HIGNN-TTS: HIERARCHICAL PROSODY MODELING WITH GRAPH NEURAL NETWORKS FOR EXPRESSIVE LONG-FORM TTS

Dake Guo[1], Xinfa Zhu[1], Liumeng Xue[2], Tao Li[1], Yuanjun Lv[1], Yuepeng Jiang[1], Lei Xie[1]*

[1]Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, Xi'an, China
[2] School of Data Science, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), China

ABSTRACT

Recent advances in text-to-speech, particularly those based on Graph Neural Networks (GNNs), have significantly improved the expressiveness of short-form synthetic speech. However, generating human-parity long-form speech with high dynamic prosodic variations is still challenging. To address this problem, we expand the capabilities of GNNs with a hierarchical prosody modeling approach, named HiGNN-TTS. Specifically, we add a virtual global node in the graph to strengthen the interconnection of word nodes and introduce a contextual attention mechanism to broaden the prosody modeling scope of GNNs from intra-sentence to inter-sentence. Additionally, we perform hierarchical supervision from acoustic prosody on each node of the graph to capture the prosodic variations with a high dynamic range. Ablation studies show the effectiveness of HiGNN-TTS in learning hierarchical prosody. Both objective and subjective evaluations demonstrate that HiGNN-TTS significantly improves the naturalness and expressiveness of long-form synthetic speech[1].

Index Terms— Expressive long-form TTS, hierarchical prosody modeling, graph neural network, semantic representation enhancing

1. INTRODUCTION

Text-to-Speech (TTS), aiming to generate human-like speech from text, has achieved dramatically advanced in naturalness with the proliferation of sequence-to-sequence (seq2seq) based neural approaches [1, 2, 3]. With the growing demand for anthropomorphic human-computer interaction, there has been increasing interest in generating speech with high expressiveness [4, 5, 6, 7], which exhibits high dynamic range in prosodic variation, including pitch and duration, etc. Human speech is expressive in nature, and proper expression rendering affects overall speech perception, which is essential for many TTS applications, such as newsreaders and voice assistants, etc. However, there is still a noticeable gap between

synthetic speech and human speech in terms of expressiveness, particularly in long-form speech synthesis scenarios like audiobooks. Conventional TTS usually works on the prosody modeling of a single sentence. However, long-form speech usually contains multiple semantically coherent sentences, where the prosody of each individual sentence is also affected by its corresponding context. Therefore, in long-form speech generation, besides ensuring the coherence of the overall rhythm, it is also necessary to model the local fine-grained and global prosody in each sentence as well as the cross-sentence contextual prosody [8].

To address sentence-level global prosody, several works have brought attention to the extraction of a global style embedding from a given reference speech which has successfully captured the global style features of speech [9, 10]. In order to synthesize expressive speech without the need of auxiliary reference speech during inference, some methods attempt to obtain prosodic variations directly from text, which is more practical. By incorporating text-only predictions, the Text-Predicted Global Style Token (TP-GST) [11] extends the capabilities of GST [10], enabling the generation of style embeddings or style token weights solely based on textual input.

However, only sentence-level global prosody modeling lacks local fine-grained prosody information within a sentence, such as pauses and emphasis. By aligning the extracted prosody embedding sequence to the phoneme sequence, some works manage to model fine-grained prosody [12, 13, 14]. Moreover, the tight interplay of prosody and semantic information within sentences has led to a growing focus on leveraging pre-trained language models (such as BERT [15], XLNET [16]) to enhance fine-grained prosody representation [17, 18]. Further, to fully utilize the tree-structured syntactic information, some works have leveraged graph-based methods to model word-level prosody. Some studies incorporate BERT embeddings and employ Relational Gated Graph Networks (RGGN) to extract semantic information, and consequently enriches prosodic variations and improves expressiveness in the generated speech [19]. Instead of using BERT embedding, other work utilizes the average pooling of hidden layer representations from the acoustic model as

---

* Corresponding author.
[1]Speech samples: https://anonymous-asru.github.io/HiGNN-TTS/

node features to enrich prosodic variation [20]. By capturing the intricate relations between linguistic units and prosodic cues, these methods facilitate the generation of speech with improved naturalness and expressiveness.

