



Expressive paragraph text-to-speech synthesis with multi-step variational autoencoder

Xuyuan Li^{♠†*}, Zengqiang Shang^{♠*}, Peiyang Shi[♠], Hua Hua^{♠†}, Ta Li^{♠†}, Pengyuan Zhang^{♠†}

♠ Laboratory of Speech and Intelligent Information Processing, Institute of Acoustics, CAS, China
† University of Chinese Academy of Sciences, China

[lixuyuan, shangzengqiang, shipeiyang, huahua, lita, zhangpengyuan]@hcccl.ioa.ac.cn

Abstract

Neural networks have been able to generate high-quality single-sentence speech. However, it remains a challenge concerning audio-book speech synthesis due to the intra-paragraph correlation of semantic and acoustic features as well as variable styles. In this paper, we propose a highly expressive paragraph speech synthesis system with a multi-step variational autoencoder, called EP-MSTTS. EP-MSTTS is the first VITS-based paragraph speech synthesis model and models the variable style of paragraph speech at five levels: frame, phoneme, word, sentence, and paragraph. We also propose a series of improvements to enhance the performance of this hierarchical model. In addition, we directly train EP-MSTTS on speech sliced by paragraph rather than sentence. Experiment results on the single-speaker French audiobook corpus released at Blizzard Challenge 2023 show EP-MSTTS obtains better performance than baseline models.

Index Terms: Paragraph speech synthesis, Expressive speech synthesis, Text-to-speech, Multi-step VAE

1. Introduction

Thanks to the development of neural networks, current text-to-speech (TTS) models [1, 2, 3, 4] have achieved near-human synthesis results on customized datasets recording independent sentences. However, it remains a challenge for synthesizing audiobooks. Different from the customized datasets, audiobooks have two special characteristics: 1) The speech of audiobooks is usually given in the form of paragraphs because the semantic and acoustic features (unrelated to semantics such as acoustic environment, hoarse voice, and nasality) of different sentences in the same paragraph are strongly correlated. 2) There are dramatic shifts in style within a single paragraph, as the reader needs to switch between styles to play the role of narrator and different characters. Recent research [5] points out that TTS methods still have obvious drawbacks when used to synthesize long stories in audiobooks compared with humans.

There are many works have been devoted to capturing semantic dependence. [6, 7, 8] used the embedding of the context sentences extracted by pre-trained language model [9, 10] to predict the global style of the current sentence. In addition to the global style, [11] additionally models local styles based on context semantics. StyleSpeech [12] followed this two-stage style modeling and used unpaired text and audio data to pre-train the context encoder and reference encoder to improve the generalization. However, these works are trained at the sentence level and ignore modeling the acoustic features which are

independent of context semantics. These acoustic features are generally consistent across paragraph-level speech but change throughout the dataset due to the change of record location and the physiological state of speakers. Thus, given different sentences in a paragraph, these sentence-level models could generate speech with very different acoustic features. To generate speech with consistent acoustic features [13, 14, 15] considered previously synthesized speech when synthesizing the current sentence. Another solution is training the model directly on the whole paragraph speech. ParaTTS [16] first performed this paragraph-level training on a modified Tacotron2 [17]. [18] applied it on Fastspeech2 [1] and found paragraph-level training has a huge advantage over single-sentence training in the following two cases: 1) Paragraphs containing many very short sentences. 2) Paragraphs with dramatic shifts in style. Another recent work [19] found that increasing the duration limit of this long-form training can further enhance the performance of paragraph speech synthesis model.

However, these models are based on FastSpeech or Tacotron, which performs not as well as VITS [3]. In addition, they model the style of paragraph speech at no more than two levels of hierarchy, leading to average rhythms. In this paper, we work on adapting the VITS to generate paragraph speech with highly expressive. To achieve this, we expand the single-layer variational autoencoder (VAE) of the original VITS model into five layers, inspired by HierTTS [20]. We describe how it captures the variable style of paragraph speech at five levels: frame, phoneme, word, sentence, and paragraph in Section 2. After that, in Section 3, we propose a novel training strategy to alleviate the posterior collapse problem of deep VAE. Finally, we apply paragraph-level training to our proposed model to generate paragraph speech with smooth transitions in style and acoustic features. The subjective and objective evaluation results on a French audiobook dataset show that our proposed model outperforms all baseline models. The results of ablation studies also demonstrate the effectiveness of our proposed module and training strategy. Audio can be found at <https://lixuyuan102.github.io/EP-MSTTS/>.

