



# VQTTS: High-Fidelity Text-to-Speech Synthesis with Self-Supervised VQ Acoustic Feature

Chenpeng Du, Yiwei Guo, Xie Chen, Kai Yu\*

MoE Key Lab of Artificial Intelligence, AI Institute  
X-LANCE Lab, Department of Computer Science and Engineering  
Shanghai Jiao Tong University, Shanghai, China

{duchenpeng, cantabilekwok, chenxie95, kai.yu}@sjtu.edu.cn

## Abstract

The mainstream neural text-to-speech(TTS) pipeline is a cascade system, including an acoustic model(AM) that predicts acoustic feature from the input transcript and a vocoder that generates waveform according to the given acoustic feature. However, the acoustic feature in current TTS systems is typically mel-spectrogram, which is highly correlated along both time and frequency axes in a complicated way, leading to a great difficulty for the AM to predict. Although high-fidelity audio can be generated by recent neural vocoders from ground-truth(GT) mel-spectrogram, the gap between the GT and the predicted mel-spectrogram from AM degrades the performance of the entire TTS system. In this work, we propose VQTTS, consisting of an AM txt2vec and a vocoder vec2wav, which uses self-supervised vector-quantized(VQ) acoustic feature rather than mel-spectrogram. We redesign both the AM and the vocoder accordingly. In particular, txt2vec basically becomes a classification model instead of a traditional regression model while vec2wav uses an additional feature encoder before HifiGAN generator for smoothing the discontinuous quantized feature. Our experiments show that vec2wav achieves better reconstruction performance than HifiGAN when using self-supervised VQ acoustic feature. Moreover, our entire TTS system VQTTS achieves state-of-the-art performance in terms of naturalness among all current publicly available TTS systems.

**Index Terms:** speech synthesis, vq-wav2vec, txt2vec, vec2wav

## 1. Introduction

Text-to-speech (TTS) synthesis is a process that transforms a transcript into its corresponding speech. Compared with traditional statistical parametric speech synthesis [1], neural TTS model [2, 3] based on deep neural network shows a better performance. The mainstream neural text-to-speech(TTS) pipeline is a cascade system, including an acoustic model(AM) that predicts acoustic feature from the input transcript and a vocoder that generates waveform according to the given acoustic feature. Two well-known AMs are Tacotron 2 [4] based on encoder-attention-decoder architecture and FastSpeech 2 [5] based on Transformer blocks. As for the vocoders, generative adversarial network (GAN) [6] based vocoders, such as multi-band MelGAN [7] and HifiGAN [8], are widely used for their high quality of speech and fast generation speed. Another important type of vocoders is neural source-filter model [9, 10] based on the mechanism of human voice production.

However, the acoustic feature in current models is typically mel-spectrogram, which is highly correlated along both time and frequency axes in a complicated way, leading to a great

difficulty for the AM to predict. Although high-fidelity audio can be generated by the neural vocoders from ground-truth(GT) mel-spectrogram, the gap between the GT mel-spectrogram and the predicted one from AM degrades the performance of the entire TTS system.

Generally, there are two approaches to alleviate this problem. The first one is to provide stronger conditions for the AM in addition to the input transcript, such as prosody and linguistic feature. In this way, the AM could be more certain in acoustic feature prediction and often provides a better speech quality. The prosodies in different granularities have been provided to TTS in the literature. For example, [11, 12, 13] uses phoneme-level prosodies, [14] uses word-level prosodies and [15, 16] uses hierarchical prosodies. The effectiveness of VQ-VAE is also explored in the related literature [17, 18]. In addition, [19, 20] introduces syntactic graph and [21, 22] introduces word embeddings to TTS models. Another approach to address the problem is to exploit better training criterion. The most common training criterion for the AM is the L1 or L2 loss, which assumes the distribution of the acoustic feature is unimodal. However, the real distribution is much more complicated. Hence, some research uses normalizing flow [23] in AM, such as FlowTTS [24] and GlowTTS [25]. The normalizing flow transforms the data distribution into a known simple distribution and is optimized via maximum log likelihood. However, the flow models should be carefully designed to ensure invertibility, which greatly restricts the capabilities of such models.