Considering that the prosody in each individual sentence is also affected by its neighboring sentences, some methods improve the naturalness and prosody consistency of the synthetic speech by harnessing contextual information. The method of using context-aware text embeddings obtained by encoding BERT embeddings of continuous multiple sentences was explored to improve the prosody of individual speech [21, 22]. Additionally, other methods leverage acoustic contextual information to improve the expressiveness of the sentence. One of work utilizes the preceding acoustic feature to improve the prosody coherence and expressiveness [23]. Furthermore, a few works use both textual context information and previous acoustic prosodic representations to estimate the style embedding of the single sentence [24], while others utilize separate text and audio context encoders to encode the respective contexts into global prosody representations [25].

Nevertheless, the aforementioned methods focus on modeling prosody on a single scale. While the expressiveness of human speech can be perceived as a compound of hierarchical factors rather than a mono scale. By employing linguistic-aware, prosody-aware, and sentence position networks, ParaTTS [26] achieves paragraph-level speech synthesis with smooth prosody and coherence. A few approaches incorporate a hierarchical context encoder to capture syllable structural features for both inter-phrase and inter-sentence contexts [27]. Another method, known as multi-scale style modeling, focuses on modeling context at cross-sentence, sentence, and subword levels [28]. They both obtain a global prosody representation from different scales, thereby improving the expressiveness of the synthetic speech.

Although tremendous effort has been made to improve expressiveness in both intra-sentence and inter-sentence scales, listeners still feel obvious fatigue with long-form synthetic speech. In this paper, we propose *HiGNN-TTS*, a graph-based hierarchical prosody modeling approach for long-form expressive speech synthesis. Specifically, we expand GNNs with *hierarchical* prosody modeling capacity covering word, sentence, and cross-sentence levels. First, we add a virtual global node in the syntax tree to improve the connections between word nodes. Second, we design a hierarchical graph encoder to explore the hierarchical structural relationship in the text at the word level, sentence level, and cross-sentence context level, which broadens the scope of graph-based speech synthesis from intra-sentence to inter-sentence. Third, we introduce prosody supervision signals from speech and propagate them to each word node of the graph through the message-passing capability of GNN, utilizing prosodic variations from the speech to guide the prosody modeling from the text. Ablation studies show that our proposed approach, HiGNN-TTS, is effective in modeling hierarchical prosody. Objective and subjective experiments demonstrate that HiGNN-TTS outperforms several competitive models, achieving more natural and expressive short-form and long-form synthetic speech.

2. METHODOLOGY

As illustrated in Fig.1, the overall architecture of HiGNN-TTS is a Fastspeech2-based model with a newly-introduced hierarchical graph prosody encoder and a pre-trained Mel-encoder. The hierarchical graph encoder models hierarchical prosody based on syntax graph. It takes in the graphs of previous, current, and next sentences and learns word-level, sentence-level, and cross-sentence context prosody representations. The pre-trained mel-encoder supervises the hierarchical prosody modeling via the acoustic prosody signals, helping capture a wide range of prosodic variations. As the backbone, the acoustic model produces a mel spectrogram from phoneme sequence input on the condition of hierarchical prosody, synthesizing expressive speech.
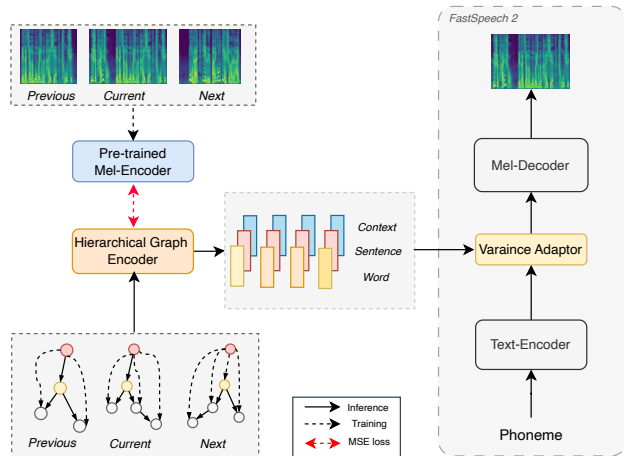

