of *PromptStyle* as follows.

- We propose a two-stage TTS approach for cross-speaker style transfer with natural language descriptions, which is more user-friendly and controllable than previous works.
- The proposed two-stage approach first uses a large amount of data without annotations to train a reference-based style transfer TTS model, and then leverages only a small amount of labeled data with style prompts to fine-tune a prompt en-

* Corresponding author.

¹<https://PromptStyle.github.io/PromptStyle>

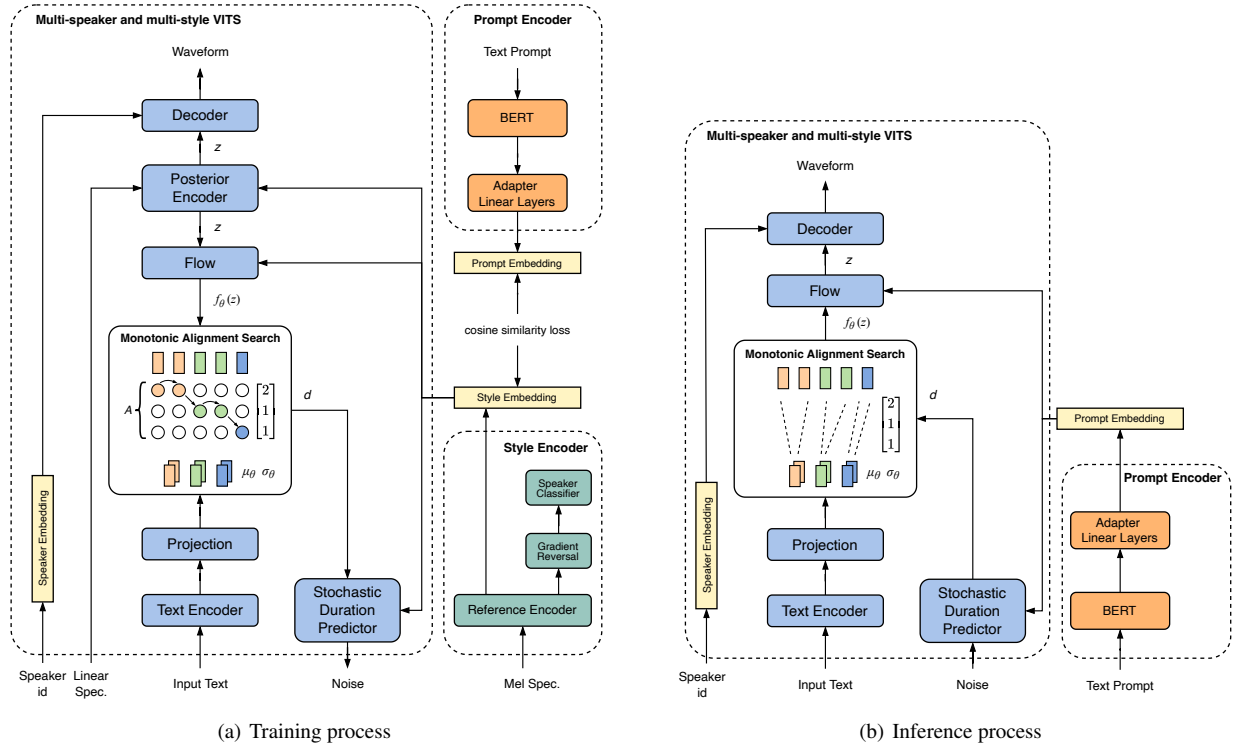


Figure 1: Architecture of PromptStyle

coder and the acoustic model.

- With the generalization capability of the pre-trained language model, our approach can generate stylistic speech for the target speaker from an unseen style prompt.

Experiments show that our proposed method can conduct cross-speaker style transfer with flexible text prompts and achieve high speech quality.

2. PromptStyle

2.1. Model Overview

The overall architecture of the model we proposed is shown in Fig 1. The proposed model consists of three parts, including a style encoder, a prompt encoder, and a TTS system. Our TTS framework is based on VITS[x], an end-to-end TTS model that directly converts the phoneme sequence to the speech waveform. VITS can generate high-quality speech attributed to its end-to-end learning manner without the possible mismatch between the acoustic model and the vocoder in the conventional paradigm. The style encoder is used to extract style information from reference speech and the prompt encoder is used to extract semantic information from natural language descriptions.