2. Model Structure

For convenience, we refer to our proposed model as EP-MSTTS. The overall architecture of EP-MSTTS can be divided into three modules: Multi-step Audio Encoder (MSAE), Multi-step Text Encoder (MSTE), and Multi-step Decoder (MSD), as illustrated in Fig. 1. All modules have five layers including L_1 , L_2 , L_3 , L_4 , and L_5 , which represent the frame, phoneme, word, sentence, and paragraph layer, from fine to coarse.

* Equal contribution.

♠ Corresponding author.

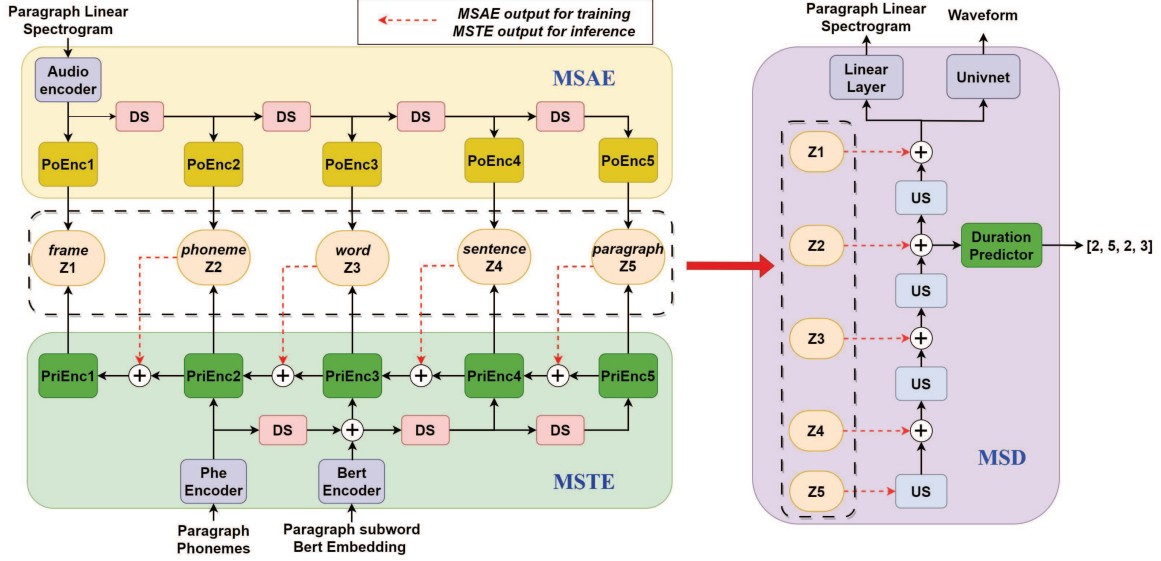


Figure 1: EP-MSTTS architecture including MSAE, MSTE, and MSD. The DS and US denote "Downsampling" and "Upsampling" operations. The PoEnc and PriEnc denote "Posterior Encoder" and "Prior Encoder". " \oplus " represents concatenating two tensors.

2.1. Multi-step Audio Encoder

Within a paragraph of audiobooks, different characters usually have different styles of sentences. The speaker usually expresses these styles by regularly changing the pitch and volume of words. To capture the variable style and express it well, the MSAE breaks down a paragraph speech at five grammatical levels. Specifically, as shown in the upper left (yellow) part of Fig. 1, we downsample the hidden state of paragraph linear-spectrogram step by step according to the pre-given alignment. The acoustic hidden state h_i^{au} at each step is encoded into a posterior latent variable by posterior encoder G_{post} :

$$u_i, \sigma_i = G_{\text{post}}(h_i^{au}) \quad (1)$$

where u_i and σ_i represent the means and variances of the posterior latent variables. Since the information in the initial step is much larger than in the subsequent steps. We only employ the Posterior Encoder mentioned in VITS [3] as the audio encoder of the frame layer. For the other levels, we follow the HierTTS [20], which uses a bidirectional GRU with attention-pooling as the downsample operation to extract coarser-grained information from the hidden state of the upper layer. All the posterior encoders are composed of a convolution layer whose output is activated by Softplus. We emphasize that the MSAE is different from the Hierarchical Audio Encoder (HAE) in HierTTS, which uses text information as an additional input. To avoid posterior latent variables learning useless information directly from the text, our MSAE only learns information from audio.