Fig. 1. The overall architecture of HiGNN-TTS

2.1. Graph Construction with Syntactic Information

We construct the graph from a syntax tree, which is built by performing dependency parsing at the word level due to the fact that the smallest unit of prosody is a word in Chinese [29] [2]. Within a syntax tree, bi-lexicalized dependencies are utilized to depict the relationships between words for semantic analysis. Moreover, with syntactic information in the text, we can effectively leverage syntactic relationships to capture prosodic variation. Based on the syntax tree, we in-

---

[2]In modern Chinese, most words are compounds written with two or more characters.

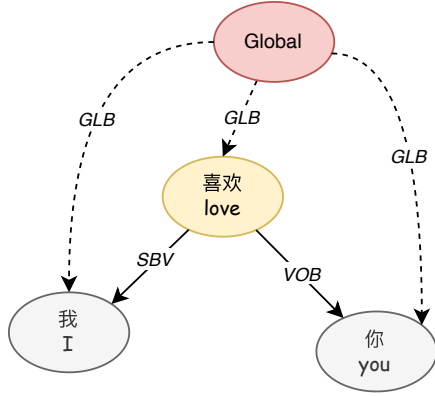troduce a virtual global node to connect all nodes in the tree, forming a syntax graph, as shown in Fig.2.

Fig. 2. The syntax graph example of "I love you." Nodes with different colors represent different constituents: red indicates virtual global nodes, yellow indicates the head of the dependence tree, and gray indicates normal word nodes. GLB, SBV, and VOB represent global relations, subject-verb relations, and verb-object relations, respectively.

The syntax graph can be represented by using a directed graph, defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Each node is denoted as $v \in \mathcal{V}$, representing a word in the corresponding sentence. Each edge is denoted as $e = (v_i, v_j) \in \mathcal{E}$, representing a specific dependency relation from node v_i to node v_j. In this work, we introduce a virtual global node, denoted as v_g, to establish connections to all existing nodes. The connections can be defined as the pair $e_{v_g,i} = (v_g, v_i) \mid \forall v_i \in \mathcal{V}$. The global node improves the model's perception by providing a larger receptive field, enabling it to take into account distant dependencies and effectively capture the prosody patterns present in sentences.

We initialize graph node features using the BERT embedding extracted from a pre-trained Chinese BERT [3]. For each word node, we first extract BERT embeddings of all Chinese characters in the word and then perform average pooling to get the word nodes' BERT embedding. Additionally, the global node is initialized using BERT CLS embedding.

2.2. Hierarchical Prosody Modeling based on Graph

We propose to model hierarchical prosody based on the graph via the hierarchical graph encoder, as shown in Fig.3. The hierarchical graph encoder takes the graphs $\mathcal{G}$ of the previous, current, and next sentences as input to capture the word-level prosody within the current sentence, the sentence-level prosody in the current sentence, and the cross-sentence context from the neighboring sentences.
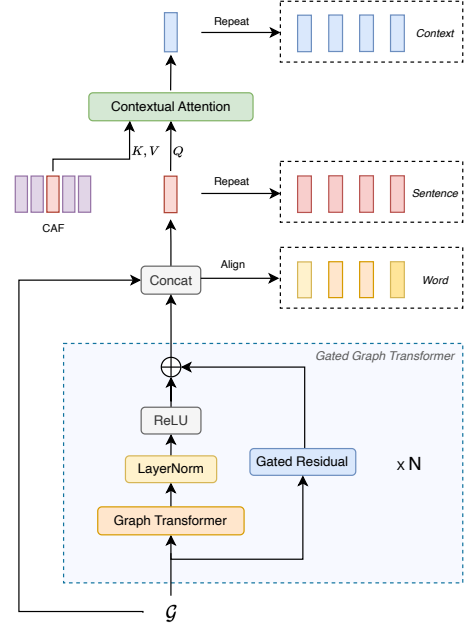
Fig. 3. The architecture of hierarchical graph encoder

