

QS-TTS: Towards Semi-Supervised Text-to-Speech Synthesis via Vector-Quantized Self-Supervised Speech Representation Learning

Haohan Guo, *Student Member, IEEE*, Fenglong Xie, *Member, IEEE*,
 Jiawen Kang, *Student Member, IEEE*, Yujia Xiao, *Student Member, IEEE*,
 Xixin Wu, *Member, IEEE*, Helen Meng, *Fellow, IEEE*

Abstract—This paper proposes a novel semi-supervised TTS framework, QS-TTS, to improve TTS quality with lower supervised data requirements via Vector-Quantized Self-Supervised Speech Representation Learning (VQ-S3RL) utilizing more unlabeled speech audio. This framework comprises two VQ-S3R learners: first, the principal learner aims to provide a generative Multi-Stage Multi-Codebook (MSMC) VQ-S3R via the MSMC-VQ-GAN combined with the contrastive S3RL, while decoding it back to the high-quality audio; then, the associate learner further abstracts the MSMC representation into a highly-compact VQ representation through a VQ-VAE. These two generative VQ-S3R learners provide profitable speech representations and pre-trained models for TTS, significantly improving synthesis quality with the lower requirement for supervised data. QS-TTS is evaluated comprehensively under various scenarios via subjective and objective tests in experiments. The results powerfully demonstrate the superior performance of QS-TTS, winning the highest MOS over supervised or semi-supervised baseline TTS approaches, especially in low-resource scenarios. Moreover, comparing various speech representations and transfer learning methods in TTS further validates the notable improvement of the proposed VQ-S3RL to TTS, showing the best audio quality and intelligibility metrics. The trend of slower decay in the synthesis quality of QS-TTS with decreasing supervised data further highlights its lower requirements for supervised data, indicating its great potential in low-resource scenarios.

I. INTRODUCTION

Text-to-Speech (TTS) synthesis is a technology aiming to convert text to speech signals with correct pronunciation, natural prosody, and high audio fidelity. It has been widely used in various intelligent products, e.g. Human-Computer Interaction (HCI) [1, 2], Speech-to-Speech Translation (S2ST) [3, 4], and Artificial Intelligence Generated Content (AIGC) [5]. Meanwhile, with the popularization and penetration of AI

Hanhan Guo, Jiawen Kang, and Helen Meng are with the Human-Computer Communications Laboratory (HCCL), Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Hong Kong SAR, China (e-mail: hguo@se.cuhk.edu.hk; jwkang@se.cuhk.edu.hk; hmmeng@se.cuhk.edu.hk). Xixin Wu is with Stanley Ho Big Data Decision Analytics Research Centre, The Chinese University of Hong Kong, Hong Kong SAR, China (e-mail: xixinwu@cuhk.edu.hk). Yujia Xiao is with the Digital Signal Processing and Speech Technology Laboratory (DSP-STL), Department of Electronic Engineering, The Chinese University of Hong Kong, Hong Kong SAR, China (e-mail: yujiaxiao@link.cuhk.edu.hk). Fenglong Xie is with Xiaohongshu, Beijing, China (e-mail: fenglongxie@xiaohongshu.com). Corresponding author: Helen Meng.

This research was supported by the Center for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission’s InnoHK scheme.

technology into various fields, the capability of a TTS system for personalization and customization has also been paid more attention, to better serve all people around the world. However, the high cost of creating a TTS dataset with sufficient high-quality speech audio and accurate transcripts hinders the development of TTS in this aspect. Hence, reducing the supervised data requirements of TTS is becoming more urgent for low-resource scenarios, which is also the goal of our work.

