

# A Multi-Stage Multi-Codebook VQ-VAE Approach to High-Performance Neural TTS

Haohan Guo\*, Fenglong Xie<sup>‡</sup>, Frank K. Soong<sup>†</sup>, Xixin Wu\*, Helen Meng\*

\*The Chinese University of Hong Kong, Hong Kong SAR, China

<sup>‡</sup> Xiaohongshu Inc., Shanghai, China

<sup>†</sup>Microsoft Research Asia, Beijing, China

{hguo, wuxx, hmmeng}@se.cuhk.edu.hk, fenglongxie@xiaohongshu.com, frankkps@microsoft.com

## Abstract

We propose a Multi-Stage, Multi-Codebook (MSMC) approach to high performance neural TTS synthesis. A vector-quantized, variational autoencoder (VQ-VAE) based feature analyzer is used to encode Mel spectrograms of speech training data by down-sampling progressively in multiple stages into MSMC Representations (MSMCRs) with different time resolutions, and quantizing them with multiple VQ codebooks, respectively. Multi-stage predictors are trained to map the input text sequence to MSMCRs progressively by minimizing a combined loss of the reconstruction Mean Square Error (MSE) and “triplet loss”. In synthesis, the neural vocoder converts the predicted MSMCRs into final speech waveforms. The proposed approach is trained and tested with an English TTS database of 16 hours by a female speaker. The proposed TTS achieves a MOS score of 4.41, which outperforms the baseline with an MOS of 3.62. Compact versions of the proposed TTS with much less parameters can still preserve high MOS scores. Ablation studies show that both multiple stages and multiple codebooks are effective for achieving high TTS performance.

**Index Terms:** Multi-stage Multi-codebook (MSMC), Speech Representation, VQ-VAE, Neural TTS

## 1. Introduction

The performance of text-to-speech synthesis is affected by the deviations between ground-truth acoustic features and the predicted features [1, 2]. This problem is usually alleviated by enhancing the acoustic model with more advanced structures [3, 4, 5, 6] and more complicated algorithms [7, 8, 9], or adapting the vocoder to the predicted features [10, 11]. However, these approaches either offer limited improvement [12], or demand higher modeling complexity to TTS. To search for a more effective solution, we attempt to address this problem from the perspective of the feature representation.

For TTS, a good speech representation offers high completeness, i.e. captures sufficient information for high-quality reconstruction of the target audio, which generally calls for higher dimensionality or larger space volume. However, this also increases the difficulty in modeling [13]. The feature with an overly large space may lead to over-fitting as inputs [14] or under-fitting as outputs [15]. Hence, it is also crucial to maintain compactness of the representation with the minimal feature space. The vector-quantized variational autoencoders (VQ-VAE) offers a discrete speech representation with notable compactness in speech codecs [16, 17]. This inspires us to explore the use of a VQ-VAE based feature analyzer for TTS.

We first examine the vector-quantized speech representation, in terms of the trade-off for the lower completeness due to higher compactness. We then proceed to use a multi-stage multi-codebook vector-quantized representation (MSMCR), where multiple codebooks are used to quantize a vector to achieve a richer representation, and to quantize the speech sequence into multiple sub-sequences for preservation of information at different time resolutions. MSMC-TTS is proposed to generate speech from the text via MSMCR. After obtaining MSMCRs from Mel spectrograms via the feature analyzer, a neural vocoder is used to up-sample them to the audio. Meanwhile, the multi-stage predictor, as the acoustic model, maps the input text to MSMCRs progressively in resolution. It is trained with the combination of MSE and triplet loss to better obtain the expected codes. This study compares the proposed approach against baselines in various scenarios to show its effectiveness to TTS and the lower requirement to modeling complexity. Ablation studies validate that multiple stages and multiple codebooks are both crucial for MSMC-TTS.

## 2. Multi-Stage Multi-Codebook Representation

In this section, we will describe the VQ-VAE based feature analyzer and the proposed multi-stage multi-codebook VQ-VAE.