2.2. Multi-step Text Encoder

Coarse-to-fine iterative generation can produce better results than single-step generation [21]. Based on this principle, we followed the structure of the Hierarchical Audio Decoder from [20] and modified its layer depth to fit the grammatical hierarchy of French paragraphs, as shown in the lower left (green) part of Fig. 1. Starting from the paragraph level, the MSTE encodes semantic information to the prior latent variables. The specific operation of each step is written as follows:

$$s_i, u_i, \sigma_i = G_{\text{pri}}(\text{US}(s_{i-1} + z_{i-1}), h_i^{\text{txt}}) \quad (2)$$

where prior encoder G_{pri} predicts the means u_i and variances σ_i of prior latent variables from semantic information h_i^{txt} , posterior latent variables z_{i-1} from the previous step, and the hidden state of prior predictor s_{i-1} from the previous step. US means Upsampling operation, which will be described in Sec. 2.3. We use a lookup table to obtain the hidden state of paragraph phonemes, two convolution layers with a kernel size of 3, and a Downsampling module to extract the hidden state of words from the Bert subword embedding. The Downsampling operation is implemented in the same way as in MSAE. Because semantic information cannot be split to the frame level, we do not set semantic information input for the frame layer. All prior encoders consist of N feed-forward transformer blocks [22] and a linear layer. Following the order from the frame layer to the paragraph layer, we set N to [4,4,3,3,2] respectively.

2.3. Multi-step Decoder

We hypothesize that generating speech waveforms using only z from the frame layer leads to redundancy of information between layers and accumulates inference errors across layers, since MSTE inference z in a recursive manner. To address this, we upsample the hidden states h of the upper step and concatenate it with z at each step, as illustrated in the right (purple) part of Fig. 1. The Upsampling operation is implemented by the Length Regulator in [22]. The hidden states at each step are given by equation (3).

$$h_i = \text{US}(h_{i-1}) + z_i \quad (3)$$

We employ the Univnet [23] to generate speech waveforms from the hidden states in the last step. Since the z of each layer can contribute directly to the reconstruction of the speech waveform. Encoders tend to encode independent information into different layers. For vocoder, this decoding method also mitigates the mismatch between training and inference caused by multi-layer cumulative errors. In addition, we employ a linear layer as the parallel structure of the vocoder to predict the linear-spectrogram of the paragraph speech. The specific effect of this parallel structure is described in section 3.

3. Training Strategy

We employ KL loss to make the output of MSTE consistent with that of MSAE. The KL loss \mathcal{L}_{kl} of all grammatical layers can be expressed as:

$$\mathcal{L}_{kl} = \alpha_1 \mathcal{L}_1 + \alpha_2 \mathcal{L}_2 + \alpha_3 \mathcal{L}_3 + \alpha_4 \mathcal{L}_4 + \alpha_5 \mathcal{L}_5 \quad (4)$$

We set the $\alpha_1, \alpha_2, \alpha_3, \alpha_4,$ and α_5 as a decreasing sequence: [1, 0.25, 0.07, 0.01, 0.005] because the information from the finer-grained layer contributes more to speech reconstruction. For MSD, we apply the adversarial training on vocoder with Multi-Period Discriminator [2] and Multi-Resolution Discriminator [23]. The adversarial loss is denoted by \mathcal{L}_g . We apply the L1 loss of mel-spectrogram \mathcal{L}_m and multi-resolution STFT loss [24] \mathcal{L}_{s1} on the output of the vocoder. In addition, another multi-resolution STFT loss \mathcal{L}_{s2} is used to constrain the output of the linear layer parallel to the vocoder. Finally, we use the L2 loss \mathcal{L}_d to train the duration predictor.