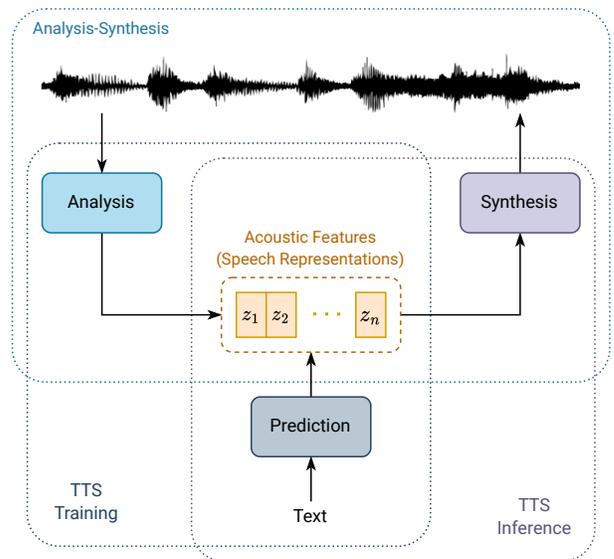


Fig. 1. A mainstream TTS framework: the acoustic feature or speech representation is extracted from the waveform via the analysis module as the output of the prediction module, and decoded to the waveform by the synthesis module.

Fortunately, a TTS system does not rely entirely on supervised data. It is feasible to utilize more unlabeled speech data that is easier to collect to compensate for insufficient supervised data in TTS training. For example, in a mainstream modular TTS framework, as shown in Fig. 1, we can employ unlabeled speech data to

- Enhance the analysis module to provide practical data-derived speech representations with sufficient phonetic information and easier to predict from the text.
- Enhance the synthesis and prediction modules by pre-training them on relevant tasks.

In this way, we may reduce supervised data requirements by

building a semi-supervised TTS system using more unlabeled speech data. The key is seeking an appropriate approach or task utilizing unlabeled speech audio well to provide the expected features and pre-trained models to TTS.

In this regard, Self-Supervised Speech Learning (S3RL) has shown excellent capability in providing profitable speech representations and pre-trained models for various supervised speech tasks. Especially in speech synthesis [6, 7], various vector-quantization-based [8] S3RL approaches have already demonstrated their eminent performance, providing compact, discrete speech representations to reduce modeling complexity while keeping high-quality speech reconstruction. However, there is still a lack of a systematic framework covering these approaches to improve low-resource TTS comprehensively. To address this issue, in this paper, we propose QS-TTS, a novel semi-supervised TTS framework based on Vector-Quantized Self-Supervised Speech Representation Learning (VQ-S3RL).

This framework conducts VQ-S3RL on unlabeled speech data for two goals:

- Providing profitable speech representations for TTS.
- Providing effective pre-trained models to enhance the acoustic model and the vocoder.

It comprises two learners. First, we train a principal VQ-S3R learner, which combines the contrastive S3RL model, HuBERT [9], and the proposed generative S3RL model, Multi-Stage Multi-Codebook (MSMC) VQ-GAN in cascade. It converts the waveform into the MSMCR, a compact generative S3R comprising multiple sequences at different time resolutions and quantized by multiple codebooks, which is then decoded to high-quality speech audio via adversarial speech generation. Then, we train an associate VQ-S3R learner to abstract the MSMCR into a highly compact VQ representation through a VQ-VAE-based model with multi-sequence encoding and decoding. These two generative VQ-S3R learners are applied in TTS training using supervised data, including playing the role of the analysis module providing the MSMCR, and providing pre-trained models for the prediction and synthesis module training.

In experiments, in addition to MOS tests to evaluate QS-TTS subjectively, we also measure the performance of QS-TTS objectively on audio quality using Fréchet Distances (FD) in various embedding spaces and the intelligibility using the Character Error Rate (CER) and Phoneme Error Rate (PER). We first evaluate the overall performance of QS-TTS by comparing it with mainstream TTS approaches, e.g. FastSpeech, VITS, and their semi-supervised versions, in standard-resource and various low-resource scenarios. Then, we investigate the effect of the proposed VQ-S3RL on TTS by comparing it with different speech representations and transfer learning methods. Finally, we evaluate the performance of QS-TTS with different sizes of supervised data to further validate its effectiveness in reducing supervised data requirements.

