



MaskedSpeech: Context-aware Speech Synthesis with Masking Strategy

Ya-Jie Zhang¹, Wei Song¹, Yanghao Yue, Zhengchen Zhang, Youzheng Wu, Xiaodong He

JD Technology Group

{zhangyajie23, songwei11, yueyanghao, zhangzhengchen1, wuyouzheng1, hexiaodong}@jd.com

Abstract

Many speech synthesis systems only consider the information within each sentence and ignore the contextual semantic and acoustic features. This makes it inadequate to generate high-expressiveness paragraph-level speech. In this paper, a context-aware speech synthesis system named MaskedSpeech is proposed, which considers both contextual semantic and acoustic features. Inspired by the masking strategy in speech editing research, the acoustic features of the current sentence are masked out and concatenated with those of contextual speech, and further used as additional model input. Furthermore, cross-utterance coarse-grained and fine-grained semantic features are employed to improve the prosody generation. The model is trained to reconstruct the masked acoustic features with the augmentation of both the contextual semantic and acoustic features. Experimental results demonstrate that the MaskedSpeech outperformed the baseline systems significantly in terms of naturalness and expressiveness.

Index Terms: TTS, Speech synthesis, prosody, speech editing, context, masked speech

1. Introduction

With the development of deep learning technology, current text-to-speech (TTS) systems [1, 2, 3, 4] could generate high-quality speech. However, most speech synthesis systems only consider the information in the current sentence to be synthesized and ignore the contextual text and speech. With the rising requirements on the naturalness and expressiveness of TTS systems, such as synthesizing audiobooks, these speech synthesis systems lack the ability to generate paragraph-level speech with adequate naturalness and expressiveness, as human speech is context coherent in a paragraph.

Thus, recent works [5, 6, 7, 8, 9, 10] incorporated semantic features of the current sentence or contextual sentences to improve the prosody of synthesized speech, especially for paragraph speech generation. Xiao *et al.* [9] extracted semantic information from the current sentence by a pretrained BERT [11] model and showed that the incorporation of semantic features could improve the prosody of synthesized speech. Then works [5, 6, 7, 10] further employed cross utterances semantic features in the speech synthesis systems, and all of them showed significant rhythm-enhancing effectiveness. Although contextual semantic information could improve the prosodic performance to a certain extent, these works do not take into account the variation of contextual acoustic features, which is also

very important for generating expressive and natural speech for a paragraph.

Some works [6, 12, 13] used acoustics features to improve the prosody modeling, but most of them extracted sentence-level prosody representation and lost the fine-grained information of contextual acoustic features. In addition, for the reference-based models [12, 13], it's non-trivial to select proper reference audio. [14] and [15] introduced an acoustic encoder to learn fine-grained prosody features instead of utterance-level features and claimed that their models could improve prosody generation and relieve the one-to-many mapping issue. However, [14] and [15] only consider speech from the current sentence and the prosody features need to be predicted from text when reference audio is unavailable during inference, this led to another one-to-many mapping issue, as the same text can potentially be used to predict acoustic features from different speech.

To improve prosody generation for a paragraph, we proposed MaskedSpeech in this paper, which takes FastSpeech2 [16] as network backbone and incorporates both contextual semantic and acoustic features. The main contributions of MaskedSpeech contains:

- Inspired by the speech editing works [17, 18, 19], we concatenate the acoustic features of previous and current sentences and mask out those from the current sentence, then those features are used as additional decoder input for the proposed network. The network could learn fine-grained prosody features from the contextual speech at the decoder part, and no prosody prediction from text is required.
- Phonemes of both previous and current sentences are concatenated and sent into the phoneme encoder of the proposed model, which helps the model learn fine-grained semantic features from the previous sentence.
- Furthermore, a cross-utterance (CU) encoder is used to extract coarse-grained sentence-level semantic representation from neighboring sentences, this allows the proposed model to capture semantic correlation between contextual sentences.

The proposed MaskedSpeech is trained by reconstructing the masked acoustic features, by utilizing the contextual text, contextual speech, and sentence-level semantic representation, which makes MaskedSpeech able to improve prosody modeling and alleviate the one-to-many mapping issue naturally.

2. Related work

Unsupervised speech representation learning [20, 21, 22] has demonstrated its superiority in the field of speech recognition and related downstream tasks. Although the success of speech

¹ Equal contributions.