### 2.1. Vector-Quantized Variational AutoEncoder

VQ-VAE aims to learn a discrete latent representation from target data using an encoder-decoder model. The encoder embeds the input sequence  $x$  to the hidden sequence  $h$ , and then quantizes it by replacing each vector  $h_i$  with the nearest prototype vector in the codebook  $C^{M \times N}$  (composed of  $M$  vectors with the dimension of  $N$ ) according to the Euclidean distance. The decoder maps the quantized sequence  $q$  to  $\hat{x}$  for reconstruction. Since quantization is a non-differentiable operation, we usually update the codebook with the exponential moving average [18], and train the model with the loss function as follows:

$$L_{VQ} = MSE(x, \hat{x}) + \alpha * MSE(h, sg(q)) \quad (1)$$

where  $MSE$  and  $sg$  denotes the mean square error and stop-gradient operation, and  $\alpha$  is a weight coefficient.

In this framework, an acoustic feature sequence can be represented by an index sequence. When a vector composed of  $N$  float-point numbers (32 bits) is compressed to one integer with the range of  $[0, M)$ , the compression ratio  $CR$  can be calculated as follows:

$$CR = \frac{N}{1} * \frac{32}{\log_2(M)} \quad (2)$$

---

Work performed during the first author’s internship at Xiaohongshu

In this way, the 80-dim Mel spectrogram can be compressed  $2560/\log_2(M)$  times. The high compression ratio brings higher compactness. However, higher compactness reduces the completeness due to information loss, and the preserved information may be insufficient for speech reconstruction and prediction. To address this problem, we propose the multi-stage multi-codebook representation.

## 2.2. Multi-Stage Multi-Codebook VQ-VAE

The MSMC-VQ-VAE based feature analyzer aims to encode a speech sequence stage-wise into the multi-stage multi-codebook representation (MSMCR), i.e. a set of sub-sequences at different time resolutions, and quantized by multiple codebooks respectively. This approach is composed of two parts – multi-head vector quantization, and multi-stage autoencoder.

### 2.2.1. Multi-Head Vector Quantization

Reducing the compression ratio may give a richer representation, but may be difficult in VQ-VAE. To half the compression ratio, we need to square the codebook size, which leads to exponentially increasing data requirements, and also computation for training to sufficiently update each code. To alleviate this problem, we propose to quantize each vector with multiple codebooks, and hence reduce the compression ratio further:

$$CR = \frac{N}{1} * \frac{32}{\sum_{k=1}^H \log_2(M_k)} \quad (3)$$

where  $H$  denotes the number of codebooks,  $M_k$  denotes the number of codes in the  $k$ -th codebook. To obtain this multi-codebook representation, we propose multi-head vector quantization (MHVQ), as inspired by multi-head attention mechanism [19]. In this approach, one VQ codebook  $C^{M \times N}$  is splitted into  $H$  codebooks  $\{C_1^{M \times \frac{N}{H}}, \dots, C_H^{M \times \frac{N}{H}}\}$ . The hidden vector  $h_i$  is segmented into  $H$  sub-vectors with the dimension of  $\frac{N}{H}$ , and quantized by their corresponding codebooks, finally concatenated as the output  $q_i$ . In this way, without introducing more parameters, the compression ratio can be halved by doubling the number of codebooks.

### 2.2.2. Multi-Stage AutoEncoder

In addition to improving completeness, it is also crucial to well-represent all speech attributes with different time resolutions, such as low-resoluntional prosody and high-resoluntional pronunciation. Otherwise, any overly compressed attribute may lead to serious degradation in overall quality, as the ‘‘Cask Effect’’ reveals. In this work, we propose the multi-stage autoencoder to stably model the speech sequence at different resolutions [20].

As shown in Fig.1, this model aims to encode the input sequence  $x_{1:L}$  into the multi-stage vector-quantized representation  $\{q_{1:L_1}^{(1)}, \dots, q_{1:L_S}^{(S)}\}$ , where  $S$  and  $L_S$  denote the number of stages and the sequence length at  $S$ -th stage. It first embeds  $x_{1:L}$  to hidden sequences with encoders progressively. Before each encoder block, a strided convolutional layer with the down-sampling rate  $d_j$  is applied for down-sampling. Hidden sequences are quantized stably from the highest stage with the lowest resolution. Firstly, the output sequence from  $E_S$  is transformed by a projection layer and quantized to  $q_{1:L_S}^{(S)}$  in the quantizer  $Q_S$ . It is then processed by the residual module  $R_S$ , and up-sampled by repeating each frame for  $d_S$  times. The residual output is concatenated with the next encoder output to calculate the next quantized sequence  $q_{1:L_{S-1}}^{(S-1)}$ , and added with it using a

