



Text-aware and Context-aware Expressive Audiobook Speech Synthesis

Dake Guo¹, Xinfa Zhu¹, Liუმeng Xue², Yongmao Zhang¹, Wenjie Tian¹, Lei Xie^{1,*}

¹Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, Xi'an, China

²School of Data Science, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), China

guodake@mail.nwpu.edu.cn, lxie@nwpu.edu.cn

Abstract

Recent advances in text-to-speech have significantly improved the expressiveness of synthetic speech. However, a major challenge remains in generating speech that captures the diverse styles exhibited by professional narrators in audiobooks without relying on manually labeled data or reference speech. To address this problem, we propose a text-aware and context-aware (TACA) style modeling approach for expressive audiobook speech synthesis. We first establish a text-aware style space to cover diverse styles via contrastive learning with the supervision of the speech style. Meanwhile, we adopt a context encoder to incorporate cross-sentence information and the style embedding obtained from text. Finally, we introduce the context encoder to two typical TTS models, VITS-based TTS and language model-based TTS. Experimental results demonstrate that our proposed approach can effectively capture diverse styles and coherent prosody, and consequently improves naturalness and expressiveness in audiobook speech synthesis¹.

Index Terms: audiobook speech synthesis, text-aware, context-aware, style modeling

1. Introduction

With the development of sequence-to-sequence (seq2seq) neural approaches [1, 2, 3], current text-to-speech (TTS) systems can generate human-like natural speech, with wide applications in voice assistant, navigation system, and audiobook production. Particularly, audiobook speech synthesis aims to synthesize expressive long-form speech from literary books, achieving efficient and cost-saving automated audio-content production. However, audiobook synthesis is more challenging due to the *rich speaking styles* performed by professional narrators, *context-aware expressiveness*, and *long-form prosody coherence* [4].

To synthesize speech with various styles for audiobooks, a straightforward method is to train a TTS model with large-scale style-labeled datasets [5, 6, 7]. However, this method heavily depends on the quality and precision of the pre-defined labels. To mitigate reliance on annotated data, some works have tried to capture a global style embedding from a given style reference speech rendering a desired style [8, 9, 10] and then generate speech with the corresponding style. Furthermore, CLAM [11] proposes to select multiple references focusing on the style-related information in the text. Other works have focused on decoupling a style representation from speech, achieving style controllable speech synthesis [12, 13, 14]. However, the styles derived from either labeled tags or reference speech are still limited, making it difficult to cover the wide range of rich styles in

audiobooks. Additionally, it is also impractical to manually assign a specific style tag or style reference for each sentence in an audiobook.

A more practical method for style extraction for speech synthesis is to predict the style embedding from text because there are only text can be accessible during inference. To achieve such text-adaptive style learning for speech synthesis, Text-Predicted Global Style Token (TP-GST) extends the capabilities of GST[8, 15] to generate style embeddings or style token weights based on textual input solely. CLAPSpeech [16] utilizes contrastive learning to learn the prosody variance of the same text under different contexts. However, it focuses on local prosody, such as word-level or even phoneme-level prosody. Even though TP-GST produces global style embedding, it relies on a limited number of style tokens, which is not enough for diverse style variance in audiobooks with abundant textual content.

Moreover, the style and prosody in each sentence are affected by its neighboring sentences in the case of long text in the book, and thus many methods improve the prosody consistency of the synthetic speech by learning contextual information [17]. Some of them use context-aware text embeddings of continuous multiple sentences to improve the prosody of speech [18, 19]. Other works leverage acoustic contexts to improve the prosody coherence and expressiveness of the current sentence [20]. Moreover, some studies utilize both textual and acoustic contextual knowledge to generate expressive long-form speech [21, 22]. Besides, given that the expressiveness of human speech is perceived as a composition of hierarchical factors rather than a singular scale, several studies leverage the hierarchical structure of discourse to further enrich the prosody variations in long-form synthetic speech [4, 23, 24, 25].