In this work, we propose VQTTS, consisting of an AM txt2vec and a vocoder vec2wav, which uses self-supervised vector-quantized(VQ) acoustic feature rather than mel-spectrogram. In particular, txt2vec basically becomes a classification model rather than a traditional regression model while vec2wav uses an additional feature encoder before HifiGAN generator for smoothing the discontinuous quantized feature. Instead of predicting the complicated mel-spectrogram which is highly correlated along both time and frequency axes, txt2vec only needs to consider the correlation along time axis in feature prediction, which narrows the gap between GT and predicted acoustic feature dramatically. Our experiments show that vec2wav achieves better reconstruction performance than HifiGAN when using self-supervised VQ acoustic feature. Moreover, our entire TTS system VQTTS achieves state-of-the-art performance in terms of naturalness among all current publicly available TTS systems.

## 2. Self-Supervised VQ Acoustic Feature

Recently, the acoustic features extracted by deep neural networks have been found to be superior to traditional acoustic

\* Corresponding author.

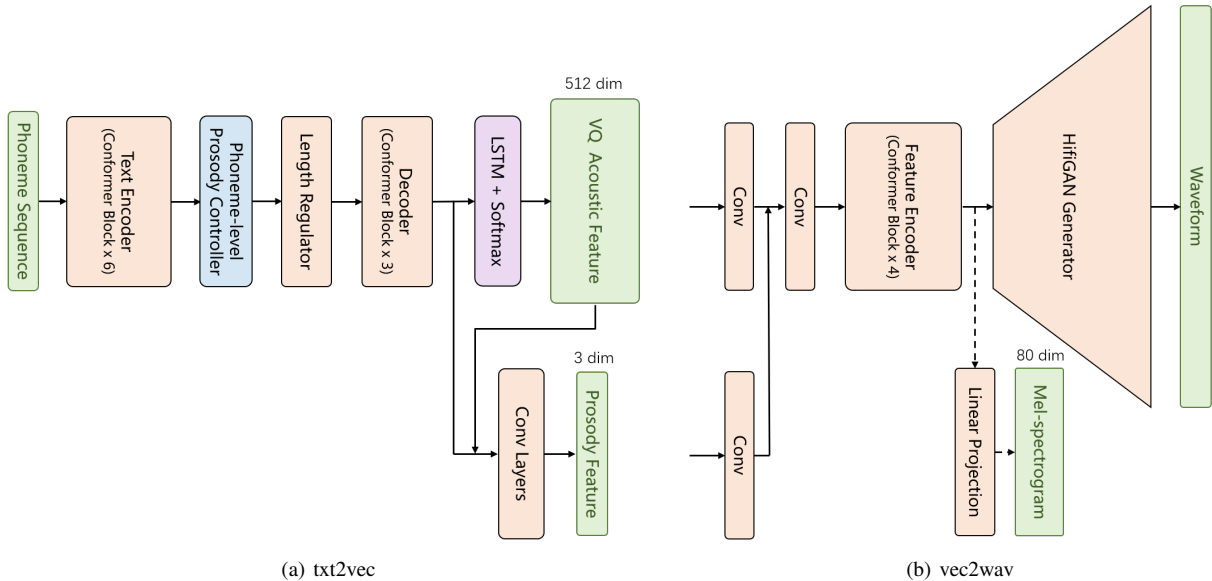


Figure 1: Model architecture of VQTTS, consisting of *txt2vec* and *vec2wav*. The two parts are connected with VQ acoustic feature together with prosody feature.

features in automatic speech recognition (ASR). These types of neural networks are typically trained with only speech data in a self-supervised manner. They take the raw speech waveform  $\mathcal{X}$  as input and generate the features  $\mathcal{Z}$  that represents the characteristics of the speech segments. For example, wav2vec [26] trains a multi-layer convolutional network optimized via a contrastive loss. Specifically, it tries to extract the features where we can predict several successive frames from the current and several previous frames.