projection layer to get the next residual output. The remaining sequences are generated in the same way recursively until we obtain  $\hat{x}_{1:L}$ . Besides, the correlation between two adjacent sequences is enhanced through predicting  $q_{1:L_j}^{(j)}$  from the higher-stage sequence  $q_{1:L_{j+1}}^{(j+1)}$ . Hence, the prediction module  $P_j$ , i.e. a projection layer, is employed to generate the predicted sequence  $p_{1:L_j}^{(j)}$  from the residual output. It is helpful to train deeper modules effectively [21]. Finally, the complete loss function of this model is written as follows:

$$L_{MS} = MSE(x, \hat{x}) + \frac{\alpha}{S} \sum_{j=1}^S MSE(h_{1:L_j}^{(j)}, sg(q_{1:L_j}^{(j)})) + \frac{\beta}{S-1} \sum_{j=1}^{S-1} LC(p_{1:L_j}^{(j)}, q_{1:L_j}^{(j)}) \quad (4)$$

where  $h_{1:L_j}^{(j)}$  denotes  $j$ -th latent sequence before quantization,  $LC$  refers to an arbitrary loss function between two sequences,  $\alpha$  and  $\beta$  are weight coefficients.

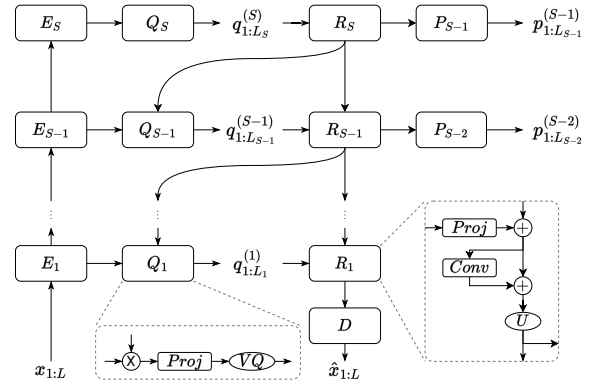


Figure 1: The architecture of the MSMC-VQ-VAE.  $E_S$ ,  $Q_S$ ,  $R_S$ ,  $P_S$  denote the  $S$ -th encoder, quantizer, residual module, prediction module.  $D$ ,  $Proj$ ,  $Conv$  denote the decoder, the projection and convolutional layer.  $U$ ,  $VQ$ , ‘‘ $\times$ ’’, ‘‘ $+$ ’’ are up-sampling, vector quantization, concatenating, and adding operations.

Finally, MSMCR  $\{q_{1:L_j}^{(j,k)} | 1 \leq j \leq S, 1 \leq k \leq H\}$  is extracted by multi-stagely modeling speech with multi-head vector quantization, where  $q_{1:L_j}^{(j,k)}$  denotes the  $j$ -th quantized sequence at  $k$ -th head with the length of  $L_j$ .

## 3. MSMC-TTS

In this section, we will introduce the MSMC-TTS system, including analysis-synthesis, and multi-stage prediction.

### 3.1. Analysis-Synthesis

In analysis, the waveform is first processed to become the acoustic features  $x_{1:L}$  via signal analysis, i.e. Mel spectrograms, which has shown respectable completeness in neural vocoders. Then, the feature analyzer is used to process them into MSMCRs. In synthesis, the neural vocoder convert MSMCRs into the corresponding speech waveform. Sequences at different resolutions in MSMCR are up-sampled to the same resolution as  $x_{1:L}$  by repetition, and then concatenated together to form the input.

### 3.2. Multi-Stage Prediction

We propose a multi-stage predictor based on FastSpeech [5] as the acoustic model of this TTS system.