This paper proposes a text-aware and context-aware (TACA) style modeling approach for expressive audiobook speech synthesis. To address the issues of limited styles and achieve more practical text-adaptive style prediction, we propose a text-aware style modeling method to learn a style space encompassing a wide range of styles. Specifically, we employ a semi-supervised contrastive learning strategy to learn style directly from the text using a large number of datasets. This process is assisted by the style learned from speech, as it allows for the capture of more detailed style variations in speech. To further integrate context, a context encoder is crafted to extract cross-sentence information from adjacent sentences and to embed the learned style directly from the text, assisting TTS in generating speech with context-aware expressiveness and long-form coherent prosody. Since the context encoder is a plug-and-play, universal module, it can be readily integrated into most TTS frameworks. In this study, we build text-aware and context-aware (TACA) TTS models based on two different TTS

* Corresponding author.

¹Speech samples: <https://dukguo.github.io/TACA-TTS/>

frameworks, i.e., a classical VITS-based TTS and a popular language model-based TTS. Experimental results demonstrate that the text-aware style modeling method effectively captures the diverse and detailed style from text and the context encoder helps to learn coherent prosody.

2. Methodology

The overall architecture of our proposed method is shown in Figure 1. We first construct an extensive style space from speech and then help to establish text-aware style space using cross-modal supervision. Then we build text-aware and context-aware (TACA) TTS models that incorporate the context encoder with the style embedding obtained from the text, to introduce contextual information and diverse styles into TTS for improved audiobook speech synthesis in expressiveness and naturalness.

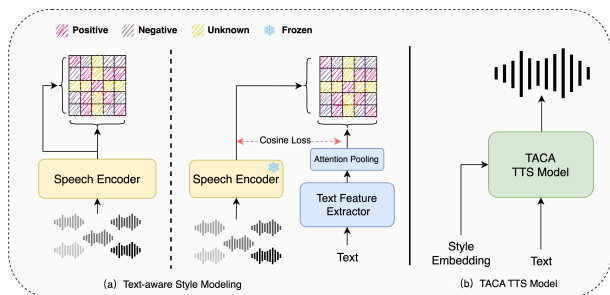


Figure 1: Overview of our proposed method

2.1. Text-aware Style Modeling

Given speech-text pairs P_1, P_2, \dots, P_n , where speech-text pair P_i contains speech S_i and text T_i , we aim to build a style space from text to speech. Considering the style expressions mainly expressed through speech, we first establish a speech style space and then help the text style space learning under cross-modal supervision. To capture rich style expressions in speech S_1, S_2, \dots, S_n , we utilize the speech encoder from SRL [26] to extract style representations. Specifically, this speech encoder leverages HuBERT [27] to handle various speech inputs and uses contrastive learning to capture style representation in a semi-supervision fashion. Compared to classifier-guided representation learning models that categorize styles into predefined classes, the style representations learned from stylistic differences between speech samples offer greater diversity, contributing to a more detailed and comprehensive style space.

Under the guidance of the established speech style space, we build the text-aware style space on a larger amount of datasets. Inspired by CLIP [28], we leverage contrastive learning to assist text-style space learning with the learned speech-style space. The key to contrastive learning lies in constructing positive and negative samples. We construct these samples for text-style learning based on the similarity between the learned speech representations. If the similarity between S_i and S_j is less than α , the pairs P_i and P_j are denoted as negative samples. Conversely, if the similarity between S_i and S_j is more than β , the pairs P_i and P_j are denoted as positive samples. Any pairs with similarity scores ranging between α and β are termed unknown samples, reflecting ambiguity in their positive or negative correlation. As shown in Figure 1 (b), we employ a powerful pre-trained text feature extractor, such as the T5 model, and attention pooling to obtain global style represen-

tations from T_1, T_2, \dots, T_n . Then we fine-tune the text feature extractor but freeze the speech encoder to learn the text-aware style via calculating the contrastive loss of P_1, P_2, \dots, P_n in a semi-supervised manner. Furthermore, we use cosine similarity loss on the style embeddings of the two modalities to accelerate convergence. Such a style modeling method not only learns rich and diverse styles via contrastive learning and semi-supervision modes but also makes the style text adaptive, facilitating the varied and expressive expression in audiobook speech synthesis where only text is available.