We observed that when directly training the total model, Univnet tends to reconstruct waveform only from the z at low levels due to its overpowering reconstruction capability. This phenomenon, known as the posterior collapse, underscores the necessity for a less potent vocoder to force the encoder to supply more useful information. However, restructuring the waveform is difficult, and we cannot risk compromising the quality of the speech reconstruction by choosing a simple vocoder structure. Inspired by MSMC-TTS [25], we designed another task: predicting the linear-spectrogram with a simple linear layer parallel to the vocoder. We also use the weight annealing [26] on the weight of \mathcal{L}_{kl} to ensure that our model has enough time to extract information from audio without text constraints. Specifically, our training strategy is divided into the following three stages:

1. First, only the MSAE and MSD are trained with the goal of reconstructing the linear spectrogram features. The weight λ_{kl} of \mathcal{L}_{kl} is set to 0.00001. As a result, the total loss can be given as:

$$\mathcal{L}_{tot} = 2.5\mathcal{L}_{s2} + 5\mathcal{L}_d + \lambda_{kl}\mathcal{L}_{kl} \quad (5)$$

2. Then the λ_{kl} starts to increase slowly, which means the latent variable of each layer begins to adjust the style information it contains with the constraints of the MSTE. We change the λ_{kl} as following equation:

$$\lambda_{kl} = 1e^{-5}(Step_g - Step_w) \quad (6)$$

$Step_g$ is the global training iterations. $Step_w$ is the total number of iterations for the first stage.

3. Finally, the network gives up reconstructing the linear spectrogram features and instead reconstructs the speech waveform. Meanwhile, λ_{kl} keeps increasing until 1. As a result, the total loss is given as:

$$\mathcal{L}_{tot} = \mathcal{L}_g + 1.5\mathcal{L}_{s1} + 2.5\mathcal{L}_m + \lambda_{kl}\mathcal{L}_{kl} + \mathcal{L}_d \quad (7)$$

The first and second stages were set to 10,000 and 30,000 iterations, respectively.

4. Experiments

4.1. Dataset

In this work, we used a portion of the French audiobook dataset [27]. This part consists of 289 chapters of 5 audiobooks from

Librivox read by a female native French speaker. After slicing, we compiled 18,155 paragraphs¹ totaling 51.6 hours. The length distribution of these paragraph speech is shown in Fig 2. All the audios are encoded in the PCM format at a sampling rate of 22.05 kHz with transcribed text and partial aligned phoneme sequences. We used a transformer-based sequence-to-sequence model to generate the remaining phonemes and used the MFA [28] to align these phonemes and audios. In our experiments, 150 paragraphs were selected for testing, 150 paragraphs for validation, and the remaining paragraphs for training. We extracted the 512-dimensional linear-spectrograms with a hop length of 300 points and a window length of 800 points. A pre-trained Bert released on Huggingface² was used to extract the subword embedding for paragraph text.




Figure 2: Probability distribution of paragraph duration length.

4.2. Experiment settings

Four models were implemented as baseline models. 1) VITS-U: To demonstrate that our VITS-based model is more suitable than the original VITS for synthesizing paragraph speech in audiobooks. We cloned the official implementation³ of VITS as one of the baseline models. To eliminate the bias introduced by the vocoder, we replaced the vocoder structure in VITS with Univnet⁴. 2) VITS-U-S: We split the audios in the dataset by sentence and trained the VITS-U on them, to compare the performance of sentence-level training with paragraph-level training on VITS. 3) VITS-U-B: We add a style encoder for VITS to represent models that utilize semantic information to predict the style of paragraph speech at a limited number of hierarchical levels. The style encoder first upsample the Bert subword embedding to the word level, as described in Sec 2.2, then downsample it to the phoneme level, and finally converts it to a style embedding utilizing the same structure as the text encoder in VITS. VITS-U-B concatenates this style embedding and text embedding together to predict the prior latent variables. 4) HierTTS: To demonstrate that our proposed improvements enhance the hierarchical modeling. We cloned the official implementation⁵ of HierTTS and modified it to the same five-

¹We cut 26 paragraphs that exceeded 218 seconds into smaller segments that did not exceed 218 seconds.

²<https://github.com/huggingface/transformers>

³<https://github.com/jaywalnut310/vits>

⁴<https://github.com/rishikksh20/UnivNet-pytorch>

⁵<https://github.com/shang0712/HierTTS>

layer structure as EP-MSTTS. All models except VITS-U-S are trained at the paragraph level using a dynamic batch with a total duration of no more than 218 seconds.




Figure 3: Speech of three continuous sentences generated by VITS-U-S (top) and VITS-U (bottom).