#### 3.2.1. Model Architecture

As shown in Fig.2, the textual features are encoded and up-sampled to  $c_{1:L}$  by the duration upsampler, i.e. the length regulator in FastSpeech.  $c_{1:L}$  is down-sampled by strided convolutional layers to different resolutions for decoding. First,  $c_{1:L_S}^{(S)}$  is fed to the decoder  $D_S$  to obtain the first predicted sequence  $p_{1:L_S}^{(S)}$ . The last hidden output in the decoder is up-sampled by repetition, and concatenated with  $q_{1:L_{S-1}}^{(S)}$  and  $c_{1:L_{S-1}}^{(S-1)}$  to calculate the next sequence  $p_{1:L_{S-1}}^{(S-1)}$ . The remaining sequences are generated recursively in the same way. In training,  $q_{1:L_j}^{(j)}$  is extracted from the target audio using the analyzer. During inference, we need to replace it with  $\hat{q}_{1:L_j}^{(j)}$  quantized from  $p_{1:L_j}^{(j)}$  using the corresponding codebooks.

#### 3.2.2. Loss Function

To map the input text to the expected prototype vectors in codebooks, the model is trained to estimate them in the continuous space directly. For the predicted vector  $p_i^{(j,k)}$ , i.e. the  $k$ -th head of  $i$ -th vector in the  $j$ -th sequence, we first adopt MSE loss to make it approach the target  $q_i^{(j,k)}$  by reducing the L2 distance:

$$L_{MSE} = \frac{1}{S} \sum_{j=1}^S \frac{1}{H} \sum_{k=1}^H \frac{1}{L_j} \sum_{i=1}^{L_j} MSE(p_i^{(j,k)}, q_i^{(j,k)}) \quad (5)$$

We also adopt the triplet loss [22], an effective ranking loss used in metric learning, to ensure that the L2 distance between the input vector and target code is the smallest among all codes in the codebook  $C_{j,k}$ . The triplet loss is written as:

$$L_{TPL} = \frac{1}{S} \sum_{j=1}^S \frac{1}{H} \sum_{k=1}^H \frac{1}{L_j} \sum_{i=1}^{L_j} TPL(p_i^{(j,k)}, q_i^{(j,k)})$$

$$TPL(p_i^{(j,k)}, q_i^{(j,k)}) = \frac{1}{M} \sum_{\substack{e \neq q_i^{(j,k)} \\ e \in C_{j,k}}} \max(0, \left\| p_i^{(j,k)} - q_i^{(j,k)} \right\|_2 - \left\| p_i^{(j,k)} - e \right\|_2 + \epsilon) \quad (6)$$

where  $\epsilon$  is a constant value called margin number and  $e$  refers to a non-target code in the codebook  $C_{j,k}$ . In this way, the output vector will not only be closer to the target, but also lie farther away from non-target codes. The complete loss function is written as:

$$L_P = L_{MSE} + \gamma * L_{TPL} \quad (7)$$

where  $\gamma$  is a weight coefficient.

## 4. Experiments

This section describes the experiments that evaluate the performance of the proposed TTS system.

### 4.1. Experimental Setup

The standard English single-speaker TTS dataset, Nancy [23], is used in our experiments. It includes 16 hours of clean

recordings of a female voice with corresponding transcripts. The phoneme sequence with punctuations extracted from the text is used as input features. Audios with the sampling rate of 16kHz are converted to 80-dim log-scale Mel spectrograms with a frameshift of 12.5ms and window length of 50ms. Phoneme-level durations are extracted from a pre-trained Tacotron2 model with stepwise attention [24].

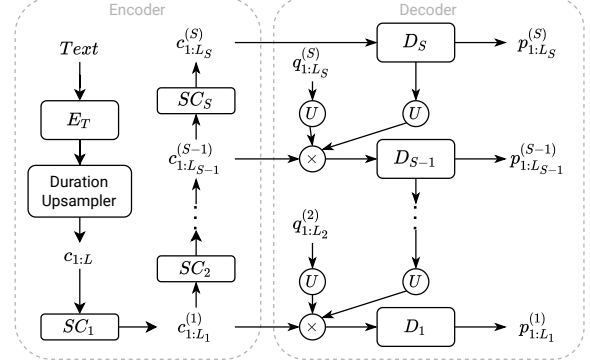


Figure 2: The framework of multi-stage predictor.  $E_T$  refers to the text encoder.  $D_S$ ,  $SC_S$  denote the  $S$ -th decoder and strided convolutional layer respectively.