2.2. Context Encoder

We design a context encoder to integrate contextual information with style representations, as shown in Fig 2. The prosody of a given sentence can be influenced by its position within a discourse and the relationships with adjacent sentences. Considering this, we broaden the receptive field of the text input in TTS models from one sentence to a couple of sentences. It explicitly incorporates context from neighboring sentences. Furthermore, we introduce the style embedding obtained from the text-aware style space to the context encoder to encourage the TTS model to generate expressive speech. It is worth noting that we apply Vector Quantization (VQ) to the style embedding. It enables the model to effectively learn and capture the commonalities in style expression across a wide range of samples, generating more stable expressive speech.

To make full use of the available dataset including the dataset with and without contextual annotation which is the order of the utterance, we adopt a basic pre-training and *context-aware fine-tuning* training strategy to use the context encoder. Specifically, the pre-training stage is to train the model without contextual information, and the fine-tuning stage is to train the model with context information. During the fine-tuning process, style embeddings are randomly sourced from either speech or text to bolster the model's generalization.




Figure 2: The architecture of context encoder

2.3. Text-aware and Context-aware TTS

We apply the context encoder with text-aware style and contextual information to two typical types of TTS models, i.e., the VITS-based model and the language model-based model, to achieve expressive audiobook speech synthesis.

2.3.1. VITS-based TTS

We employ BERT-VITS 2² as one of the TTS models, acknowledged for its exceptional expressiveness in speech synthesis. Bert-VITS 2, building upon the VITS 2 [29] architecture, incorporates BERT embeddings, enriching the synthetic speech with semantic information and significantly improving its prosody. We replaced the original text encoder with the con-

²<https://github.com/fishaudio/Bert-VITS2>

text encoder, achieving a larger receptive field of context to better model long-form speech, as shown in Figure 3. We apply basic pre-training and *context-aware fine-tuning* training strategy to obtain a Base-VITS and a TACA-VITS, where the TACA-VITS is augmented with rich style and contextual information for more expressive speech synthesis.

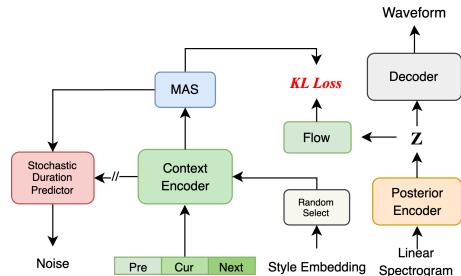


Figure 3: Text-aware and context-aware VITS-based TTS

2.3.2. Language model-based TTS

Leveraging language models for speech synthesis based on discrete tokens has evolved into a new research paradigm. We build a language model-based TTS as another TTS model, as shown in Figure 4, which is similar to AR-VITS³. As shown in the right part of Figure 4, it begins to extract audio features using the Hubert model with WN Blocks [3], followed by the application of VQ for semantic tokens, and the tokens are transformed into waveforms using VITS, which is called Hubert-VITS. In the Hubert-VITS model, the input consists of frame-level tokens, thus eliminating the necessity for duration modeling to stretch the input to a frame-level representation. The left part of Figure 4 is an autoregressive (AR) language model, which aims to predict the semantic tokens from the text. In this work, we first train the Hubert-VITS model using high-quality audio samples, aiming to enhance its waveform reconstruction capabilities. Subsequently, we pre-train a transform-based language model to predict semantic tokens on a large dataset because language models perform better on big datasets. After that, we apply *context-aware fine-tuning* (TACA-LM) to the pre-trained language model (Base-LM) to encourage the language model with a larger receptive field of context and more detailed style information, improving the expressiveness of the synthesized speech.