Based on the VITS framework, we propose a two-stage training procedure to achieve style transfer with text prompts. In the first stage, multi-style expressive data with speaker annotations are used to train a style transfer system based on VITS with the style encoder, aiming to learn a style embedding space. In the second stage, we use expressive speech with natural language description prompts to fine-tune the whole model to construct the relationship between the semantic embedding space and the style embedding space, so that we can achieve style transfer with text prompts.

2.2. Cross-modal Style Encoder

We introduce a cross-modal style encoder consisting of a style encoder and a prompt encoder for constructing a shared space of stylistic and semantic representation.

The style encoder is used to extract the style information from Mel-spectrograms but not linear spectrograms, because linear spectrograms contain too much information which may prevent capturing the style information accurately. The style encoder consists of a reference encoder and an auxiliary adversarial component. The reference encoder almost has a similar structure to the Global Style Tokens (GSTs) [24] discarding style token. The auxiliary adversarial component consists of a gradient reversal layer, feed-forward layers, and a speaker classifier, to disentangle the speaker and style information. The style encoder aims to construct a style embedding space for style transfer from large amounts of expressive speech without style annotations.

A prompt encoder is used to build a prompt embedding space for controllable style transfer with natural language descriptions. The prompt encoder consists of a pre-trained BERT [25] and an adapter layer. First, the text prompt is fed into the BERT model to extract semantic features. Then the adapter layer converts the semantic features into the prompt embedding of the same size as style embedding from the style encoder, and captures the relationship between the prompt embedding and the style embedding space.

2.3. Training and Inference

We introduce a two-stage training procedure to obtain a style transfer model with text prompts. This bi-modal training approach has the following two advantages. 1) By establishing the relationship between the style and prompt spaces, either reference speech or text descriptions can be used to control the style transfer. 2) The type and scale of the annotated data can




Figure 2: Visualization of the style embeddings from different models – (a) *CST-TTS*, (b) *GST-MLTTS* and (c) *PromptStyle*.

be determined specifically according to the actual domain requirements, and the fine-tuning process is time-saving.

2.3.1. Stage 1: Style transfer by reference Mel

In the first stage, a style transfer system with reference speech is developed. We particularly study the role of each main module in VITS for the style transfer task. The text encoder processes the input phonemes, independent of the speaker and style information. The normalizing flow improves the flexibility of the prior distribution. The stochastic duration predictor estimates the distribution of phoneme duration from text embedding. We believe that duration is an important aspect of style, so only style embedding is added to the input of the stochastic duration predictor. For simplicity, the normalizing flow is designed to be a volume-preserving transformation with the Jacobian determinant of one. So we only add style embedding to the normalizing flow through global conditioning, considering that it is hard to take into account both speaker and style embedding in the construction of the prior distribution. The posterior encoder produces the normal posterior distribution from a linear spectrogram and the decoder reconstructs the latent variables z to spectrogram. Referring to DelightfulTTS [26], SSIM loss L_{ssim} is added, and the speaker classifier in the adversarial component is trained with a cross-entropy loss with gradient inversion L_{adv} . Specifically, the loss of the first stage is

$$L_{stage1} = L_{vits} + L_{ssim} + L_{adv}. \quad (1)$$

2.3.2. Stage 2: Style transfer by text prompt

The second stage aims to achieve controllable style transfer with text prompts. To this end, we use a prompt encoder to extract prompt embeddings from natural language descriptions, and the prompt embedding learns style embedding by cosine similarity loss referring to L_{cons} . The cosine similarity loss is not only used by a separate optimizer for the prompt encoder but also used to train the TTS system. Note that, we only train the last attention layer of BERT, and the parameters of the style encoder are frozen throughout this training process. The loss of the second stage is

$$L_{stage2} = L_{vits} + L_{cons}. \quad (2)$$

3. Experiments

3.1. Datasets

We use 2 hours of an audiobook corpus recorded from a single speaker as the target speaker. Most recordings are common reading style without much expressive speech. As there

is no public expressive dataset with abundant natural language prompts, we use an internally expressive Mandarin Chinese speech corpus as source speakers. The corpus contains 12 hours of speech data from 8 female speakers and about 6 hours of speech is associated with natural language descriptions. To obtain the natural language descriptions (text prompts), we invite 9 professional annotators to use a phrase or a sentence to describe the style of each utterance. They are told to focus on the speaking style and ignore the linguistic content.