Later, vector quantization is applied to self-supervised feature extraction. Vq-wav2vec [27] quantizes the acoustic feature  $\mathcal{Z}$  to VQ acoustic feature  $\hat{\mathcal{Z}}$  with gumbel-softmax or k-means. Then  $\hat{\mathcal{Z}}$  is used to train a BERT [28] model for ASR initialization. In avoid of mode collapse where only a little amount of vectors in the codebook are actually used, vq-wav2vec divides the dimension of  $\hat{\mathcal{Z}}$  into 2 groups and quantizes them separately. Based on that, wav2vec 2.0 [29] jointly trains the feature extractor and the BERT via a contrastive loss and a diversity loss that encourages the model to utilize more vectors in the codebook. HuBERT [30] introduces another quantization strategy. Instead of training the codebook jointly with the feature extractor, HuBERT clusters the features with k-means in advance.

In addition to the ASR task, self-supervised VQ acoustic feature has been also proven to be effective in a range of other tasks, such as voice conversion [31, 32], speech translation [33] and speech separation [34]. In this paper, we further investigate the use of VQ acoustic feature in the TTS task.

### 3. VQTTS

VQTTS uses self-supervised VQ acoustic feature for speech synthesis. It has been found in the literature [32, 31] that waveform reconstruction from VQ acoustic feature needs additional prosody feature. Therefore, in this work, we use 3 dimensional prosody feature, including log pitch, energy and probability of voice(POV)[35]. The prosody features are then normalized to zero means and unit variances. For simplicity, we abbreviate the combination of the VQ acoustic feature and the 3 di-

mensional prosody feature to VQ&pros in the following sections. VQTTS contains two parts, the acoustic model *txt2vec* which predicts VQ&pros from input phoneme sequence and the vocoder *vec2wav* which generates waveform from VQ&pros. Here we elaborate the two parts in the following two subsections respectively.

#### 3.1. *txt2vec*

##### 3.1.1. Model architecture

Before training *txt2vec*, we label the phoneme-level(PL) prosodies for all phonemes in advance, which will be detailed in the next subsection. Here we first demonstrate the overall model architecture of *txt2vec* in Figure 1(a). The text encoder consists of 6 Conformer blocks [36], which encodes the input phonemes into hidden states  $\mathbf{h}$ . The hidden states are then sent to a PL prosody controller which predicts the PL prosody labels and a duration predictor that predicts the duration for each phoneme. After that, we repeat the hidden states according to the corresponding phoneme durations as in FastSpeech 2. The decoder is composed of 3 Conformer blocks and its output is passed through an LSTM layer followed by a Softmax activation function for VQ acoustic feature classification. Then the decoder output and the VQ acoustic feature are further concatenated and passed into 4 convolution layers, each followed by layer normalization and a dropout layer, for prosody feature prediction. The phoneme duration and prosody feature are trained with L2 and L1 loss respectively while the PL prosody label and the VQ acoustic feature are trained with cross entropy loss. The overall training criterion is

$$\mathcal{L}_{\text{txt2vec}} = \mathcal{L}_{\text{PL,lab}} + \mathcal{L}_{\text{dur}} + \mathcal{L}_{\text{VQ}} + \mathcal{L}_{\text{pros}}. \quad (1)$$

##### 3.1.2. Phoneme-level prosody labelling

As mentioned earlier, we have 3 dimensional normalized prosody features  $\mathbf{p}$  and then we calculate their dynamic features  $\Delta\mathbf{p}$  and  $\Delta^2\mathbf{p}$ . The total 9 dimensional prosody features  $[\mathbf{p}, \Delta\mathbf{p}, \Delta^2\mathbf{p}]$  are averaged over the frames within each

phoneme, so that we can represent the prosody of each phoneme with one vector. Then we cluster all the PL prosody representations into  $n$  classes with  $k$ -means and take the cluster index as the PL prosody label.

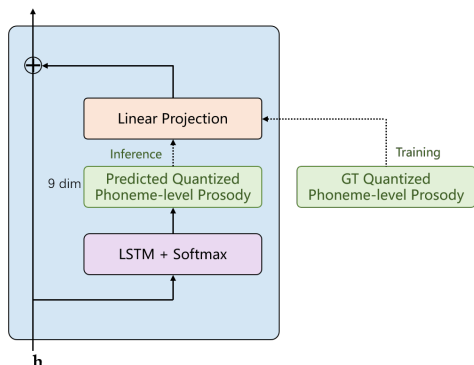


Figure 2: The detail of phoneme-level prosody controller. The PL prosodies are quantized with the  $k$ -means.