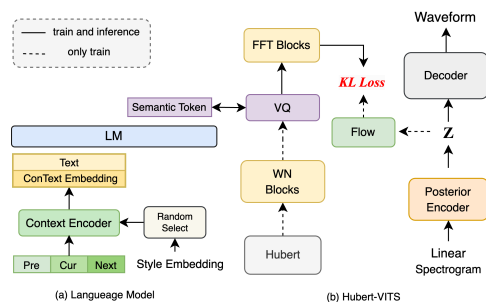


Figure 4: Text-aware and context-aware LM-based TTS

³<https://github.com/inmny/ar-vits>

3. Experimental Setup

3.1. Datasets and Preprocessing

Our validation methodology for the framework involves the application of audiobooks and an expressive speech corpus. There are three datasets involved in experiments.

- **20H-Audiobook-HQ:** A Mandarin audiobook dataset with utterance order annotated contains about 20 hours of speech with the corresponding text.
- **100H-Multi-Style:** A multi-style Mandarin dataset contains about 100 hours of speech with the corresponding text, of which 66 hours are labeled with the style tags.
- **6kH-Audiobook:** An audiobook dataset from the Internet contains about 6000 hours of speech without text. We transcribe the audio to text using Paraformer [30].

We select 5 chapters, about 100 audio clips from 20H-Audiobook-HQ as the test set.

3.2. Training Setup and Model Details

During text-aware style modeling, we use the Chinese-Hubert-Large⁴ to extract features from speech with layer 6. We use the same configuration as SRL in Speech Encoder. For text feature extractor, we use a pre-trained Chinese T5, Randeng-T5-784M⁵. The α and β in text-aware style modeling are set to 0.60 and 0.95, respectively. The dimension of style embedding is set to 384. All embeddings from either Speech Encoder or T5 undergo L2 normalization before calculating loss. We train the Speech Encoder on 100H-Multi-Style with a batch size of 96 and T5 on 6kH-Audiobook with a batch of 64.

For TACA-VITS, the model configuration is consistent with the original repository, which is the version of 2.0. The Context Encoder takes contextual phonemes and aligned Bert embeddings as input. The vector quantization of style in Context Encoder uses a codebook with 64×32 . We train the base model using 100H-Multi-Style with a batch size of 48. When *context-aware fine-tuning*, we train the model in the same setup.

For TACA-LM, we extract semantic tokens with vector quantization from layer 9 of the Chinese-Hubert-Base. The codebook size and dim is set to 1024 and 128, respectively. The configuration of VITS is the same as the original paper. The Hubert-VITS is trained with a batch size of 48 using 300 hours of high-quality speech selected from 6kH-Audiobook. We employ nanoGPT⁶ as the backbone of the language model, which contains 12 transformer layers with 12 heads. The input sequence applies only one position embedding, and its embedding dimension is set to 768. Different from traditional AM, we employ text as input rather than phonemes. The language model is trained on 6kH-Audiobook with a batch size of 24. Also, 10 times gradient accumulation is used during training. The context encoder is trained on the same configuration as TACA-VITS.

4. Experimental Results

To ascertain the efficacy of the style embedding learned via cross-modal supervision, we first conduct a detailed analysis of the style space from the text style encoder. Subsequently, we further investigate the effectiveness of the two TTS models on

⁴https://github.com/TencentGameMate/chinese_speech_pretrain

⁵<https://huggingface.co/IDEA-CCNL/Randeng-T5-784M>

⁶<https://github.com/karpathy/nanoGPT>

audiobook synthetic speech in expressiveness and naturalness compared with vanilla TTS models without context encoder.




Figure 5: *t*-SNE visualization of the style representation extracted from the text in 9 different chapters.