Firstly, we adopt Gated Graph Transformer [30] to encode $\mathcal{G}$ into $\mathcal{G}'$, in which each node features are updated with the consideration of correlated node features, enhancing the connections between each word and its surrounding words within a sentence. Specifically, the Gated Graph Transformer contains N layers of Graph Transformer with a gated residual connection, as shown in the bottom part of Fig. 3. Graph Transformer is based on the Transformer model and performs weighted aggregation of node features using a self-attention mechanism. The update of the i-th node feature can be formulated by the following equation:

$$\mathbf{x}_i' = \mathbf{W}_1\mathbf{x}_i + \sum_{j \in \mathcal{N}(i)} \alpha_{i,j}\left(\mathbf{W}_2\mathbf{x}_j + \mathbf{W}_3\mathbf{e}_{ij}\right), \quad (1)$$

where $\mathbf{x}_i'$ represents the updated node feature and $\mathbf{x}_i$ denotes the original node feature. $\mathcal{N}(i)$ denotes the set of neighboring nodes of node i, $\mathbf{e}_{ij}$ corresponds to the edge between node i and its neighboring node j, and $\alpha_{i,j}$ represents the learnable attention weight between node i and its neighboring node j. $\mathbf{W}_1$, $\mathbf{W}_2$, and $\mathbf{W}_3$ are learnable weight matrices applied for linear transformations.

After encoding $\mathcal{G}$ to $\mathcal{G}'$, we fusion $\mathcal{G}$ and $\mathcal{G}'$ via a concatenate operation as the final graph and consumed by the following operations. We believe that the encoded $\mathcal{G}'$ involves acoustic prosody because the hierarchical graph encoder is jointly trained with the acoustic model and constrained by the pre-trained Mel-encoder outputs (as described in Section 2.3). Hence, the node features in the final fusion graph incorporate not only semantic information in the text but also acoustic prosody in the speech. Similar to GST, we apply a tanh acti-

vation to each node in the fusion graph to further enhance the prosody diversity of each node feature.

Additionally, the introduction of the global node in $\mathcal{G}$ allows the Gated Graph Transformer to connect all nodes in the graph, thus better understanding the overall sentence structure and providing sentence-level prosody. Consequently, in this work, we regard the global node feature in the fusion graph of the current sentence as the corresponding sentence-level prosody representation and the other node features as word-level prosody representation.

To capture cross-sentence contextual prosody, we first concatenate the sentence-level representations of previous, current, and next sentences, forming the aggregated features at the long-form level, referred to as the *Context Aggregation Features* (CAF). And then, we utilize a contextual attention mechanism to learn the cross-sentence prosody, where CAF is used as the key and value, and the current sentence-level representation acts as the query. In fact, the contextual attention mechanism can be considered as a certain type of GNN, in which we treat each sentence representation as a node and the entire concatenated representation as a fully connected graph. By using the attention mechanism, we calculate the attention weights between each sentence and other sentences, effectively aggregating contextual information. This process is similar to massage propagation in GNN.

To facilitate the incorporation of the hierarchical prosody representations learned from the hierarchical graph encoder into the acoustic model, we align each word-level representation to the length of the phoneme sequence of each corresponding word, and repeat the sentence-level and cross-sentence prosody representations, forming the same length as the phoneme sequence of the current sentence.

2.3. Hierarchical Supervision from Acoustic Prosody

In order to capture the prosodic variations in speech during hierarchical prosody modeling, we employ a pre-trained mel-encoder model to supervise the prosody representations learned in the hierarchical graph encoder. The mel-encoder is pre-trained on a multi-speaker multi-style dataset to extract sentence-level acoustic prosody embedding from the mel spectrogram of a given reference speech. To obtain a pure prosody embedding, we decouple the speaker-related information in the speech from the extracted prosody representation by using the Domain Adversarial Training (DAT) [31] strategy. Specifically, a speaker classifier with a gradient reversal layer is adopted on the extracted prosody embedding, making the prosody embedding to be speaker indistinguishable.