This work was supported by the National Key R&D Program of China under Grant No. 2020AAA0108600.

representation learning in speech understanding tasks, it's not suitable for the speech synthesis task due to its training strategy. MAM [23] and FAT-MLM [24] masked several acoustic segments randomly and used the masked acoustic features as network input, then Transformer-based [25] network is employed to reconstruct the masked regions. But the quality of the reconstructed acoustic features could not meet the requirement of speech synthesis. Bai *et al.* [17] proposed a network that uses masked acoustic features as input and the network is trained to predict high-quality acoustic features, then their proposed network could be further applied to speech synthesis and speech editing tasks. CampNet [18] introduced a speech editing model, the acoustic features are randomly masked out by a prior ratio and then sent to the decoder for masked speech reconstruction. Each decoder layer used an attention mechanism to learn contextual features from unmasked part.

All the mentioned works for speech reconstruction try to predict high-quality and contextual consistent acoustic features by learning the context relationship of masked and unmasked speech. Inspired by these works, we incorporate the masking strategy into cross-utterance speech synthesis task to improve the prosody modeling, with the assumption that both the textual and acoustic context could facilitate the prosody improvement of speech synthesis. Instead of masking on word-level speech region for speech editing tasks, the acoustic features of the current sentence are concatenated with those of previous sentence, and then acoustic features from the current sentence are masked out. In other words, we proposed a novel context-aware speech synthesis method that reconstructs speech for the current sentence with improved prosody by utilizing both the semantic and acoustic features of contextual sentences.

3. Proposed Method

The proposed MaskedSpeech utilizes FastSpeech2 [16] as the network backbone, which further fuses contextual semantic and acoustic features to improve prosody generation. MelGAN [26] is used to convert acoustic features to speech in the experiments.

3.1. The MaskedSpeech Model

Fig. 1 illustrates the network architecture of the proposed MaskedSpeech model. The backbone network consists of a phoneme encoder, a variance adaptor, a decoder, and a Post-Net module. To improve prosody for the speech in a paragraph by contextual features, concatenated sentences are used as the input of the phoneme encoder, a CU text encoder is employed to obtain CU semantic features from neighboring sentences. Furthermore, a masked mel-encoder is designed to extract local prosody features from the contextual speech, and then the decoder could learn global prosody dependency from the contextual speech. By conditioning on contextual semantic and acoustic features, the proposed MaskedSpeech alleviates the one-to-many mapping issue and could generate speech with improved prosody for a paragraph, and the generated speech is more coherent and consistent with contextual speech.

3.1.1. Learning from contextual text

The prosody of the current sentence could be impacted by the relative sentence position, discourse relations (DRs) in neighboring sentences, the emotion of contextual sentences, etc. To improve the prosody generation by contextual text, we incorporate both fine-grained contextual phonemes and coarse-grained sentence-level semantic features into the proposed model.

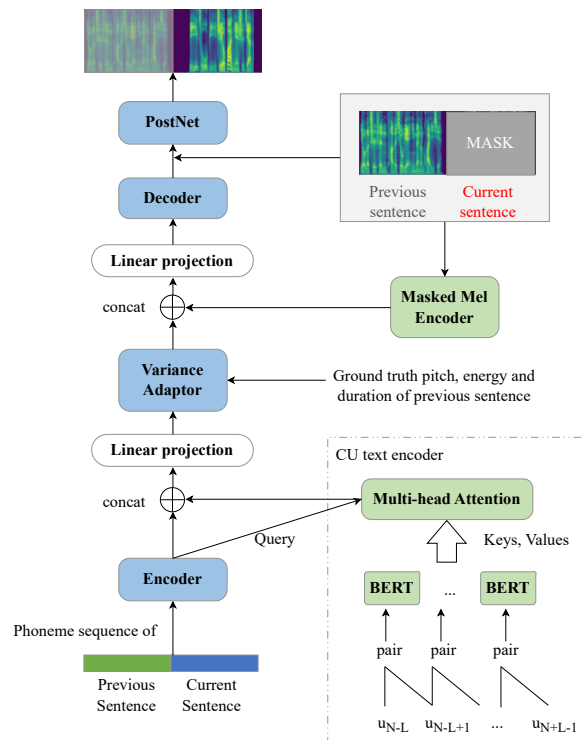


Figure 1: The architecture of the MaskedSpeech model.