4.1. Style Space Analysis

The speaking style is mainly expressed through speech. In the text-speech paired dataset, the style predicted from the text should be aligned with that of the corresponding speech. To verify whether the style space built by our cross-modal style modeling method is shared between text and speech, we calculate the cosine similarity between the style embeddings extracted from speech and text on 1000 sentences. The higher cosine similarity of 0.93 indicates that our method is effective in aligning the text to the speaking style in the speech. Moreover, the T5 encoder used for text feature extraction is more powerful than BERT which obtains the cosine similarity of 0.82, benefitting to the style embedding predicted from text and then resulting in expressive generated speech.

We further visualize the style representations predicted from the text of 9 chapters, including 450 sentences in total. The *t*-SNE visualization result is presented in Figure 5, in which points with the same color are derived from the same chapters. From the visualization, we can see that (1) The style embedding of text from the same chapter tends to cluster a group, indicating that our method captures a relatively consistent style within a chapter. It helps to produce coherent prosody in audiobook speech synthesis. (2) The points in a cluster group are relatively dispersed, which means the nuance in style among sentences within a chapter can be learned deliberately. (3) The points in the whole style space are diversely distributed, showing that the style space covers diverse styles. It benefits audiobook speech synthesis with extensive expressiveness.

4.2. Objective Evaluation

We evaluate the performance of TTS models on the reserved test set with character error rate (CER) and mel-cepstral distortion (MCD). We use Paraformer [30] to evaluate CER and employ dynamic time warping (DTW) to align the predicted mel-spectrogram with the ground truth to calculate the MCD. The objective results are shown in Table 1, where VITS and LM refer to the vanilla methods, and both of them are fine-tuned on 6kH-Audiobook for a fair comparison. As we can see, LM is somewhat lacking in quality and stability compared to VITS reflected in higher MCD and CER. However, after context-aware fine-tuning of VITS (TACA-VITS), the overall changes in MCD and WER are not significant, indicating that our training strategy doesn't degrade model performance. For LM, applying the same fine-tuning strategy (TACA-LM) results in reduced MCD and CER, which implies that for LMs with stronger model-

ing capabilities, integrating conditions associated with semantics and style can slightly improve speech quality and stability. Compared to the two models of the VITS, the LM-based methods demonstrate higher CER, chiefly due to the inherent problems of autoregressive models, including missed readings and repetition.

Table 1: *Evaluations results for different TTS models.*

TTS Model	MCD ↓	CER(%) ↓	NMOS ↑	EMOS ↑
VITS	4.23	5.80	3.84±0.110	3.61±0.143
TACA-VITS	4.21	5.99	3.90±0.098	3.93±0.105
LM	8.92	13.9	2.91±0.087	3.80±0.112
TACA-LM	8.15	13.1	3.22±0.099	4.05±0.104

4.3. Subjective Evaluation

We evaluate synthetic audiobook speech in terms of naturalness and expressiveness through Mean Opinion Score (MOS) tests, denoted as EMOS and NMOS, respectively. NMOS assesses perceptual quality, including clarity and intelligibility, while EMOS focuses on prosodic expression, ignoring quality issues. To align closely with the audiobook scenario, we generate five audio clips with lengths ranging from 1 to 2 minutes constructed by combining multiple sentences for each MOS test. In each MOS test, a group of 20 native Chinese Mandarin listeners is asked to listen to synthetic speech and rate on a scale from 1 to 5 with a 0.5-point interval.