Our contributions are summarized as follows: 1) we propose QS-TTS, a novel semi-supervised TTS approach, significantly outperforming supervised / semi-supervised baseline systems in low-resource scenarios; 2) we apply VQ-S3RL with only unlabeled data to provide profitable speech representations and

effective pre-trained model for low-resource TTS, more effectively reducing the supervised data requirements over other semi-supervised approaches. 3) we conduct comprehensive subjective and objective evaluations to demonstrate the effectiveness of QS-TTS across different low-resource scenarios, validating its lower requirements for supervised data. In the rest of this paper, Section II introduces the background of this work. Section III illustrates the framework of QS-TTS in detail. Experiments are described in Sections IV and V. Finally, Section VI gives the conclusion to this paper.¹

II. BACKGROUND

A. Semi-Supervised Text-to-Speech Synthesis

The semi-supervised TTS aims to utilize both supervised and unsupervised data in training to improve TTS quality. It is usually achieved by transfer learning, i.e. pre-training TTS modules with only audio or text, and fine-tuning it using supervised data. For example, in Tacotron-based TTS [10, 11], the auto-regressive decoder can be pre-trained with unlabeled audio [12, 13], then applied in supervised training to achieve better prediction quality. In [14], the phoneme encoder is also pre-trained using only text in a BERT-like [15] way. We can also tag unlabeled speech audio using limited supervised data. For example, speech chain [16], i.e. back-translation [17], can train Automatic Speech Recognition (ASR) and TTS iteratively using only a few paired data and much-unpaired data. In [18, 19], a weak ASR system is trained with a few minutes of supervised data, then decodes all unlabeled audio to generate transcripts to create a low-accuracy supervised dataset. TTS is directly pre-trained by this dataset, and then fine-tuned with the target supervised data. Besides, we can also create pseudo-labels via unsupervised acoustic unit discovery [20, 21] to form a pseudo-supervised dataset for TTS pre-training. Eventually, employing more unlabeled data in training can enhance the TTS system with lower requirements for supervised data.

However, with the development of representation learning, data-derived speech representations show better performance than conventional, signal-processing-based acoustic features in TTS, which can be better predicted from the text while keeping sufficient speech information for high-quality audio reconstruction. It effectively reduces the requirement for supervised data, indicating a new direction for semi-supervised TTS.

B. Self-Supervised Speech Representation Learning

S3RL [22] aims to learn useful representations from unlabeled speech data to serve downstream supervised speech tasks. It is usually divided into two categories: contrastive and generative [23]. Contrastive S3RL models, e.g. Wav2Vec [24–26], HuBERT [9], WavLM [27], are usually encoder-based models trained with contrastive loss functions [28]. This kind of model is more robust to noisy data in training, hence can be applied with massive low-quality speech data to

¹Audio samples are available at <https://hhguo.github.io/DemoQSTTS/>. The code is available at <https://github.com/hhguo/MSMC-TTS>