The architecture of the PL prosody controller is illustrated in Figure 2, which is trained to predict the PL prosody labels from the text encoder output  $\mathbf{h}$  with an LSTM. Then the quantized PL prosodies, i.e. the centers of the corresponding  $k$ -means clusters, are then projected and added to  $\mathbf{h}$  for controlling the following acoustic feature generation. Note that we use GT quantized PL prosodies in training and predicted ones in inference.

### 3.1.3. Beam search decoding

There are two LSTMs in txt2vec, which are used for the autoregressive predictions of PL prosody label and VQ acoustic feature respectively. During training, both the LSTMs are conditioned on their inputs and the ground-truth previous outputs. During inference, we apply beam search decoding. In particular, the decoding starts with an all-zero vector  $\langle \text{sos} \rangle$ . We denote the beam size as  $k$  here. At each decoding step, we consider the top  $k$  classes for all current hypotheses and take the results with the top  $k$  probabilities as the new  $k$  hypotheses. Compared with greedy search decoding that always selects the best result at each step based on the history, beam search decoding considers both the history and the future.

## 3.2. vec2wav

### 3.2.1. Model architecture

The model architecture of vec2wav is demonstrated in Figure 1(b). Both the VQ acoustic feature and the prosody feature are transformed with a convolution layer whose channels are 92 and 32 respectively and the kernel sizes are 5. The two outputs are then concatenated and passed to a convolution layer, a feature encoder, and a HifiGAN generator successively. The feature encoder here is designed for smoothing the discontinuous quantized acoustic feature. It contains 4 Conformer blocks, each of which uses 2 attention heads and 384 dimensional self-attention. The output of the HifiGAN generator is the corresponding waveform. The training criterion for HifiGAN is used for optimizing the vec2wav model.

### 3.2.2. Multi-task warmup

We find that vec2wav is hard to converge when we train the model from scratch with only HifiGAN loss. Therefore, we

propose a multi-task warmup trick, which additionally uses a linear projection layer to predict the mel-spectrogram from the feature encoder output. Formally, we can write the training criterion during the warmup as

$$\mathcal{L}_{\text{vec2wav}} = \mathcal{L}_{\text{HifiGAN}} + \alpha \mathcal{L}_{\text{mel}}. \quad (2)$$

After the warmup, we remove the mel-spectrogram prediction task, which means the  $\alpha$  is set to 0.

## 4. Experiments and Results

### 4.1. Experimental setup

We use LJSpeech [37] dataset in our experiments, which is an English dataset containing about 24 hours speech recorded by a female speaker. We leave out 100 utterances for validation and 150 utterances for testing. All the speech data in this work is resampled to 16kHz for simplicity. We use a publicly available pretrained  $k$ -means-based vq-wav2vec model<sup>1</sup> for VQ acoustic feature extraction. The frame shift of vq-wav2vec is 10ms and the number of possible VQ acoustic vectors is 21.5k. The 3-dimensional prosody feature is extracted by Kaldi [35]. Audio samples are available online<sup>2</sup>.

### 4.2. Speech reconstruction with vocoders

We train both the vec2wav and HifiGAN on the training set with VQ&pros using Adam optimizer for 1 million steps. In vec2wav training, we set  $\alpha$  to 60 at the first 200k training steps for warmup. A HifiGAN with mel-spectrogram is also trained for comparison. Then we evaluate the performance of speech reconstruction on the test set given GT acoustic feature in both subjective and objective ways. In particular, we perform a mean opinion score(MOS) listening test where 15 listeners are asked to rate each utterance from 1 to 5 in terms of speech quality. Each listener is presented with 20 utterances randomly selected from the test set. For objective evaluations, we compute PESQ [38] which measures the similarity between synthetic speech and the corresponding recordings. We also analyze the Gross Pitch Error (GPE) [39] which calculates the proportion of frames whose pitch discrepancy in the recording and synthetic speech is less than 20% among voiced frames. The results are listed in Table 1.

Table 1: Speech reconstruction performance of vocoders on the test set.