The MOS results presented in Table 1 show that our proposed method achieves improvements in naturalness and expressiveness on both two TTS models. The observed gains in NMOS and EMOS metrics underscore the effectiveness of our proposed method, particularly highlighting the pivotal role of the context encoder. After incorporating context information, TACA-VITS shows a slight improvement over VITS in naturalness, and TACA-LM also exhibits an improvement in naturalness. It indicates that context-aware knowledge helps to coherent prosody and results in more natural speech. Additionally, TACA-VITS demonstrates a significant improvement in expressiveness, and TACA-LM also sees a notable improvement in expressiveness. The contextual information and text-adaptive style guidance bring context-aware expressiveness and rich styles, producing more expressive audiobook speech. It is noted that the VITS models demonstrate superior performance over LM models in naturalness, which can be attributed to the inherent backdraws of autoregressive inference in LM, such as repetition, omission, and unclear pronunciation. However, LM shows great potential in expressiveness owing to its excellent capabilities of semantic comprehension and context-aware style expression.

5. Conclusions

This paper proposes a novel method for expressive audiobook speech synthesis through text-aware and context-aware style modeling. We first establish a text-aware style space to cover diverse styles with the supervision of speech style space. Furthermore, we design a context encoder to incorporate contextual information in cross-sentence and embed style embeddings derived from text. Since text-aware style modeling and context-aware information learning methods are designed universally, they can be integrated into most TTS frameworks. We build VITS-based and LM-based TTS systems based on the two modules. Experimental results show that our method effectively captures diverse styles and coherent prosody, improving naturalness and expressiveness in audiobook speech synthesis.