To supervise hierarchical prosody, speaker-independent acoustic prosody embeddings are extracted from the speech of previous, current, and next sentences, respectively, and then used to directly constrain the corresponding sentence-level prosody representations (i.e., the global node feature) learned in the graph encoder via the mean squared error (MSE) loss. Although the supervision is directly conducted on the sentence-level prosody representation, the cross-sentence prosody representation can also be influenced because it is computed based on the sentence-level prosody representation. Furthermore, with the help of the message-passing capability of GNN, the supervision signals propagate from the sentence-level to word-level nodes, enabling the supervision of acoustic prosody in speech to hierarchical levels. Through hierarchical acoustic prosody supervision, we are able to capture the prosodic variations with high dynamic range comprehensively, thus improving the expressiveness of speech synthesis.

3. EXPERIMENTS

3.1. Training Setup

In the experiments, we use an internal multi-speaker audio-book corpus. The corpus contains roughly 8 hours of recordings collected from two professional Mandarin native speakers. The dataset has a total of 7,016 audio clips, 90% of which of the clips are used for training, and the rest of the clips are used for validation and testing. For feature extraction, 80-dimensional mel-spectrograms are extracted with a 24kHz sampling rate. The frame size is set to 1,024, and the hop size is set to 256. The auxiliary pitch contour is extracted through WORLD vocoder [32]. The phoneme duration is obtained through an HMM-based force alignment model. We use DDParser [33] as our dependency parser [4]. We train all models up to 400k steps on two 3090 GPUs with a batch size of 8. In addition, we use a DSPGAN [34]which is trained on an internal studio-quality dataset with 251 hours from 308 speakers, to convert the generated mel-spectrogram into waveforms.

3.2. Compared Methods

To evaluate the performance of HiGNN-TTS, three baseline models are implemented for comparison:

- **FS2-BERT:** a FastSpeech2-based model, which considers context information by using cross-sentence BERT embedding [17].

- **HCE:** a FastSpeech2-based model, which utilizes Hierarchical Context Encoder (HCE) to predict the sentence-level style embedding from the hierarchical context information [27].

- **ATCE:** a FastSpeech2-based model with Acoustic and Text Context Encoders (ATCE), which uses both text and audio context to obtain context prosody representations [25].

---

[4]https://github.com/baidu/DDParser

Table 1. Objective evaluations for different models.

| Model	MCD ↓	Duration MSE ↓	F0 RMSE ↓
FS2-BERT	4.283	0.1421	51.1192
HCE	4.073	0.1197	41.2083
ATCE	4.123	0.1243	42.4219
HiGNN-TTS	**3.979**	**0.1083**	**36.6086**

3.3. Objective Evaluation

We evaluate the model's performance on the reserved short-sentence testing set with the number of 50 using several widely used objective evaluation metrics as in prior studies[35]. These metrics include the root mean square error (RMSE) for pitch, mean square error (MSE) for the duration, and mel cepstral distortion (MCD). To calculate the RMSE for F0 and MCD, we employ dynamic time warping (DTW) [36] to align the predicted mel-spectrogram with the ground-truth mel-spectrogram.

The objective results are shown in Table 1. HiGNN-TTS achieves superior performance over the three baselines on all objective evaluation metrics. Lower MCD and Duration MSE indicate that HiGNN-TTS obtains better mel and duration prediction and achieves more natural speech synthesis. Notably, we observe a significant degradation in the F0 RMSE of HiGNN-TTS compared to other metrics. This can be attributed to that our hierarchical modeling approach with the supervision of acoustic prosody effectively captures local prosodic variations with high dynamics. Furthermore, we present F0 contours of the same speech synthesized by different models in Figure 4. It reveals that the F0 pattern of our model is close to that of the ground truth recording, indicating that our approach is effective in capturing hierarchical prosodic variations.

3.4. Subjective Evaluation

We conducted Mean Opinion Score (MOS) tests to evaluate the expressiveness and naturalness of the synthetic speech in a similar way to [28, 27]. Moreover, we evaluate 20 short-form sentences (S-MOS) and 10 long-form sentences (L-MOS). S-MOS consists of individual sentences with context, each lasting about 5 seconds. And L-MOS is constructed by combining multiple sentences with a duration of approximately 40 seconds. In each MOS test, a group of 20 native Chinese Mandarin listeners is asked to listen to synthetic speeches and rate them on a scale from 1 to 5 with a 0.5-point interval.