Neural vocoders in all TTS systems are implemented based on HiFiGAN [25]. UnivNet discriminator [26] is adopted to train the vocoder by both frequency-domain and time-domain discrimination for better audio quality. The MSMC feature analyzer is implemented with FeedForward Transformer blocks [19]. In each encoder or decoder, the input sequence is processed by a projection layer, and added with position encodings, and then fed to 4 Transformer blocks with the model dimension of 256. The loss function adopts  $L_P$  as  $L_C$ , and sets  $\alpha, \beta, \gamma$  to 1, 0.1, 1. All analyzers are trained for 200,000 iterations using the Adam optimizer [27] with the batch size of 64. The learning rate is exponentially decayed from  $2 \times 10^{-4}$  to  $10^{-6}$  after 20,000 iterations. The acoustic model, multi-stage predictor, is implemented based on FastSpeech, where the dimension of hidden vectors is 600, and the number of Transformer blocks in the text encoder and each decoder is 6. It is trained for 100,000 iterations with the same training configuration as the feature analyzer. Besides, except for the predictor loss  $L_P$ , the duration loss in FastSpeech [5] is also employed to train the duration predictor with a weight coefficient of 0.1.

Finally, mean opinion score (MOS) tests are conducted for subjective evaluation. In each test, 15 listeners are involved to evaluate 20 sets of samples. The listener needs to rate each audio sample according to its overall quality with the score range from 1 to 5 and an increment of 0.5<sup>1</sup>.

### 4.2. System Comparison

To evaluate the performance of MSMC-TTS, we compare it with the baseline system, a Mel spectrogram based FastSpeech (Mel-FS) with the same model configuration. For MSMC-TTS, we adopt a two-stage representation with the down-sampling rates [ $d_1 = 1, d_2 = 4$ ], and 4-head vector quantization with the codebook size of 512.<sup>2</sup>

<sup>1</sup>Samples are available at <https://hhguo.github.io/DemoMSMCTTS>

<sup>2</sup>The implementation of FastSpeech is available at <https://github.com/NVIDIA/NeMo>

#### 4.2.1. Standard TTS

We first evaluate their performance in standard single-speaker TTS. Table 1 shows the test results. The baseline system Mel-FS, obtains the MOS of 3.62, which has an obvious gap compared with the original recordings with the MOS of 4.50. Mel-FS-FT refers to Mel-FS with the vocoder fine-tuned for 100,000 more iterations on the synthetic Mel spectrograms, which are generated from the training set using ground-truth durations. This method is often used to adapt the vocoder to the predicted features. In our experiments, it only improves audio quality slightly to the MOS of 3.69. Compared with the two previous systems, MSMC-TTS shows significantly improved performance with MOS of 4.41, which is very closed to the quality of the original recordings. The generated audio presents higher fidelity and naturalness in different time resolutions, which shows the effectiveness of MSMC-TTS.

Table 1: *MOS Test: Standard single-speaker TTS ( $\pm 95\%$  CI)*

Recording	Mel-FS	Mel-FS-FT	MSMC-TTS
4.50 $\pm$ 0.08	3.62 $\pm$ 0.09	3.69 $\pm$ 0.09	4.41 $\pm$ 0.11

#### 4.2.2. Lightweight TTS

We evaluate the impact of the decreased modeling complexity on the proposed approach in lightweight TTS by using fewer parameters or simpler model structures. We conduct another MOS test to compare Mel-FS and MSMC-TTS systems with these three model configurations as the encoder or decoder of the acoustic model:

- M1: 4 layers of 600-dim Transformer blocks.
- M2: 3 layers of 128-dim Transformer blocks.
- M3: 4 layers of 128-dim 1-D convolutional layers.

The test results are shown in Table 2. For Mel-FS, the output quality is degraded seriously from 4.08 to 2.10 when the parameters are reduced from 72.50 MB to 1.75 MB. When Transformer blocks are replaced with the less complex CNNs, the MOS decreases further to 1.86. Both intelligibility and fidelity are affected seriously by the reduced modeling complexity. However, MSMC-TTS still performs well in all situations with only slight degradation as the parameters are reduced. MSMC-TTS trained with M3 (fewest parameters and simplest structure) still significantly outperforms Mel-FS with M1. It validates that the proposed approach has a much lower requirement to the modeling complexity, and great potential in lightweight TTS.