6. References

- [1] Y. Wang, R. J. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. V. Le, Y. Agiomyrgiannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.
- [2] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "Fastspeech: Fast, robust and controllable text to speech," in *Proc. NeurIPS*, 2019, pp. 3165–3174.
- [3] J. Kim, J. Kong, and J. Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *Proc. ICML*, 2021, pp. 5530–5540.
- [4] D. Guo, X. Zhu, L. Xue, T. Li, Y. Lv, Y. Jiang, and L. Xie, "HIGNN-TTS: Hierarchical Prosody Modeling With Graph Neural Networks for Expressive Long-Form TTS," in *Proc. ASRU*, 2023, pp. 1–7.
- [5] H. Choi, S. Park, J. Park, and M. Hahn, "Multi-speaker emotional acoustic modeling for cnn-based speech synthesis," in *Proc. ICASSP*, 2019, pp. 6950–6954.
- [6] T. Li, S. Yang, L. Xue, and L. Xie, "Controllable emotion transfer for end-to-end speech synthesis," in *Proc. ISCSLP*, 2021, pp. 1–5.
- [7] R. Liu, B. Sisman, G. Gao, and H. Li, "Expressive TTS training with frame and style reconstruction loss," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1806–1818, 2021.
- [8] Y. Wang, D. Stanton, Y. Zhang, R. J. Skerry-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 *Proc. ICML*, 2018, pp. 5167–5176.
- [9] Y. Bian, C. Chen, Y. Kang, and Z. Pan, "Multi-reference tacotron by intercross training for style disentangling, transfer and control in speech synthesis," *arXiv preprint arXiv:1904.02373*, 2019.
- [10] D. Min, D. B. Lee, E. Yang, and S. J. Hwang, "Meta-stylespeech: Multi-speaker adaptive text-to-speech generation," in *Proc. ICML*, 2021, pp. 7748–7759.
- [11] Y. Meng, X. Li, Z. Wu, T. Li, Z. Sun, X. Xiao, C. Sun, H. Zhan, and H. Meng, "CALM: contrastive cross-modal speaking style modeling for expressive text-to-speech synthesis," in *Proc. Interspeech*, 2022, pp. 5533–5537.
- [12] X. Zhu, Y. Lei, K. Song, Y. Zhang, T. Li, and L. Xie, "Multi-speaker expressive speech synthesis via multiple factors decoupling," in *Proc. ICASSP*, 2023, pp. 1–5.
- [13] R. Fu, J. Tao, Z. Wen, J. Yi, T. Wang, and C. Qiang, "Bi-level style and prosody decoupling modeling for personalized end-to-end speech synthesis," in *Proc. ICASSP*, 2021, pp. 6568–6572.
- [14] W. Song, Y. Yue, Y.-j. Zhang, Z. Zhang, Y. Wu, and X. He, "Multi-speaker multi-style speech synthesis with timbre and style disentanglement," in *Proc. NCMSC*, 2023, pp. 132–140.
- [15] D. Stanton, Y. Wang, and R. J. Skerry-Ryan, "Predicting expressive speaking style from text in end-to-end speech synthesis," in *Proc. SLT*, 2018, pp. 595–602.
- [16] Z. Ye, R. Huang, Y. Ren, Z. Jiang, J. Liu, J. He, X. Yin, and Z. Zhao, "CLAPSpeech: Learning prosody from text context with contrastive language-audio pre-training," in *Proc. ACL*, 2023, pp. 9317–9331.
- [17] Y.-J. Zhang, W. Song, Y. Yue, Z. Zhang, Y. Wu, and X. He, "Maskedspeech: Context-aware speech synthesis with masking strategy," *arXiv preprint arXiv:2211.06170*, 2022.
- [18] G. Xu, W. Song, Z. Zhang, C. Zhang, X. He, and B. Zhou, "Improving prosody modelling with cross-utterance bert embeddings for end-to-end speech synthesis," in *Proc. ICASSP*, 2021, pp. 6079–6083.
- [19] Y.-J. Zhang, C. Zhang, W. Song, Z. Zhang, Y. Wu, and X. He, "Prosody modelling with pre-trained cross-utterance representations for improved speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, pp. 1–11, 2023.
- [20] P. O. Gallegos and S. King, "Using previous acoustic context to improve text-to-speech synthesis," vol. abs/2012.03763, 2020.
- [21] S. Lei, Y. Zhou, L. Chen, Z. Wu, S. Kang, and H. Meng, "Context-aware coherent speaking style prediction with hierarchical transformers for audiobook speech synthesis," in *Proc. ICASSP 2023*, 2023, pp. 1–5.
- [22] D. Xin, S. Adavanne, F. Ang, A. Kulkarni, S. Takamichi, and H. Saruwatari, "Improving speech prosody of audiobook text-to-speech synthesis with acoustic and textual contexts," in *Proc. ICASSP*, 2023, pp. 1–5.
- [23] L. Xue, F. K. Soong, S. Zhang, and L. Xie, "ParaTTS: Learning linguistic and prosodic cross-sentence information in paragraph-based TTS," *IEEE ACM Trans. Audio Speech Lang. Process.*, pp. 2854–2864, 2022.
- [24] S. Lei, Y. Zhou, L. Chen, Z. Wu, S. Kang, and H. Meng, "Towards expressive speaking style modelling with hierarchical context information for mandarin speech synthesis," in *Proc. ICASSP 2022*, 2022, pp. 7922–7926.
- [25] S. Lei, Y. Zhou, L. Chen, J. Hu, Z. Wu, S. Kang, and H. Meng, "Towards multi-scale speaking style modelling with hierarchical context information for mandarin speech synthesis," in *Proc. ICASSP*, 2022, pp. 5523–5527.
- [26] X. Zhu, Y. Li, Y. Lei, N. Jiang, G. Zhao, and L. Xie, "Multi-speaker expressive speech synthesis via semi-supervised contrastive learning," *arXiv preprint arXiv:2310.17101*, 2023.
- [27] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [28] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *Proc. ICML*, 2021, pp. 8748–8763.
- [29] J. Kong, J. Park, B. Kim, J. Kim, D. Kong, and S. Kim, "VITS2: Improving Quality and Efficiency of Single-Stage Text-to-Speech with Adversarial Learning and Architecture Design," in *Proc. INTERSPEECH 2023*, 2023, pp. 4374–4378.
- [30] Z. Gao, S. Zhang, I. McLoughlin, and Z. Yan, "Paraformer: Fast and Accurate Parallel Transformer for Non-autoregressive End-to-End Speech Recognition," in *Proc. INTERSPEECH*, 2022, pp. 2063–2067.