As presented in Table 2, our proposed method, HiGNN-TTS, achieves the best S-MOS of 4.09 and L-MOS of 3.86. By effectively leveraging semantic information, HiGNN-TTS achieves superior performance compared to the other models. Moreover, it is worth noting that, due to potential listener fatigue when listening to long-form speech, the L-MOS score is consistently lower than the S-MOS score. Even so, our hier-
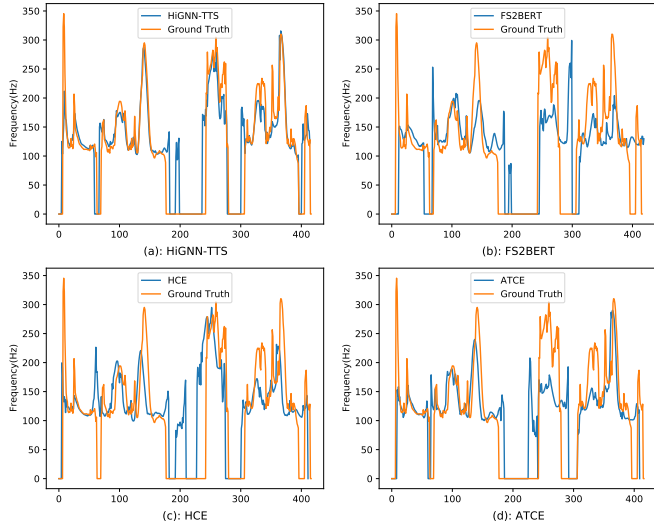
Fig. 4. F0 contours of the same speech synthesized by different models.

Table 2. The S-MOS and L-MOS of different models with 95% confidence intervals.

Model	S-MOS ↑	L-MOS ↑
FS2-BERT	3.52 ± 0.124	3.34 ± 0.113
HCE	3.85 ± 0.116	3.46 ± 0.101
ATCE	3.91 ± 0.102	3.44 ± 0.121
HiGNN-TTS	$\mathbf{4.09 \pm 0.098}$	$\mathbf{3.86 \pm 0.113}$

archical prosody modeling approach still exhibits remarkable superiority over all baseline models in long-form speech. The higher S-MOS and L-MOS scores demonstrate our model's capability to model hierarchical prosody and produce more expressive speech.

3.5. Ablation Study

We perform ablation studies by removing acoustic prosody supervision, contextual attention, and global node (sentence-level representation) in HiGNN-TTS. Then we conduct comparison mean opinion score (CMOS) tests of expressiveness and naturalness on both long-form and short-form speech.

Table 3. Results of ablation studies.

Model	CMOS
HiGNN-TTS	0
- speech supervision	-0.931
- contextual attention	-0.510
- global node	-0.483

The results, listed in Table 3, show that all the designs in HiGNN-TTS yield a significant positive effect on the ex-

pressiveness of generated speech. Specifically, When removing the acoustic prosody supervision strategy from HiGNN-TTS, we find it results in the largest CMOS decrease. This indicates that acoustic prosody supervision helps hierarchical graph prosody encoder to capture more large dynamic range of prosodic variations. Moreover, removing contextual attention also leads decrease in CMOS because of the lack of cross-sentence context information. Based on the removal of contextual attention, we further remove the global node in the syntax graph, and it results in CMOS reductions due to the absence of the overall rhythm in the individual sentence.

4. CONCLUSION

In this work, we propose HiGNN-TTS, a graph-based hierarchical prosody modeling for expressive long-form speech synthesis. We modify the syntax tree by the addition of a virtual global node, which enhances the relationships with the word nodes. Additionally, a hierarchical graph encoder is designed to capture the hierarchical structural relationship in the text at the word level, sentence level, and cross-sentence context level, which broadens the scope of graph-based speech synthesis from intra-sentence to inter-sentence. Furthermore, we introduce the signal of acoustic prosody in the speech to hierarchically supervise the prosody learned from the text in the graph. Ablation studies show the effectiveness of the hierarchical graph encoder and hierarchical supervision from acoustic prosody. Experimental results demonstrate that our proposed HiGNN-TTS significantly improves prosodic variations and generates more natural and expressive short-form and long-form speech.