Table 2: *MOS Test: Lightweight TTS ( $\pm 95\%$  CI)*

		M1	M2	M3
Mel-FS	Param (MB)	72.50	1.75	0.52
	MOS	4.08 $\pm$ 0.10	2.10 $\pm$ 0.09	1.86 $\pm$ 0.09
MSMC-TTS	Param (MB)	116.58	4.81	3.12
	MOS	4.65 $\pm$ 0.08	4.58 $\pm$ 0.08	4.47 $\pm$ 0.09

#### 4.3. Ablation Studies

To verify the effectiveness of multi-head vector quantization and multi-stage representation, MSMC-TTS systems based on three versions of MSMCRs are compared in terms of completeness, compactness, and TTS quality:

- V1: 1-stage representation with 1-head quantization.
- V2: 1-stage representation with 4-head quantization.
- V3: 2-stage representation with 4-head quantization and down-sample rates of [1, 4].

Compactness is represented by the compression ratio over the 80-dim Mel spectrogram. Completeness is evaluated in terms of the MCD, F0-RMSE, and V/UV error rates between the ground-truth audio and the reconstructed audio in analysis-synthesis. Finally, an MOS test is conducted to evaluate their performance in standard single-speaker TTS.

Table 3 shows the experimental results. V1, quantized by only one codebook, has the highest compactness with the highest compression ratio of 284. But its completeness is also the worst with the highest MCD of 2.79 dB, and the reconstructed audio cannot even present accurate pronunciations. Hence, its performance in TTS obtains the lowest MOS of 3.74. V2 using multi-head vector quantization has a 4x lower compression ratio of 71, which is still quite high. But the completeness has significantly recovered with much lower MCD, F0-RMSE, and V/UV error rates. Pronunciation and prosody are expressed accurately in reconstruction. The TTS performance is much higher at MOS of 4.36. It shows that improving completeness while keeping high compactness is helpful to raise TTS performance.

Compared with V2, the two-stage representation V3 has slightly reduced compactness at compression ratio of 57. Its completeness is equivalent to V2 with a slightly higher MCD but somewhat reduced loss in F0. However, it achieves the highest MOS of 4.61. It shows obvious improvement on prosody in both coarse-grain and fine-grain levels, which further validates the effectiveness of multi-stage modeling. Hence, both multi-codebooks and multi-stages are crucial for MSMC-TTS.

Table 3: *Comparison of different representations on compactness, completeness, and TTS performance.*

	Compression Ratio	Analysis-Synthesis			TTS
		MCD (dB)	F0-RMSE (Hz)	F0-VUV (%)	MOS (95% CI)
V1	284	2.79	3.53	3.09	3.74 $\pm$ 0.09
V2	71	2.19	2.43	2.38	4.36 $\pm$ 0.08
V3	57	2.22	2.42	2.33	4.61 $\pm$ 0.07

## 5. Conclusion

This paper proposes a MSMC VQ-VAE approach to improve neural TTS. The proposed feature analyzer converts the Mel spectrogram to MSMCRs that are composed of sequences at different time resolutions, and quantized by multiple codebooks, respectively. A multi-stage predictor is trained as the acoustic model to better predict MSMCRs by minimizing the combined loss of MSE and triplet loss. In experiments, MSMC-TTS achieves the highest MOS score of 4.41, which significantly outperforms the baseline system with an MOS of 3.62. Moreover, the proposed approach shows a lower requirement to the modeling complexity, which still performs well with fewer model parameters. Ablation studies validate that multiple stages and multiple codebooks are both crucial for MSMC-TTS.

## 6. Acknowledgments

This work is supported by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Fund. Thanks Guanghui Xu for conducting subjective tests.