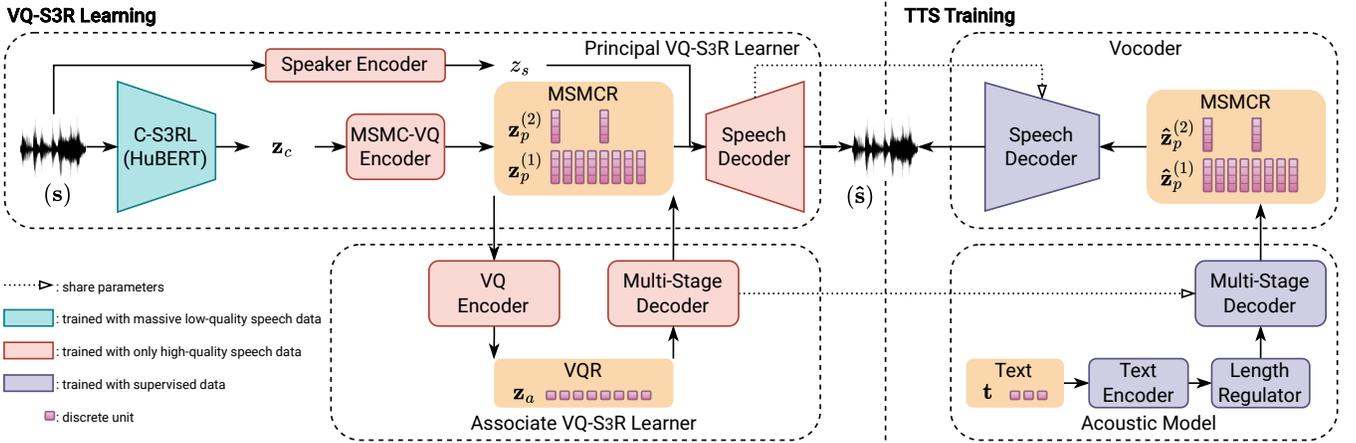


Fig. 2. The framework of QS-TTS. In vector-quantized self-supervised speech representation learning (VQ-S3RL), the principal VQ-S3R learner first converts the speech signal s into z_c via a pre-trained contrastive S3RL model (HuBERT), and then feed it to the Multi-Stage Multi-Codebook (MSMC) VQ-GAN and ECAPA-TDNN-based speaker encoder to obtain the MSMCR z_p and the speaker embedding z_s , finally reconstruct the waveform \hat{s} via the speech decoder. The associate VQ-S3R learner compresses z_p into the VQ sequence z_a via a VQ-VAE model. In TTS training, the vocoder and acoustic model are trained based on the pre-trained speech decoder and multi-stage decoder to map the text to the MSMCR \hat{z}_p , and then synthesize the waveform \hat{s} .

learn a general speech representation [29] to enhance various downstream speech classification [30–34] and synthesis [35, 36] tasks. Generative S3RL aims to reconstruct speech in training while applying restrictions to the latent space, which is more compatible with TTS intuitively due to their overlapped goal of speech generation. Hence, it is widely applied in TTS tasks, such as the Auto-regressive model, VAE [37], and VQ-VAE [8]. It cannot only provide an effective speech representation [38], but provide a good pre-trained model for TTS training [12]. However, this kind of approach is more sensitive to noisy data due to the reconstruction objective [23], hence having a higher requirement for audio quality, lacking good generalization for low-quality audio.

C. Vector-Quantized Representation Learning

As a generative self-supervised learning method, vector-quantized representation learning aims to compress the target via vector quantization into a compact, discrete representation, while keeping high-quality reconstruction, e.g. VQ-VAE [8] and its enhanced version, VQ-GAN [39] and VQ-Diffusion [40]. It has been widely applied in speech generation, e.g. speech coding [41–43], Voice Conversion (VC) [44], and TTS [7, 45–47]. To extract a better VQ speech representation with a balance between compactness and completeness, MSMC-TTS [7] proposes a Multi-Stage Multi-Codebook Representation (MSMCR), comprising multiple sequences with different time resolutions and quantized by multiple codebooks. It can be better predicted from the text via multi-stage modeling, significantly improving TTS performance with lower supervised data requirements, further showing the great potential of VQ representations in semi-supervised TTS.

III. QS-TTS

QS-TTS is a semi-supervised TTS framework based on vector-quantized self-supervised speech representation learning (VQ-S3RL). As shown in Fig. 2, it trains two VQ-S3R

learners on unlabeled speech data to provide more profitable speech representations and effective pre-trained models to enhance supervised TTS training, thereby improving synthesis quality while reducing the supervised data requirement. In this section, we will illustrate each module in detail.

A. The Principal VQ-S3R Learner

We first propose a principal VQ-S3R learner combining contrastive and generative S3RL models to extract an effective speech representation, which is easier to predict from the text and well-reconstructed into high-quality audio. It first employs a contrastive S3RL model, HuBERT [9], trained with massive speech audio to extract an effective general speech representation z_c from the speech signal s . Then, it conducts the generative VQ-S3RL based on MSMC-VQ-GAN using only high-quality speech data to convert z_c into the generative VQ-S3R, MSMCR, while decoding it with the speaker embedding z_g to high-quality audio \hat{s} .