The phoneme sequences of the previous sentence and current sentence are concatenated as the input of the phoneme encoder, rather than only using the phonemes of the current sentence in general methods. Thus, the phoneme encoder could learn phoneme representation which takes previous sentence into consideration implicitly, and then the variance adaptor could generate more expressive prosody features.

Further, a CU encoder is used to incorporate sentence-level semantic representation. Following the work [5], paired BERT embedding (PBE) is used to extract sentence-level coarse-grained semantic features between sentence pairs in the CU encoder. As shown in Fig. 1, the neighboring sentences $\{u_{N-L}, \dots, u_N, u_{N+1}, \dots, u_{N+L-1}\}$ can be used to derive $2L$ sentence pairs, where u_N means the current sentence and L is the number of sentences taken into account in the past or following context. Then PBE is extracted by a pre-trained BERT model from the $[CLS]$ token of a sentence pair. A multi-head attention module is used to incorporate PBEs into the proposed model, where PBEs are taken as keys and values, and the output of the phoneme encoder is considered as the query. The output of the CU encoder is concatenated with the output of the phoneme encoder, then linear projected.

3.1.2. Learning from contextual speech

Human speech is naturally coherent in a paragraph, the prosody and emotion of the current sentence could be affected by the contextual speech. Contextual speech usually contains rich, expressive, and detailed prosody information, which could be utilized to improve the prosodic performance of the current sentence in speech synthesis systems, especially for the speech generation in a paragraph.

A masked mel-encoder is incorporated into the FastSpeech2 network in the proposed MaskedSpeech. The acoustic features from contextual sentences, *i.e.* the mel-spectrograms of

the previous and current sentences are concatenated and those from the current sentence are masked with a special *MASK* token, then those features are taken as the input of the masked mel-encoder. The masked mel-encoder consists of two 1-dimensional convolutional (Conv1d) layers, it could learn to extract local prosody features from the acoustic features in the previous sentence. ReLU activation function is used in each Conv1d layer, followed by layer normalization [27] and dropout [28].

The output of the masked mel-encoder is concatenated with the output of the variance adaptor and then projected to the dimension of the decoder. The decoder consists of several self-attention-based Conformer [29] blocks, so the decoder could learn both local and global prosody dependency implicitly from the contextual acoustic features. As shown in Fig. 1, the predicted acoustic features of the previous sentence are replaced with ground truth features at the decoder’s output part, then a Post-Net module with 5 Conv1d layers is used to refine the predicted acoustic features. As the PostNet consists of Conv1d layers and learns local features, it could improve the boundary prosody transition for the current sentence.

3.2. Training and Inference Strategies

3.2.1. Training

During training, the model is trained to reconstruct the masked acoustic features and predict the prosody features of the current sentence. In other words, we only compute training objectives for the current sentence. The losses include phoneme-level duration prediction loss, pitch and energy prediction losses, and mel-spectrogram reconstruction loss. Mean squared error (MSE) is used as the training objectives for pitch, energy, and duration prediction, and mean absolute error (MAE) is used for the mel-spectrogram reconstruction.

3.2.2. Inference

Different from the training procedure, where ground truth contextual text and speech are available and consecutive with the current sentence, there is no ground truth previous speech in the inference procedure, although the contextual text is available for a paragraph speech synthesis. To solve this discrepancy, we tried to select text-audio pair randomly from the training corpus and use it as the contextual speech and phonemes for the proposed model. As this random selection would make a difference compared to the training strategy, an AB preference test is conducted to compare the system performance with ground truth contextual speech and randomly selected contextual speech, the AB test result showed that no significant difference was observed for these two different kinds of contextual speech. We could also select contextual speech with the expected style and prosody deliberately, to generate speech with the desired prosody.

In addition, when a single sentence and corresponding masked speech (*i.e.* a few words’ acoustic features are masked out) are used as the network input, the proposed MaskedSpeech turns into a speech editing model, which can achieve inserting, replacing, and deleting part of audio by editing the corresponding text, as the model is trained to recover the masked acoustic region given the contextual text and speech.