For our experiments, all the speech recordings are down-sampled to 24 kHz. Linear spectrograms obtained from raw waveforms through the Short-time Fourier transform (STFT) are used as the input of VITS. Eighty-dimensional Mel-frequency spectrograms are extracted with 12.5 ms frame shift and 50 ms frame length. The text prompts are tokenized by a pre-trained BERT.

3.2. Model Configuration

Our work focuses on style transfer with text prompts based on the audiobook datasets. Thus two models of style transfer on audiobook datasets are established as baselines – one with text prompt and the other using reference speech.

- **CST-TTS**: Following [27], a cross-speaker style transfer framework is built based on non-autoregressive feedforward structure which can control the expressiveness by style tags.
- **GST-MLTTS**: A cross-speaker emotion transfer model [28] based on semi-supervised training and SCLN.
- **PromptStyle**: The proposed style transfer system. The hyperparameters of VITS are set to be the same as in [23]. The reference encoder has the same architecture as GSTs without style tokens. The prompt encoder uses a pre-trained BERT to extract 768-dimensional global semantic embedding, and an adapter layer consisting of three linear layers with ReLU activation converting the semantic embedding to 256-dimensional prompt embedding.

All models are trained with a batch size of 32 for 200k steps, while the fine-tuning stage of *StylePrompt* is trained with 16 batches for 10k steps. HiFi-GAN [29] is used in *CST-TTS* and *GST-MLTTS* as the neural vocoder for reconstructing the waveform.

3.3. Performance on the Style Transfer

3.3.1. Subjective Evaluation

The performances of three comparison models on style transfer are evaluated in terms of speech quality, speaker similarity, and style similarity with a Mean Opinion Score (MOS) test.

Table 1: Similarity/Quality MOS results (Average score and 95% confidence interval)

Models	Speakers Similarity	Style Similarity	Speech Quality
CST-TTS	3.01	3.23	3.14
GST-MLTTS	3.46	3.42	3.53
PromptStyle	3.63	3.78	4.06

Table 2: The setup of the ablation experiments.

Models	Posterior Encoder		Flow		Decoder	
	style	speaker	style	speaker	style	speaker
M1	✓	✓	✓	✓	✓	✓
M2	✓	✓	✓		✓	✓
PromptStyle	✓	✓	✓			✓
M3		✓	✓			✓

Specifically, 50 speech samples are synthesized for each model, and 15 listeners who are native Chinese take part in this experiment. They are asked to score on a 1~5 point with 0.5 intervals for the speech quality, speaker similarity with the target speaker, and style similarity with the source speaker. As shown in Table 1, PromptStyle has clearly better performance in terms of speech quality, speaker similarity, and style similarity. We believe such good performance is due to the end-to-end structure and the style embedding as well as the speaker embedding added to different modules as conditions.

3.3.2. Objective evaluation

To visualize style embedding space by t-SNE [30], We manually selected 13 common styles from the training dataset for evaluation. Style embedding is extracted from the selected speech using PromptStyle, ST-TTS, and GST-MLTTS respectively. As shown in Fig 2, each point indicates a style embedding and points with the same color are from the same style category. The distance between the two points indicates the relative similarity of the embeddings. Smaller distances indicate more similar embeddings. An ideal style encoder should make the style embeddings from the same category closer to each other while embeddings from different categories should be separated apart. From Fig 2, we observe that MST-TTS gets the worst performance and GST-MLTTS performs better but the points from the same category are still diversely distributed. In contrast, in Fig. 2(c), the style embeddings from the same category are clustered together, and different clusters are apart from each other. It indicates that PromptStyle presents significant superiority in extracting style embeddings.