In this section, we will introduce the model architecture and training method of MSMC-VQ-GAN.

1) *Model architecture*: The model architecture of MSMC-VQ-GAN is mainly composed of an MSMC-VQ encoder and a speech decoder. Fig. 3 shows an example of a two-stage four-codebook VQ-GAN model. In the MSMC-VQ encoder, the input z_c is first processed by a Transformer encoder composed of a linear layer and a feedforward Transformer block, then quantized in two stages. In the higher stage, i.e. stage 2, the input sequence is down-sampled 4 times along the time axis via the down-sample module, which has an average pooling layer and a feedforward Transformer block. Then, the higher-stage sequence $\tilde{z}_p^{(2)}$ is quantized by a four-head codebook $c_p^{(2)}$ via Multi-Head Vector Quantization (MHVQ) [7], i.e. product quantization [48], which chunks the codebook into multiple sub-codebooks to quantize the input vector chunked in the same way, respectively. The quantized output $z_p^{(2)}$ is further processed by an up-sample block comprising two MLP layers

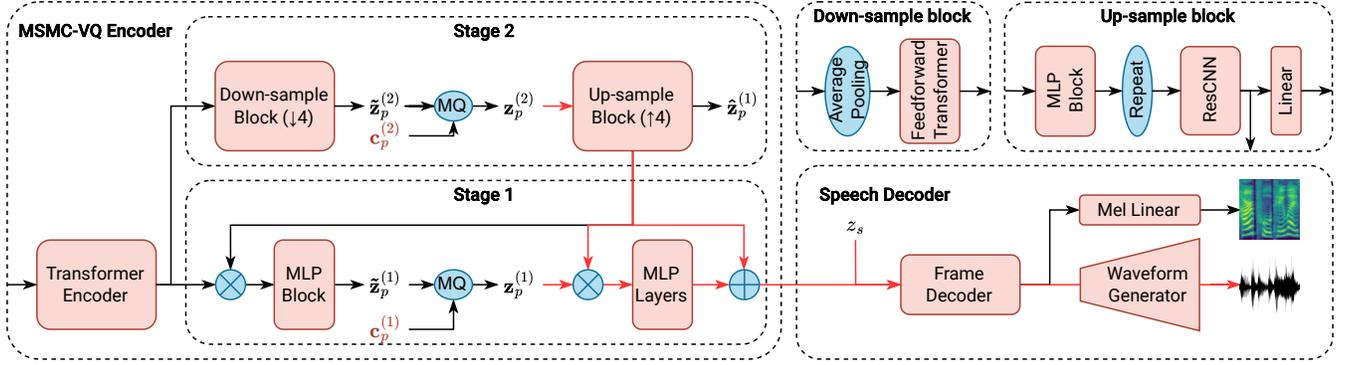


Fig. 3. The model architecture of a two-stage four-codebook VQ-GAN generator. “MQ”, “ \oplus ”, “ \otimes ” denote multi-head vector quantization, addition, and concatenation operations. The red arrows emphasize the process of converting the MSMCR into the waveform at inference.

with a LeakyReLU activation function in between, repetition operation for up-sampling, and four residual convolutional layers, to help the following quantization, and predict the lower-stage (stage-1) quantized sequence $\hat{z}_p^{(1)}$. In the lower stage, we obtain the quantized sequence $z_p^{(1)}$ with the guidance of the high-stage information, and add it with the high-stage residual output as the encoder output, which is then fed to the speech decoder to generate the waveform.

The speech decoder is composed of a frame decoder and a waveform generator. To better support multi-speaker waveform generation, we extract an utterance-level embedding from the Mel spectrogram of the input audio via the global encoder based on ECAPA-TDNN [49] as the speaker representation z_s . We add the input sequence with z_s , and then employ the frame decoder with a feedforward Transformer block to process it, finally generating the speech waveform \hat{s} via the waveform generator based on the Hifi-GAN model. Meanwhile, like MSMC-VQ-VAE, we still predict the Mel spectrogram from the output of the frame decoder using a Mel linear layer.