4.3. Evaluation

4.3.1. Training Evaluation

We visualized the speech of three continuous sentences generated by VITS-U-S and VITS-U. The speech of VITS-U-S is generated sentence by sentence and is concatenated together, while the speech of VITS-U is generated in a single inference. As shown in Fig. 3, besides its inability to predict the pauses between sentences, VITS-U-S generates speech with discontinuities in loudness (the second sentence) and high frequency (the third sentence). On the contrary, the acoustic features of the speech synthesized by VITS-U are consistent. In addition, we further computed the mel cepstral distortion (MCD) and Log-F0 root mean square error (log-F0 RMSE) of generated speech with the scripts in ESPnet⁶. As shown in Tab. 1, compared to VITS-U-S, VITS-U decreased by 0.034 on MCD and remained unchanged on log-F0 RMSE. These results illustrate that paragraph-level training leads to better paragraph audiobook synthesis on VITS, which is consistent with the findings of previous experiments on FastSpeech-based and Tacotron-based models.

Table 1: Subjective and objective comparison on test set

Metric	MOS ($\pm 95\%CI$) \uparrow	MCD \downarrow	log-F0 RMSE \downarrow
GT	4.21 \pm 0.07	N/A	N/A
VITS-U-S	-	5.692	0.266
VITS-U	3.35 \pm 0.09	5.658	0.266
VITS-U-B	3.52 \pm 0.07	5.574	0.246
HierTTS	3.80 \pm 0.08	5.458	0.250
EP-MSTTS	3.98\pm0.07	4.876	0.215

4.3.2. Model Evaluation

We perform subjective and objective evaluations on models that are trained at the paragraph level. 25 native French speakers

⁶<https://github.com/espnet/espnet/tree/master/egs2/TEMPLATE/tts1>

were invited to score as the role of the 5-scale mean opinion score (MOS) in terms of sound quality, naturalness, and expressiveness. Each listener was randomly assigned 20 paragraphs from the test set. As Tab. 1 shown, VITS-U-B outperforms VITS-U on both subjective and objective metrics, which demonstrates the importance of contextual semantic information for capturing the style of paragraph speech. In addition, HierTTS, which models the style with five grammatical levels, outperforms the single-level VITS-U-B on MOS and MCD, but its Log-F0 RMSE rises by a relative 0.004. This is reasonable because, as described in its original paper [20], the F0 of speech generated by HierTTS fluctuates over a wide range. We believe this is mainly due to the recursive five-layer decoder accumulating inference errors at each layer. Finally, our proposed model obtains the highest MOS and has significant reductions over HierTTS on MCD/log-F0 with 0.582/0.035. These results illustrate that our proposed modifications significantly improve the performance of the VITS-based hierarchical model on audiobook paragraph speech synthesis.

Table 2: CMOS results on ablation studies.

Models	CMOS
EP-MSTTS	0
+Text in MSAE	-0.12
-MSD	-0.26
-Training Strategy	-0.17

4.3.3. Ablation Study

EP-MSTTS differs from previous Hierarchical VAE models mainly in MSAE, MSD, and the training strategy. We conducted the comparative MOS (CMOS) test (on a 7-point scale) to demonstrate the effectiveness of these proposed modifications. As shown in Tab. 2: 1) After introducing the text input in MASE like HierTTS, the CMOS is reduced by 0.12. We also observe that the KL loss converges faster at the word, sentence, and paragraph levels, which means that less information is compressed into these levels. 2) Removing the MSD and using only the z from the frame layer reduces CMOS by 0.26. This most significant decrease proves our hypothesis in Sec. 2.3. 3) Finally, replacing our training strategy with the staged KL weighted annealing in HierTTS reduces the CMOS by 0.17. It is worth noting that, unlike the staged KL weighted annealing, our training strategy does not require extensive experience in deciding when to increase the KL weights of a particular level.

5. Conclusion

In this paper, we introduce multi-step VAE into the VITS to enhance its performance on audiobook paragraph synthesis. In addition, we propose a series of modifications to improve the multi-step VAE. Experiment results demonstrate that the paragraph speech synthesized by EP-MSTTS outperforms all the baseline models. Ablation studies confirm the effectiveness of our proposed improvements. In the future, we will further analyze the information that has been decoupled into different grammatical layers and explore to realize the smooth transition of synthesized speech between paragraphs, using coarse-grained information of the paragraph layer.