Table 1: *The naturalness MOS results of different models with 95% confidence intervals.*

Model	MOS
Ground truth	4.43±0.07
Fastspeech2 [16]	4.05±0.08
MaskedSpeech	4.21±0.07

4. Experiments

4.1. Experimental Setup

A Chinese Mandarin audiobook corpus with 29 hours of recordings was used in the experiments, this corpus was recorded by a professional male speaker. The text and recordings in this corpus are consecutive. 200 sentences were selected for testing, and 200 sentences for validation, the remaining sentences were used for training. The test set contains several paragraphs of consecutive sentences for evaluation with ground truth contextual text and speech. The audio files were down-sampled to 16kHz sampling rate. 80-dimensional mel-spectrograms were used as the acoustic features in the experiment and were extracted with a frame shift of 12 ms and frame length of 48 ms. The phoneme boundaries and duration were extracted by Montreal forced alignment (MFA) [30] tool. Pitch was extracted the PyWORLD¹ toolkit.

The pre-trained “bert-base-chinese” model [31] released by Huggingface² was used to extract cross utterance sentence embedding in the CU encoder. The BERT model consists of 12 Transformer layers, each layer contains 768 hidden units and 12-heads are used for each attention layer. It is worth mentioning that the parameters of the BERT model are fixed during training. The multi-head attention in the CU encoder has 4 heads and 512 hidden units. In our experiments, PBEs are generated by 2 preceding sentences, 2 following sentences, and the current sentence. A filter size of 256 and kernel size of 3 is used in each Conv1d layer of the masked mel-encoder.

The proposed model is trained by Adam [32] optimizer, with a batch size of 16 and noam [25] learning rate schedule. The learning rate is warmed up to a maximum value of 10^{-3} in the first 4000 steps and then exponentially decayed. Adam optimizer is configured with $\beta_1=0.9$, $\beta_2=0.98$ and $\epsilon=10^{-9}$. The model is trained by a total of 200,000 steps, with L2 parameter regularization of weight $1e-5$.

4.2. Subjective Evaluation

To evaluate the quality of the proposed MaskedSpeech model, the Mean Opinion Score (MOS) evaluation is conducted to evaluate the naturalness and expressiveness of the synthesized speech³. A Fastspeech2 model was trained and used as the baseline model. 20 sentences were randomly selected from the test set and synthesized by different models, each audio is rated by 10 native Chinese listeners. The listeners are asked to rate the speech’s naturalness and expressiveness on a scale from 1 to 5 with intervals of 0.5. Score 1 means the speech is unnatural completely and score 5 means the speech is completely natural. For the proposed model, ground truth context speech and text are used for the audio generation. Tab. 1 shows the MOS

¹<https://github.com/JeremyCCHsu/Python-Wrapper-for-World-Vocoder>

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

³The audio samples can be found at <https://speech11.github.io/MaskedSpeech/>.

Table 2: Subjective preference scores (%) among different models, where N/P denotes “No Preference” and p means the p -value of t -test between two systems.

Fastspeech2	PBE	MaskedSpeech w/o PBE	Random	MaskedSpeech	N/P	p
19.5	-	-	-	49.5	31.0	<0.001
-	23.0	-	-	54.0	23.0	<0.001
-	-	26.5	-	38.0	35.5	0.0425
-	-	-	27.5	29	43.5	0.779

results. Our proposed MaskedSpeech model achieved a higher MOS score than the baseline Fastspeech2 model in terms of naturalness and expressiveness. This indicates that with the help of both contextual semantic and acoustic features, the proposed model could improve the prosody generation.

4.3. Ablation Study on Contextual Features

In this section, ablation study is conducted to evaluate the effectiveness of both contextual semantic and acoustic features. The masking strategy is expected to help the synthesized audio to have a suitable and consecutive speaking style and prosody contour under the condition of the neighboring utterances. In addition to masking strategy of acoustic features, there are some contributions corresponding to contextual semantic information which can further improve prosody modeling. To evaluate the effectiveness of these contextual features, three AB preference tests were conducted for the ablation study. First, we removed both CU encoder and masked mel-encoder from MaskedSpeech, which is a pure **Fastspeech2** model. Then, we removed only masked mel-encoder, which is a **PBE**-based model. In addition, we removed only the CU encoder, which is referred to **MaskedSpeech w/o PBE** model (“w/o” means without). The MaskedSpeech model was compared with Fastspeech2 model, PBE-based model, and MaskedSpeech w/o PBE model separately. Tab. 2 shows the results of AB preference tests. For AB preference tests, the synthesized audio is concatenated with its previous recording to let the rater judge the consecutiveness and expressiveness of the speech. As shown in Tab. 2, MaskedSpeech outperforms the other three models significantly, which proves the effectiveness of using contextual semantic and acoustic features. In addition, removing the masked mel-encoder results in a larger decline than removing the CU text encoder, this proves that fine-grained prosody features learned from the contextual speech play an important role in improving the prosody modeling.