2) *Loss function*: The training objective of MSMC-VQ-GAN is composed of multiple loss terms. First, due to the non-differentiable VQ operations, we back-propagate the gradient to the two-stage MSMC-VQ encoder via the following loss:

$$\mathcal{L}_{vq} = \|\bar{z}_p^{(1)} - sg(z_p^{(1)})\|_2^2 + \|\bar{z}_p^{(2)} - sg(z_p^{(2)})\|_2^2 \quad (1)$$

where S denotes the number of stages, and $sg(*)$ is the stop-gradient operation, i.e. no gradient is back-propagated from the variable. We adopt the exponential moving average-based method [50] to update codebooks in training. For effective multi-stage representation learning, we also enhance the relationship between adjacent stages with the following loss term:

$$\mathcal{L}_{ms} = \|\hat{z}_p^{(1)} - sg(z_p^{(1)})\|_2^2 \quad (2)$$

It can help the higher stage learn an effective representation stably [7, 51].

To reconstruct high-quality speech audio, we apply GAN training to the model with a UnivNet discriminator [52], composed of multiple sub-discriminators for multi-resolution spectrogram discriminating and multi-period waveform discriminating, to capture more discriminative information in

frequency and time domains. The loss function for the discriminator is written as:

$$\mathcal{L}_d = \frac{1}{K} \sum_{k=1}^K [(D_k(s) - 1)^2 + D_k(\hat{s})^2] \quad (3)$$

where K denotes the number of sub-discriminators. And the adversarial loss for the MSMC-VQ-GAN is written as:

$$\mathcal{L}_{adv} = \frac{1}{K} \sum_{k=1}^K [(D_k(\hat{s}) - 1)^2] \quad (4)$$

To enhance GAN training quality, the Mel-spectrogram loss and feature matching loss, widely used in GAN-based neural vocoder training, are also employed in MSMC-VQ-GAN training [53]. Mel-spectrogram loss is the L1 distance between two waveforms in the Mel-scale frequency domain, which can improve the perceptual quality of the generated audio. It is written as follows:

$$\mathcal{L}_{mel} = \|\phi(s) - \phi(\hat{s})\|_1 \quad (5)$$

where ϕ denotes the operating converting the waveform into the log-scale Mel spectrogram. Feature matching loss can further improve GAN training quality by reducing differences between the ground-truth waveform and the generated waveform in the hidden feature space of the discriminator as follows:

$$\mathcal{L}_{fm} = \frac{1}{K} \sum_{k=1}^K \frac{1}{N_k} \sum_{i=1}^{N_k} \|(D_k^{(i)}(s)) - D_k^{(i)}(\hat{s})\|_1 \quad (6)$$

where N_k denotes the number of hidden layers of k -th sub-discriminator, and $D_k^{(i)}(*)$ denotes the output feature of i -th layer of k -th sub-discriminator.

Finally, to avoid the negative impact of the unstable performance of GAN to VQ-S3RL in the early-stage training, we still apply the frame-level reconstruction loss, which is an L2 distance between the predicted Mel spectrogram and the ground-truth one:

$$\mathcal{L}_{frame} = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 \quad (7)$$

where \mathbf{x} and $\hat{\mathbf{x}}$ denote the ground-truth and predicted Mel spectrograms. Finally, the loss function for MSMC-VQ-GAN is written as follows:

$$\begin{aligned} \mathcal{L}_g = & \mathcal{L}_{adv} + \lambda_{fm} * \mathcal{L}_{fm} + \lambda_{mel} * \mathcal{L}_{mel} \\ & + \lambda_{vq} * \mathcal{L}_{vq} + \lambda_{ms} * \mathcal{L}_{ms} + \lambda_{frame} * \mathcal{L}_{frame} \end{aligned} \quad (8)$$

where $\lambda_{fm}, \lambda_{mel}, \lambda_{vq}, \lambda_{ms}, \lambda_{frame}$ are weight coefficients.