5. REFERENCES

[1] Yuxuan Wang, R. J. Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc V. Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous, "Tacotron: Towards end-to-end speech synthesis," in *Proc. Interspeech*, 2017, pp. 4006–4010.

[2] Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, "Fastspeech: Fast, robust and controllable text to speech," in *Proc. NeurIPS*, 2019, pp. 3165–3174.

[3] Jaehyeon Kim, Jungil Kong, and Juhee Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *Proc. ICML*, 2021, pp. 5530–5540.

[4] Tao Li, Shan Yang, Liumeng Xue, and Lei Xie, "Controllable emotion transfer for end-to-end speech synthesis," in *Proc. ISCSLP*, 2021, pp. 1–5.

[5] Tao Li, Xinsheng Wang, Qicong Xie, Zhichao Wang, Mingqi Jiang, and Lei Xie, "Cross-speaker emotion

transfer based on prosody compensation for end-to-end speech synthesis," in *Proc. Interspeech*, 2022, pp. 5498–5502.

[6] Tao Li, Xinsheng Wang, Qicong Xie, Zhichao Wang, and Lei Xie, "Cross-speaker emotion disentangling and transfer for end-to-end speech synthesis," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 30, pp. 1448–1460, 2022.

[7] Xinfa Zhu, Yi Lei, Kun Song, Yongmao Zhang, Tao Li, and Lei Xie, "Multi-speaker expressive speech synthesis via multiple factors decoupling," in *Proc. ICASSP*, 2023, pp. 1–5.

[8] Rob Clark, Hanna Silén, Tom Kenter, and Ralph Leith, "Evaluating long-form text-to-speech: Comparing the ratings of sentences and paragraphs," vol. abs/1909.03965, 2019.

[9] R. J. Skerry-Ryan, Eric Battenberg, Ying Xiao, Yuxuan Wang, Daisy Stanton, Joel Shor, Ron J. Weiss, Rob Clark, and Rif A. Saurous, "Towards end-to-end prosody transfer for expressive speech synthesis with tacotron," in *Proce. ICML*, 2018, pp. 4700–4709.

[10] Yuxuan Wang, Daisy Stanton, Yu Zhang, R. J. Skerry-Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Ye Jia, Fei Ren, and Rif A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *Proc. ICML*, 2018, pp. 5167–5176.

[11] Daisy Stanton, Yuxuan 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.

[12] Viacheslav Klimkov, Srikanth Ronanki, Jonas Rohnke, and Thomas Drugman, "Fine-grained robust prosody transfer for single-speaker neural text-to-speech," in *Proc. Interspeech*, 2019, pp. 4440–4444.

[13] Daxin Tan and Tan Lee, "Fine-Grained Style Modeling, Transfer and Prediction in Text-to-Speech Synthesis via Phone-Level Content-Style Disentanglement," in *Proc. Interspeech*, 2021, pp. 4683–4687.

[14] Younggun Lee and Taesu Kim, "Robust and fine-grained prosody control of end-to-end speech synthesis," in *Proc.ICASSP*, 2019, pp. 5911–5915.

[15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL-HLT*, 2019, pp. 4171–4186.

[16] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le, "Xlnet:

Generalized autoregressive pretraining for language understanding," in *Proc. NeurIPS*, 2019, pp. 5754–5764.

[17] Tomoki Hayashi, Shinji Watanabe, Tomoki Toda, Kazuya Takeda, Shubham Toshniwal, and Karen Livescu, "Pre-trained text embeddings for enhanced text-to-speech synthesis," in *Proc. Interspeech*, 2019, pp. 4430–4434.

[18] Bing Yang, Jiaqi Zhong, and Shan Liu, "Pre-trained text representations for improving front-end text processing in mandarin text-to-speech synthesis," in *Proc. Interspeech*, 2019, pp. 4480–4484.