As mentioned in section 3.2.2, when the ground truth previous recording is not available, we randomly sample an audio from the corpus instead and use it as the contextual speech. The method is referred to as **Random**. An AB preference test was conducted to evaluate its performance compared with **MaskedSpeech**. The result is shown in the fourth line of Tab. 2. The MaskedSpeech model with ground truth previous speech is slightly better than **Random** but the difference is not significant. In general, in some cases like narrations and normal-style dialogues, using random audio with a general style is comparable to using ground truth previous speech with the same style. However, we found that in some emotional dialogues, sometimes the speaker deliberately uses more distinctive prosody to align with the context speech. In this case, using ground truth context speech will be better than randomly selected one without the desired style. In conclusion, random sampling provides a substitutable way when there is no ground truth contextual speech. Furthermore, compared with random sampling, manually selecting contextual speech will derive more controllable

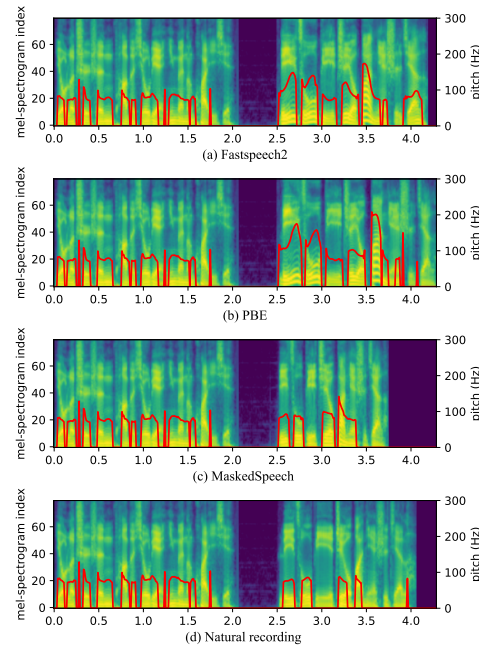


Figure 2: A comparison between the mel-spectrograms of three models and natural recordings.

and expressive synthesized speech.

4.4. A Case Study

Fig. 2 shows the mel-spectrograms of the speech synthesized by three different models and natural recording. Each sub-figure shows the concatenated mel-spectrograms of previous and current sentences with distinct silence as the boundary. The previous sentence is natural recording and the second sentence is the synthetic audio. Fig. 2 (d) shows that the speaking speed of the nature recording is relatively faster and the fundamental frequency is lower. MaskedSpeech perceives this property according to the contextual semantic and acoustic features and generates speech with a relatively faster speaking rate and lower pitch. This makes the prosodic transition of the consecutive sentences sound continuous and meaningful.

5. Conclusion

A context-aware speech synthesis system is proposed in this paper, with improved prosody generation by utilizing the contextual semantic and acoustic features from neighboring sentences. The proposed model takes the concatenated and masked mel-spectrograms as augmented input and learns to reconstruct the masked mel-spectrograms. The phoneme encoder takes into account the phonemes from context sentences and those of the current sentence, and a CU encoder is used to extract cross-utterance semantic representation. Experimental results showed the effectiveness of the proposed MaskedSpeech model.