B. The Associate VQ-S3R Learner

To better predict the MSMCR from the text, we propose an associate VQ-S3R learner to provide an effective pre-trained model for training the acoustic model. It imitates the process of converting the text, a highly-compact discrete sequence, into the MSMCR, by employing a VQ-VAE-based model to abstract the MSMCR into a more compact VQ sequence and reconstruct it back.

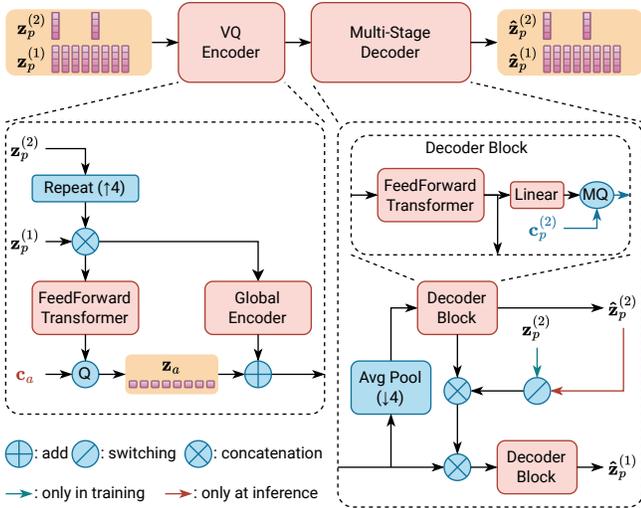


Fig. 4. The model structure of VQ-VAE for the two-stage four-codebook representation, where the second-stage sequence has a down-sample rate of 4 along the time axis.

1) *Model architecture*: Fig. 4 shows the VQ-VAE model for the two-stage representation, where the stage-2 sequence $\mathbf{z}_p^{(2)}$ has a down-sample rate of 4. It is composed of a VQ encoder and a multi-stage decoder. In the VQ encoder, all sequences are first up-sampled to the same length as the stage-1 sequence $\mathbf{z}_p^{(1)}$, e.g. repeating all vectors in $\mathbf{z}_p^{(2)}$ for 4 times. Then, aligned sequences are concatenated to be processed by a feedforward Transformer block, and quantized by one codebook to obtain a highly-compact VQ sequence \mathbf{z}_a . To force it to capture more phonetics-related information, we still apply the ECAPA-TDNN-based global encoder to extract an utterance-level speaker embedding, which is jointly trained with the VQ-VAE. This embedding is then up-sampled to the frame level by repetition and added with the quantized sequence \mathbf{z}_a for decoding.

The multi-stage decoder aims to reconstruct the MSMCR from \mathbf{z}_a . It predicts sequences from high to low stages in a cascaded way. For this two-stage representation, the decoder first down-samples the encoder output 4 times, then feeds it to the decoder block comprising a feedforward Transformer

block followed by a linear layer. The output of the linear layer refers to the predicted stage-2 sequence $\hat{\mathbf{z}}_p^{(2)}$ in training, and is quantized by the codebook $\mathbf{c}_p^{(2)}$ from the MSMC-VQ-GAN as the output. Similarly, in stage 1, $\hat{\mathbf{z}}_p^{(1)}$ is predicted by giving the encoder output and the stage-2 outputs. Notably, in training, we replace $\hat{\mathbf{z}}_p^{(2)}$ with the ground-truth quantized sequence $\mathbf{z}_p^{(2)}$ as the input.

2) *Loss function*: The model is trained with the loss function \mathcal{L}_a written as follows:

$$\begin{aligned} \mathcal{L}_{vq} = & \|\bar{\mathbf{z}}_a - sg(\mathbf{z}_a)\|_2^2 \\ \mathcal{L}_{rec} = & \|\hat{\mathbf{z}}_p^{(1)} - \mathbf{z}_p^{(1)}\|_2^2 + \|\hat{\mathbf{z}}_p^{(2)} - \mathbf{z}_p^{(2)}\|_2^2 \\ \mathcal{L}_a = & \mathcal{L}_{vq} + \lambda_{rec} * \mathcal{L}_{rec} \end{aligned} \quad (9)$$

where \mathcal{L}_{vq} is still the VQ loss for the VQ encoder training, \mathcal{L}_{rec} is the reconstruction loss between the ground-truth and reconstructed MSMCR with 2 stages, and λ_{rec} is a weight coefficient. In training, we still use the exponential moving average-based method to update the codebook \mathbf{c}_a .