[19] Yixuan Zhou, Changhe Song, Jingbei Li, Zhiyong Wu, Yanyao Bian, Dan Su, and Helen Meng, "Enhancing word-level semantic representation via dependency structure for expressive text-to-speech synthesis," in *Proc. Interspeech*, 2022, pp. 5518–5522.

[20] Zhenhui Ye, Zhou Zhao, Yi Ren, and Fei Wu, "Syntaspeech: Syntax-aware generative adversarial text-to-speech," in *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022*, 2022, pp. 4468–4474.

[21] Guanghui Xu, Wei Song, Zhengchen Zhang, Chao Zhang, Xiaodong He, and Bowen Zhou, "Improving prosody modelling with cross-utterance bert embeddings for end-to-end speech synthesis," in *Proc. ICASSP*, 2021, pp. 6079–6083.

[22] Ya-Jie Zhang, Chao Zhang, Wei Song, Zhengchen Zhang, Youzheng Wu, and Xiaodong 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.

[23] Pilar Oplustil Gallegos and Simon King, "Using previous acoustic context to improve text-to-speech synthesis," vol. abs/2012.03763, 2020.

[24] Shun Lei, Yixuan Zhou, Liyang Chen, Zhiyong Wu, Shiyin Kang, and Helen Meng, "Context-aware coherent speaking style prediction with hierarchical transformers for audiobook speech synthesis," in *Proc. ICASSP 2023*, 2023, pp. 1–5.

[25] Detai Xin, Sharath Adavanne, Federico Ang, Ashish Kulkarni, Shinnosuke Takamichi, and Hiroshi Saruwatari, "Improving speech prosody of audiobook text-to-speech synthesis with acoustic and textual contexts," in *Proc. ICASSP*, 2023, pp. 1–5.

[26] Liumeng Xue, Frank K. Soong, Shaofei Zhang, and Lei Xie, "ParaTTS: Learning linguistic and prosodic cross-sentence information in paragraph-based TTS," *IEEE*

ACM Trans. Audio Speech Lang. Process., pp. 2854–2864, 2022.

[27] Shun Lei, Yixuan Zhou, Liyang Chen, Zhiyong Wu, Shiyin Kang, and Helen Meng, "Towards expressive speaking style modelling with hierarchical context information for mandarin speech synthesis," in *Proc. ICASSP 2022*, 2022, pp. 7922–7926.

[28] Shun Lei, Yixuan Zhou, Liyang Chen, Jiankun Hu, Zhiyong Wu, Shiyin Kang, and Helen Meng, "Towards multi-scale speaking style modelling with hierarchical context information for mandarin speech synthesis," in *Proc. ICASSP*, 2022, pp. 5523–5527.

[29] Min Chu and Yao Qian, "Locating boundaries for prosodic constituents in unrestricted mandarin texts," *Int. J. Comput. Linguistics Chin. Lang. Process.*, vol. 6, 2001.

[30] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjing Wang, and Yu Sun, "Masked label prediction: Unified message passing model for semi-supervised classification," in *Proc. IJCAI*, 2021, pp. 1548–1554.

[31] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor S. Lempitsky, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, pp. 59:1–59:35, 2016.

[32] Masanori Morise, Fumiya Yokomori, and Kenji Ozawa, "World: A vocoder-based high-quality speech synthesis system for real-time applications," *IEICE Transactions on Information and Systems*, pp. 1877–1884, 2016.

[33] Shuai Zhang, Lijie Wang, Ke Sun, and Xinyan Xiao, "A practical chinese dependency parser based on a large-scale dataset," 2020.

[34] Kun Song, Yongmao Zhang, Yi Lei, Jian Cong, Hanzhao Li, Lei Xie, Gang He, and Jinfeng Bai, "DSP-GAN: A gan-based universal vocoder for high-fidelity tts by time-frequency domain supervision from dsp," in *Proc. ICASSP*, 2023.

[35] Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *Proc. ICLR*, 2021.

[36] Toni Giorgino, "Computing and visualizing dynamic time warping alignments in R: The dtw package," *Journal of Statistical Software*, pp. 1–24, 2009.