6. References

- [1] 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 *Proc. ICASSP*, 2018, pp. 4779–4783.
- [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.
- [3] C. Yu, H. Lu, N. Hu, M. Yu, C. Weng, K. Xu, P. Liu, D. Tuo, S. Kang, G. Lei, D. Su, and D. Yu, “DurLAN: Duration Informed Attention Network For Multimodal Synthesis,” in *Proc. Interspeech*, 2020.
- [4] W. Song, X. Yuan, Z. Zhang, C. Zhang, Y. Wu, X. He, and B. Zhou, “DIAN: Duration Informed Auto-Regressive Network for Voice Cloning,” in *Proc. ICASSP*. IEEE, 2021, pp. 8598–8602.
- [5] 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*. IEEE, 2021, pp. 6079–6083.
- [6] J. Cong, S. Yang, N. Hu, G. Li, L. Xie, and D. Su, “Controllable Context-Aware Conversational Speech Synthesis,” in *Proc. Interspeech*, 2021, pp. 4658–4662.
- [7] Y.-J. Zhang and Z.-H. Ling, “Learning Deep and Wide Contextual Representations Using BERT for Statistical Parametric Speech Synthesis,” in *Proc. ICSDSP*, 2021, pp. 146–150.
- [8] P. Oplustil-Gallegos, J. O’Mahony, and S. King, “Comparing acoustic and textual representations of previous linguistic context for improving Text-to-Speech,” in *Proc. SSW 11*, 2021, pp. 205–210.
- [9] Y. Xiao, L. He, H. Ming, and F. K. Soong, “Improving Prosody with Linguistic and Bert Derived Features in Multi-Speaker Based Mandarin Chinese Neural TTS,” in *Proc. ICASSP*. IEEE, 2020, pp. 6704–6708.
- [10] 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.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proc. NAACL-HLT*, Jun. 2019, pp. 4171–4186.
- [12] 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. 5167–5176.
- [13] Y.-J. Zhang, S. Pan, L. He, and Z.-H. Ling, “Learning Latent Representations for Style Control and Transfer in End-to-end Speech Synthesis,” in *Proc. ICASSP*. IEEE, 2019, pp. 6945–6949.
- [14] 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.
- [15] M. Chen, X. Tan, B. Li, Y. Liu, T. Qin, S. Zhao, and T.-Y. Liu, “AdaSpeech: Adaptive Text to Speech for Custom Voice,” *arXiv preprint arXiv:2103.00993*, 2021.
- [16] 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,” 2021.
- [17] H. Bai, R. Zheng, J. Chen, M. Ma, X. Li, and L. Huang, “A³T: Alignment-Aware Acoustic and Text Pretraining for Speech Synthesis and Editing,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 1399–1411.
- [18] T. Wang, J. Yi, L. Deng, R. Fu, J. Tao, and Z. Wen, “Context-Aware Mask Prediction Network for End-to-End Text-Based Speech Editing,” in *Proc. ICASSP*. IEEE, 2022, pp. 6082–6086.
- [19] D. Tan, L. Deng, Y. T. Yeung, X. Jiang, X. Chen, and T. Lee, “EditSpeech: A Text Based Speech Editing System Using Partial Inference and Bidirectional Fusion,” in *Proc. ASRU*. IEEE, 2021, pp. 626–633.
- [20] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations,” *Proc. NeurIPS*, vol. 33, 2020.
- [21] A. Conneau, A. Baevski, R. Collobert, A. Mohamed, and M. Auli, “Unsupervised Cross-Lingual Representation Learning for Speech Recognition,” in *Proc. Interspeech*, 2021, pp. 2426–2430.
- [22] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, 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.
- [23] J. Chen, M. Ma, R. Zheng, and L. Huang, “MAM: Masked Acoustic Modeling for End-to-End Speech-to-Text Translation,” *arXiv preprint arXiv:2010.11445*, 2020.
- [24] R. Zheng, J. Chen, M. Ma, and L. Huang, “Fused Acoustic and Text Encoding for Multimodal Bilingual Pretraining and Speech Translation,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 12 736–12 746.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention Is All You Need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [26] K. Kumar, R. Kumar, T. de Boissiere, L. Gestein, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, “MelGAN: Generative Adversarial Networks for Conditional Waveform Synthesis,” *Advances in neural information processing systems*, vol. 32, 2019.
- [27] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer Normalization,” *arXiv preprint arXiv:1607.06450*, 2016.
- [28] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A Simple Way to Prevent Neural Networks from Overfitting,” *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [29] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, “Conformer: Convolution-augmented transformer for speech recognition,” *Proc. Interspeech 2020*, pp. 5036–5040, 2020.
- [30] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal Forced Aligner: Trainable Text-Speech Alignment Using Kaldi,” in *Interspeech*, 2017, pp. 498–502.
- [31] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, “Transformers: State-of-the-Art Natural Language Processing,” in *Proc. EMNLP*, 2020, pp. 38–45.
- [32] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” in *Proc. ICLR*, San Diego, 2015.