3.4. Ablation Study

In order to verify the influence of the style embedding and speaker embedding added to different modules as conditions, we conduct the ablation experiments considering the posterior encoder, the normalizing flow, and the decoder. Intuitively, we add both style embeddings and speaker embeddings to all their modules, referring to **M1**. Then we remove speaker embeddings from the flow, referring to **M2**. On the base of M2, removing style embeddings from the decoder is our final version of StylePrompt. Finally, on the base of StylePrompt, we remove the style embeddings from the posterior encoder, referring to **M3**. Similarly, these models are evaluated in terms of speech

Table 3: Similarity/Quality MOS results of ablations (Average score and 95% confidence interval)

Models	Speakers Similarity	Style Similarity	Speech Quality
M1	3.96	2.87	4.07
M2	3.43	3.78	3.98
PromptStyle	3.62	3.82	4.05
M3	3.57	3.73	4.01




Figure 3: ABX preference results

quality, speaker similarity, and style similarity with a MOS test. The model setups and experimental results are shown in Table 2 and Table 3, respectively.

As shown in Table 3, M1 maintains well the timbre of the target speaker but fails to achieve decent style transfer. Comparing M2 and M1, the results indicate that it is hard for a normalizing flow to extract a proper style representation with both style embedding and speaker embedding as conditions. Comparing PromptStyle with M2, it is verified that PromptStyle can capture more speaker timbre when the speaker embedding is added to the decoder of VITS. Finally, comparing M3 with PromptStyle, we find that adding the style embedding to the posterior encoder as conditions is beneficial to improve the expressiveness of the synthesized speech.

3.5. Style Control by Text Prompt

To evaluate the performance of PromptStyle in style transfer by text prompts, we further conduct ABX preference tests between the prompt embedding and the style embedding in terms of speech quality, speaker similarity, and style similarity. Note that style embedding serves as the performance top line, as speech reference contains more direct style information than the text prompt. Precisely, we extract style embedding and prompt embedding from the reference speech and the text prompt, respectively, to control the transferred style. Listeners are offered a pair of randomly selected samples and asked to choose which one is as expressive as the descriptions of the provided text prompts. As the preference test results shown in Fig. 3, the listeners give higher preference to *No preference*, which means that our proposed method achieves the goal of text prompt guided style transfer with precise style control and high speech quality.

4. Conclusions

In this work, we propose PromptStyle, a style transfer model with natural language prompts, to make the style transfer more controllable and user-friendly. PromptStyle establishes the relationship between the style embedding space and the prompt embedding space through two-stage training. Experiments on expressive audiobook synthesis show that PromptStyle can achieve the goal of style transfer with text prompts while maintaining relatively high stability and speaker similarity.