6. Acknowledgements

This work is supported by the National Key Research and Development Program of China(No. 2021YFC3320102) and Post-doctoral Fellowship Program of CPSF(No. GZB20230811)

7. References

- [1] 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*, 2021.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] E. Rodero and I. Lucas, "Synthetic versus human voices in audiobooks: The human emotional intimacy effect," *New Media & Society*, vol. 25, no. 7, pp. 1746–1764, 2023.
- [6] 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 *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7922–7926.
- [7] 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 *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6079–6083.
- [8] J. Li, Y. Meng, C. Li, Z. Wu, H. Meng, C. Weng, and D. Su, "Enhancing speaking styles in conversational text-to-speech synthesis with graph-based multi-modal context modeling," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7917–7921.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [10] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "Xlnet: Generalized autoregressive pretraining for language understanding," *Advances in neural information processing systems*, vol. 32, 2019.
- [11] X. Chen, S. Lei, Z. Wu, D. Xu, W. Zhao, and H. Meng, "Un-supervised multi-scale expressive speaking style modeling with hierarchical context information for audiobook speech synthesis," in *Proceedings of the 29th International Conference on Computational Linguistics*, 2022, pp. 7193–7202.
- [12] X. Chen, X. Wang, S. Zhang, L. He, Z. Wu, X. Wu, and H. Meng, "Stylespeech: Self-supervised style enhancing with vq-vae-based pre-training for expressive audiobook speech synthesis," *arXiv preprint arXiv:2312.12181*, 2023.
- [13] Y. Xiao, S. Zhang, X. Wang, X. Tan, L. He, S. Zhao, F. K. Soong, and T. Lee, "Contextspeech: Expressive and efficient text-to-speech for paragraph reading," *arXiv preprint arXiv:2307.00782*, 2023.
- [14] 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 *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [15] 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 *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [16] L. Xue, F. K. Soong, S. Zhang, and L. Xie, "Paratts: Learning linguistic and prosodic cross-sentence information in paragraph-based tts," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2854–2864, 2022.
- [17] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [18] P. Makarov, S. Ammar Abbas, M. Łajszczak, A. Joly, S. Karlapati, A. Moinet, T. Drugman, and P. Karanasou, "Simple and Effective Multi-sentence TTS with Expressive and Coherent Prosody," in *Proc. Interspeech 2022*, 2022, pp. 3368–3372.
- [19] W. Zhang, C.-C. Yeh, W. Beckman, T. Raitio, R. Rasipuram, L. Golipour, and D. Winarsky, "Audiobook synthesis with long-form neural text-to-speech," in *12th Speech Synthesis Workshop (SSW) 2023*, 2023.
- [20] Z. Shang, P. Shi, P. Zhang, L. Wang, and G. Zhao, "Hiertts: Expressive end-to-end text-to-waveform using a multi-scale hierarchical variational auto-encoder," *Applied Sciences*, vol. 13, no. 2, p. 868, 2023.
- [21] K. Gregor, I. Danihelka, A. Graves, D. Rezende, and D. Wierstra, "Draw: A recurrent neural network for image generation," in *International conference on machine learning*. PMLR, 2015, pp. 1462–1471.
- [22] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fastspeech: Fast, robust and controllable text to speech," *Advances in neural information processing systems*, vol. 32, 2019.
- [23] W. Jang, D. Lim, J. Yoon, B. Kim, and J. Kim, "Univnet: A neural vocoder with multi-resolution spectrogram discriminators for high-fidelity waveform generation," *arXiv preprint arXiv:2106.07889*, 2021.
- [24] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6199–6203.
- [25] H. Guo, F. Xie, X. Wu, F. K. Soong, and H. MengFellow, "Msmc-tts: Multi-stage multi-codebook vq-vae based neural tts," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [26] S. R. Bowman, L. Vilnis, O. Vinyals, A. Dai, R. Jozefowicz, and S. Bengio, "Generating sentences from a continuous space," in *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 10–21. [Online]. Available: <https://aclanthology.org/K16-1002>
- [27] G. Bailly, O. Perrotin, and M. Lenglet, "Resources for End-to-End French Text-to-Speech Blizzard challenge," jan 2023. [Online]. Available: <https://doi.org/10.5281/zenodo.7560290>
- [28] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kald." in *Interspeech*, vol. 2017, 2017, pp. 498–502.