Feature	Method	MOS	PESQ	GPE(%)
-	Recording	4.86±0.04	-	-
Mel	HifiGAN	4.68±0.04	3.60	0.79
VQ&pros	HifiGAN	4.53±0.06	2.38	0.98
VQ&pros	vec2wav	<b>4.79±0.06</b>	<b>2.54</b>	<b>0.76</b>

In the objective evaluations, it can be found that vec2wav can better reconstruct the recordings than HifiGAN with VQ&pros. Also, we can see that the PESQ value of vec2wav is worse than that of the HifiGAN with mel-spectrogram. This is largely due to the information loss brought by quantization. However, a closer reconstruction does not indicate a better speech quality. Actually, the differences between the

<sup>1</sup><https://github.com/pytorch/fairseq/tree/main/examples/wav2vec>.

<sup>2</sup><https://cpdu.github.io/vqtts>.

generated speech from vec2wav and the HifiGAN with mel-spectrogram are almost imperceptible. In the subjective listening test, vec2wav performs better than the HifiGAN with VQ&pros and achieves comparable quality to the HifiGAN with mel-spectrogram. As for the HifiGAN with VQ&pros, we can sometimes hear some undesired artifacts, which could be caused by the discontinuous quantized input feature.

### 4.3. Naturalness of text-to-speech synthesis

We train the entire text-to-speech system VQTTS in which txt2vec is optimized with Adam optimizer for 1000 epochs. The number of PL prosody clusters  $n$  is set to 128 and the beam sizes in beam search decoding are set to 5 and 10 in PL prosody and VQ acoustic feature prediction respectively. Then we compare VQTTS with other current popular methods, including Tacotron 2, GlowTTS, FastSpeech 2 and the fully end-to-end TTS model VITS [40]. In the first three baseline systems, we use 80 dimensional mel-spectrogram as the acoustic feature and HifiGAN as the vocoder. Again, MOS listening test is conducted in the same way as the previous section for evaluating the naturalness of the synthetic speech. The results are shown in Table 2 with 95% confidence interval.

Table 2: Evaluation for the text-to-speech synthesis systems.

Setup	Feature	MOS
Recording	-	4.86±0.04
GT Mel + HifiGAN	Mel	4.68±0.04
GT VQ&pros + vec2wav	VQ&pros	4.79±0.06
Tacotron 2 + HifiGAN	Mel	3.67±0.05
GlowTTS + HifiGAN	Mel	3.72±0.05
FastSpeech 2 + HifiGAN	Mel	3.79±0.05
VITS	-	4.62±0.04
VQTTS (txt2vec + vec2wav)	VQ&pros	<b>4.71±0.05</b>

As is expected, quality degradation can be observed in all the cascade baseline TTS systems compared with speech reconstruction from GT mel-spectrogram. Although the fully end-to-end model VITS has a similar quality to VQTTS, it sometimes suffers from unnatural prosodies. The proposed VQTTS, however, generates high-fidelity and natural speech and shows little quality degradation compared with speech reconstruction from GT VQ&pros. Moreover, VQTTS is a cascade TTS system, which is more flexible than the fully end-to-end TTS system VITS.

### 4.4. Prosody Diversity in PL prosody hypotheses

Text-to-speech is a one-to-many mapping, since diverse prosodies are contained in speech in addition to the transcripts. VQTTS models the diversities with the PL prosody controller, which enables us to control the speech synthesis with different PL prosody hypotheses in beam search. Here we synthesize a sentence in the test set with 3 different prosody hypotheses and demonstrate their pitch tracks in Figure 3 where obviously we can see their differences.

### 4.5. Decoding algorithm

We explore the effectiveness of beam search decoding for both PL prosody label and VQ acoustic feature predictions. To this end, we exploit greedy search and beam search with a beam size of 5 and 10 in the two tasks separately. Here, VQ acoustic feature prediction is conditioned on the GT durations and PL




Figure 3: Pitch tracks of the synthetic speech with different prosodies.

prosody labels in order to make sure that the predicted feature is precisely aligned with the GT feature so that we can calculate the prediction accuracy. The results are presented in Table 3 and 4.

We can find that the accuracy in all setups is not so high because we have demonstrated the diversity of speech in the previous section. Despite that, the accuracy of beam search decoding is still slightly better than greedy search in both inference tasks. Also, the beam size of 5 is better in PL prosody label prediction while the size of 10 is better in VQ acoustic feature prediction.