5. References

- [1] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio *et al.*, “Tacotron: Towards end-to-end speech synthesis,” *Proc. Interspeech 2017*, pp. 4006–4010, 2017.
- [2] 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 *International Conference on Learning Representations*, 2020.
- [3] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, “A survey on neural speech synthesis,” *arXiv preprint arXiv:2106.15561*, 2021.
- [4] Y. Bian, C. Chen, Y. Kang, and Z. Pan, “Multi-reference tacotron by intercross training for style disentangling, transfer and control in speech synthesis,” *CoRR*, vol. abs/1904.02373, 2019.
- [5] T. Li, X. Wang, Q. Xie, Z. Wang, and L. Xie, “Cross-speaker emotion disentangling and transfer for end-to-end speech synthesis,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1448–1460, 2022.
- [6] G. Zhang, Y. Qin, W. Zhang, J. Wu, M. Li, Y. Gai, F. Jiang, and T. Lee, “iemotts: Toward robust cross-speaker emotion transfer and control for speech synthesis based on disentanglement between prosody and timbre,” *arXiv preprint arXiv:2206.14866*, 2022.
- [7] T. Li, X. Wang, Q. Xie, Z. Wang, M. Jiang, and L. Xie, “Cross-speaker Emotion Transfer Based On Prosody Compensation for End-to-End Speech Synthesis,” in *Proc. Interspeech 2022*, 2022, pp. 5498–5502.
- [8] Z. Shang, Z. Huang, H. Zhang, P. Zhang, and Y. Yan, “Incorporating cross-speaker style transfer for multi-language text-to-speech,” in *Interspeech*, 2021, pp. 1619–1623.
- [9] H. Li, Y. Kang, and Z. Wang, “Emphasis: An emotional phoneme-based acoustic model for speech synthesis system,” *Proc. Interspeech 2018*, pp. 3077–3081, 2018.
- [10] X. Zhu, S. Yang, G. Yang, and L. Xie, “Controlling emotion strength with relative attribute for end-to-end speech synthesis,” *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 192–199, 2019.
- [11] N. Tits, F. Wang, K. E. Haddad, V. Pagel, and T. Dutoit, “Visualization and interpretation of latent spaces for controlling expressive speech synthesis through audio analysis,” *arXiv preprint arXiv:1903.11570*, 2019.
- [12] N. Tits, K. El Haddad, and T. Dutoit, “Exploring transfer learning for low resource emotional tts,” in *Intelligent Systems and Applications: Proceedings of the 2019 Intelligent Systems Conference (IntelliSys) Volume 1*. Springer, 2020, pp. 52–60.
- [13] R. Skerry-Ryan, E. Battenberg, Y. Xiao, Y. Wang, D. Stanton, J. Shor, R. Weiss, R. Clark, and R. A. Saurous, “Towards end-to-end prosody transfer for expressive speech synthesis with tacotron,” in *international conference on machine learning*. PMLR, 2018, pp. 4693–4702.
- [14] 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,” *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6945–6949, 2019.
- [15] A. Kulkarni, V. Colotte, and D. Juvet, “Improving transfer of expressivity for end-to-end multispeaker text-to-speech synthesis,” *2021 29th European Signal Processing Conference (EUSIPCO)*, pp. 31–35, 2021.
- [16] Y. Jia, Y. Zhang, R. J. Weiss, Q. Wang, J. Shen, F. Ren, Z. Chen, P. Nguyen, R. Pang, I. L. Moreno *et al.*, “Transfer learning from speaker verification to multispeaker text-to-speech synthesis,” in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 2018, pp. 4485–4495.
- [17] T. Li, S. Yang, L. Xue, and L. Xie, “Controllable emotion transfer for end-to-end speech synthesis,” *2021 12th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pp. 1–5, 2021.
- [18] Z. Zhang, Y. Gu, X. Han, S. Chen, C. Xiao, Z. Sun, Y. Yao, F. Qi, J. Guan, P. Ke *et al.*, “Cpm-2: Large-scale cost-effective pre-trained language models,” *AI Open*, vol. 2, pp. 216–224, 2021.
- [19] P. H. Seo, A. Nagrani, A. Arnab, and C. Schmid, “End-to-end generative pretraining for multimodal video captioning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17 959–17 968.
- [20] M. Kim, S. J. Cheon, B. J. Choi, J. J. Kim, and N. S. Kim, “Expressive text-to-speech using style tag,” *arXiv preprint arXiv:2104.00436*, 2021.
- [21] Z. Guo, Y. Leng, Y. Wu, S. Zhao, and X. Tan, “Prompotts: Controllable text-to-speech with text descriptions,” *arXiv preprint arXiv:2211.12171*, 2022.
- [22] D. Yang, S. Liu, R. Huang, G. Lei, C. Weng, H. Meng, and D. Yu, “Instructtts: Modelling expressive tts in discrete latent space with natural language style prompt,” *arXiv preprint arXiv:2301.13662*, 2023.
- [23] 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.
- [24] Y. Wang, D. Stanton, Y. Zhang, R. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” in *Proc. ICML, 2018*, pp. 5180–5189.
- [25] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural speech synthesis with transformer network,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, no. 01, 2019, pp. 6706–6713.
- [26] Y. Liu, Z. Xu, G. Wang, K. Chen, B. Li, X. Tan, J. Li, L. He, and S. Zhao, “Delightfultts: The microsoft speech synthesis system for blizzard challenge 2021,” *arXiv preprint arXiv:2110.12612*, 2021.
- [27] Y. Shin, Y. Lee, S. Jo, Y. Hwang, and T. Kim, “Text-driven emotional style control and cross-speaker style transfer in neural tts,” *arXiv preprint arXiv:2207.06000*, 2022.
- [28] P.-F. Wu, Z. Ling, L. Juan Liu, Y. Jiang, H.-C. Wu, and L.-R. Dai, “End-to-end emotional speech synthesis using style tokens and semi-supervised training,” *2019 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (AP-SIPA ASC)*, pp. 623–627, 2019.
- [29] 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.
- [30] V. D. M. Laurens and G. Hinton, “Visualizing data using t-SNE,” *Journal of Machine Learning Research*, vol. 9, no. 2605, pp. 2579–2605, 2008.