## 7. References

- [1] Y. Ijima, T. Asami, and H. Mizuno, "Objective evaluation using association between dimensions within spectral features for statistical parametric speech synthesis." in *Proc. Interspeech*, 2016, pp. 337–341.
- [2] Y. Saito, S. Takamichi, and H. Saruwatari, "Statistical parametric speech synthesis incorporating generative adversarial networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 1, pp. 84–96, 2017.
- [3] H. R. Ihm, J. Y. Lee, B. J. Choi, S. J. Cheon, and N. S. Kim, "Reformer-tts: Neural speech synthesis with reformer network." in *Proc. Interspeech*, 2020, pp. 2012–2016.
- [4] S. Yang, H. Lu, S. Kang, L. Xie, and D. Yu, "Enhancing hybrid self-attention structure with relative-position-aware bias for speech synthesis," in *Proc. ICASSP*. IEEE, 2019, pp. 6910–6914.
- [5] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "Fastspeech: Fast, robust and controllable text to speech," in *Proc. NeurIPS*, 2019, pp. 3165–3174.
- [6] J. Kim, J. Kong, and J. Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *Proc. ICML*. PMLR, 2021, pp. 5530–5540.
- [7] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, "Neural speech synthesis with transformer network," in *Proc. AAAI*, vol. 33, no. 01, 2019, pp. 6706–6713.
- [8] J. Yang, J.-S. Bae, T. Bak, Y.-I. Kim, and H.-Y. Cho, "GANSpeech: Adversarial Training for High-Fidelity Multi-Speaker Speech Synthesis," in *Proc. Interspeech*, 2021, pp. 2202–2206.
- [9] J. Kim, S. Kim, J. Kong, and S. Yoon, "Glow-tts: A generative flow for text-to-speech via monotonic alignment search," in *Proc. NeurIPS*, 2020.
- [10] 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*. IEEE, 2018, pp. 4779–4783.
- [11] P. Neekhara, J. Li, and B. Ginsburg, "Adapting tts models for new speakers using transfer learning," *ArXiv*, vol. abs/2110.05798, 2021.
- [12] H. Guo, H. Lu, X. Wu, and H. Meng, "A multi-scale time-frequency spectrogram discriminator for gan-based non-autoregressive tts," *arXiv preprint arXiv:2203.01080*, 2022.
- [13] M. Verleysen and D. François, "The curse of dimensionality in data mining and time series prediction," in *International work-conference on artificial neural networks*. Springer, 2005, pp. 758–770.
- [14] R. Liu and D. F. Gillies, "Overfitting in linear feature extraction for classification of high-dimensional image data," *Pattern Recognition*, vol. 53, pp. 73–86, 2016.
- [15] B. Labbé, R. Héroult, and C. Chatelain, "Learning deep neural networks for high dimensional output problems," in *Proc. ICMLA*. IEEE, 2009, pp. 63–68.
- [16] C. Gărbacea, A. van den Oord, Y. Li, F. S. Lim, A. Luebs, O. Vinyals, and T. C. Walters, "Low bit-rate speech coding with vq-vae and a wavenet decoder," in *Proc. ICASSP*. IEEE, 2019, pp. 735–739.
- [17] Y. Chen, S. Yang, N. Hu, L. Xie, and D. Su, "Tenc: Low bit-rate speech coding with vq-vae and gan," in *ICMI*, 2021, pp. 126–130.
- [18] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," in *Proc. NeurIPS*, 2017.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, 2017, pp. 5998–6008.
- [20] A. Razavi, A. van den Oord, and O. Vinyals, "Generating diverse high-fidelity images with vq-vae-2," in *Proc. NeurIPS*, 2019, pp. 14 866–14 876.
- [21] R. Child, "Very deep vaes generalize autoregressive models and can outperform them on images," in *Proc. ICLR*, 2021.
- [22] X. Dong and J. Shen, "Triplet loss in siamese network for object tracking," in *Proc. ECCV*, 2018, pp. 459–474.
- [23] S. King and V. Karaiskos, "The blizzard challenge 2011," in *Blizzard Challenge Workshop*, 2011.
- [24] M. He, Y. Deng, and L. He, "Robust sequence-to-sequence acoustic modeling with stepwise monotonic attention for neural tts," *Proc. Interspeech*, pp. 1293–1297, 2019.
- [25] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," in *Proc. NeurIPS*, 2020.
- [26] 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," in *Proc. Interspeech*, 2021.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.