Table 3: Prediction accuracy for phone-level prosody label.

Decoding Algorithm	Accuracy(%)
greedy search	12.12
beam search (beam size = 5)	<b>12.60</b>
beam search (beam size = 10)	12.46

Table 4: Prediction accuracy for VQ acoustic feature.

Decoding Algorithm	Accuracy(%)
greedy search	13.96
beam search (beam size = 5)	14.09
beam search (beam size = 10)	<b>14.20</b>

## 5. Conclusions

In this paper, we propose VQTTS that utilizes self-supervised VQ acoustic feature rather than traditional mel-spectrogram, which dramatically narrows the quality gap between GT and predicted acoustic feature and consequently improves the performance of entire TTS system. The vocoder in VQTTS, vec2wav, uses an additional feature encoder for smoothing the discontinuous quantized input feature and achieves a better reconstruction performance than HifiGAN. We also find that diverse prosodies can be generated by the different PL prosody hypotheses in beam search decoding. Furthermore, beam search decoding performs better than greedy search in both PL prosody and VQ acoustic feature prediction.

## 6. Acknowledgements

State Key Laboratory of Media Convergence Production Technology and Systems Project (No.SKLMCPTS2020003), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), National Natural Science Foundation of China (Grant No.92048205)

## 7. References

- [1] H. Zen, K. Tokuda, and A. W. Black, "Statistical parametric speech synthesis," *speech communication*, vol. 51, no. 11, pp. 1039–1064, 2009.
- [2] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *arXiv preprint arXiv:1609.03499*, 2016.
- [3] Y. Wang, R. J. Skerry-Ryan, D. Stanton *et al.*, "Tacotron: Towards end-to-end speech synthesis," in *Proc. ISCA Interspeech*, 2017, pp. 4006–4010.
- [4] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Ryan, R. A. Saurous, Y. Agiomyriannakis, and Y. Wu, "Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions," in *Proc. IEEE ICASSP*, 2018, pp. 4779–4783.
- [5] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *Proc. ICLR*, 2021.
- [6] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Proc. NIPS*, 2014.
- [7] G. Yang, S. Yang, K. Liu, P. Fang, W. Chen, and L. Xie, "Multi-band MelGAN: Faster waveform generation for high-quality text-to-speech," in *Proc. IEEE SLT*, 2021, pp. 492–498.
- [8] J. Kong, J. Kim, and J. Bae, "Hifi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," *Proc. NeurIPS*, vol. 33, pp. 17 022–17 033, 2020.
- [9] X. Wang, S. Takaki, and J. Yamagishi, "Neural source-filter waveform models for statistical parametric speech synthesis," *IEEE/ACM Trans. ASLP*, vol. 28, pp. 402–415, 2019.
- [10] Z. Liu, K. Chen, and K. Yu, "Neural homomorphic vocoder," in *Proc. ISCA Interspeech*, 2020, pp. 240–244.
- [11] C. Du and K. Yu, "Rich prosody diversity modelling with phone-level mixture density network," in *Proc. ISCA Interspeech*, 2021, pp. 3136–3140.
- [12] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, and Y. Wu, "Fully-hierarchical fine-grained prosody modeling for interpretable speech synthesis," in *Proc. IEEE ICASSP*, 2020, pp. 6264–6268.
- [13] C. Du and K. Yu, "Phone-level prosody modelling with gmm-based mdn for diverse and controllable speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 190–201, 2022.
- [14] Y. Guo, C. Du, and K. Yu, "Unsupervised word-level prosody tagging for controllable speech synthesis," *arXiv preprint arXiv:2202.07200*, 2022.
- [15] Y. Hono, K. Tsuboi, K. Sawada, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Hierarchical multi-grained generative model for expressive speech synthesis," in *Proc. ISCA Interspeech*, 2020, pp. 3441–3445.
- [16] C. M. Chien and H. Y. Lee, "Hierarchical prosody modeling for non-autoregressive speech synthesis," in *Proc. IEEE SLT*, 2021, pp. 446–453.
- [17] X. Wang, S. Takaki, J. Yamagishi, S. King, and K. Tokuda, "A vector quantized variational autoencoder (vq-vae) autoregressive neural  $f_0$  model for statistical parametric speech synthesis," *IEEE/ACM Trans. ASLP*, vol. 28, pp. 157–170, 2019.
- [18] G. Sun, Y. Zhang, R. J. Weiss, Y. Cao, H. Zen, A. Rosenberg, B. Ramabhadran, and Y. Wu, "Generating diverse and natural text-to-speech samples using a quantized fine-grained VAE and autoregressive prosody prior," in *Proc. IEEE ICASSP*, 2020, pp. 6699–6703.
- [19] A. Sun, J. Wang, N. Cheng, H. Peng, Z. Zeng, and J. Xiao, "GraphTTS: graph-to-sequence modelling in neural text-to-speech," in *Proc. IEEE ICASSP*, 2020, pp. 6719–6723.
- [20] R. Liu, B. Sisman, and H. Li, "Graphspeech: Syntax-aware graph attention network for neural speech synthesis," in *Proc. IEEE ICASSP*, 2021, pp. 6059–6063.
- [21] F. Shen, C. Du, and K. Yu, "Acoustic word embeddings for end-to-end speech synthesis," *Applied Sciences*, vol. 11, no. 19, 2021.
- [22] Y. Jia, H. Zen, J. Shen, Y. Zhang, and Y. Wu, "PnG BERT: Augmented BERT on Phonemes and Graphemes for Neural TTS," in *Proc. ISCA Interspeech*, 2021, pp. 151–155.
- [23] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, "Normalizing flows for probabilistic modeling and inference," *Journal of Machine Learning Research*, vol. 22, no. 57, pp. 1–64, 2021.
- [24] C. Miao, S. Liang, M. Chen, J. Ma, S. Wang, and J. Xiao, "Flow-tts: A non-autoregressive network for text to speech based on flow," in *Proc. IEEE ICASSP*, 2020, pp. 7209–7213.
- [25] J. Kim, S. Kim, J. Kong, and S. Yoon, "Glow-tts: A generative flow for text-to-speech via monotonic alignment search," *Proc. NeurIPS*, vol. 33, pp. 8067–8077, 2020.
- [26] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," in *Proc. ISCA Interspeech*, 2019, pp. 3465–3469.
- [27] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-supervised learning of discrete speech representations," *arXiv preprint arXiv:1910.05453*, 2019.
- [28] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proc. NAACL*, 2019, pp. 4171–4186.
- [29] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. NeurIPS*, 2020.
- [30] W. Hsu, B. Bolte, Y. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Trans. ASLP*, vol. 29, pp. 3451–3460, 2021.
- [31] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, "Neural analysis and synthesis: Reconstructing speech from self-supervised representations," *Proc. NeurIPS*, 2021.
- [32] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W. Hsu, A. Mohamed, and E. Dupoux, "Speech resynthesis from discrete disentangled self-supervised representations," in *Proc. ISCA Interspeech*, 2021, pp. 3615–3619.
- [33] A. Lee, P.-J. Chen, C. Wang, J. Gu, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino *et al.*, "Direct speech-to-speech translation with discrete units," *arXiv preprint arXiv:2107.05604*, 2021.
- [34] J. Shi, X. Chang, T. Hayashi, Y.-J. Lu, S. Watanabe, and B. Xu, "Discretization and re-synthesis: an alternative method to solve the cocktail party problem," *arXiv preprint arXiv:2112.09382*, 2021.
- [35] P. Ghahremani, B. BabaAli, D. Povey, K. Riedhammer, J. Trmal, and S. Khudanpur, "A pitch extraction algorithm tuned for automatic speech recognition," in *Proc. IEEE ICASSP*, 2014, pp. 2494–2498.
- [36] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. ISCA Interspeech*, 2020, pp. 5036–5040.
- [37] K. Ito, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [38] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *Proc. IEEE ICASSP*, vol. 2, 2001, pp. 749–752 vol.2.
- [39] L. Rabiner, M. Cheng, A. Rosenberg, and C. McGonegal, "A comparative performance study of several pitch detection algorithms," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 24, no. 5, pp. 399–418, 1976.
- [40] 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.