C. TTS Synthesis

In TTS synthesis, we aim to convert the text to its corresponding MSMCR via the acoustic model, and then generate the waveform from the MSMCR via the vocoder. The acoustic model has the same architecture as MSMC-TTS [7], based on the FastSpeech [54]. It encodes the text sequence using a Transformer encoder, then up-samples it via repetition according to the predicted duration, finally generating the MSMCR $\hat{\mathbf{z}}_q$ via a multi-stage decoder. In training, it inherits parameters of the pre-trained multi-stage decoder in the associate VQ-S3R learner, and then is trained with the supervised data. The training loss function is written as follows:

$$\begin{aligned} \mathcal{L}_{dur} = & \|\hat{\mathbf{d}} - \mathbf{d}\|_2^2 \\ \mathcal{L}_{am} = & \mathcal{L}_{rec} + \lambda_{dur} * \mathcal{L}_{dur} \end{aligned} \quad (10)$$

where \mathcal{L}_{dur} denotes the duration loss between the ground-truth duration \mathbf{d} and the predicted duration $\hat{\mathbf{d}}$, and λ_{dur} is a weight coefficient.

In vocoder, to convert the predicted $\hat{\mathbf{z}}_p$ composed of multiple sequences to the waveform, we first feed it to the MSMC-VQ encoder in the pre-trained MSMC-VQ-GAN to obtain the encoder output sequence, as indicated by the red arrows shown in Fig. 3. Then, we synthesize the waveform via the speech decoder fine-tuned with the audio in the supervised dataset to adapt the target speaker better. In fine-tuning, we only update the parameters of the speech decoder using the same training configurations of MSMC-VQ-GAN.

IV. EXPERIMENTAL PROTOCOL

A. Dataset

In VQ-S3RL, we employ AISHell-3, a Mandarin multi-speaker high-quality speech dataset, as the training set. This dataset contains roughly 85 hours of emotion-neutral recordings spoken by 218 native Chinese Mandarin speakers, showing a rich coverage of phonetic and speaker information. In TTS training, we employ multiple supervised TTS datasets to evaluate TTS systems under various scenarios. The first

dataset is *CSMSC*², a single-speaker Mandarin TTS corpus with 10-hour high-quality supervised data, which is widely applied in training the standard single-speaker Mandarin TTS system. We also conduct low-resource scenarios by extracting subsets from *CSMSC*. Besides, a test set with 200 utterances is also extracted from this dataset but has no overlap with all training sets. Then, we construct a more challenging low-resource scenario using only 10 minutes of child speech data spoken by a five-year-old girl in Mandarin³. It has a test set with 24 utterances out of the training set. Finally, we use an internal Cantonese dataset with 15 minutes of supervised data to evaluate the performance of TTS systems in low-resource languages. The Cantonese test set comprises 134 utterances out of the training set.

B. Feature

First, in audio processing, all single-channel audio used in our work is down-sampled to the sample rate of 16kHz. Then, we extract Mel spectrograms for all datasets in the following way: first, pre-emphasize the audio with the coefficient of 0.97, and then convert it to the 1025-dim magnitude spectrograms by STFT with a window length of 50ms, a frameshift of 12.5ms, and an FFT size of 2048, finally compress the spectrogram into the 80-dim log-scale Mel spectrogram. We employ a HuBERT⁴ [9], the contrastive S3RL model, pre-trained on WenetSpeech [55], a Mandarin dataset with around 10,000 hours of speech audio, to extract the general speech representation. The HuBERT feature is a sequence of 1024-dim vectors with a frameshift of 20ms. To align it to the Mel spectrogram, we up-sample it via nearest neighbor interpolation to a frameshift of 12.5ms.

In text processing, we convert the text to phonemes as the input of the acoustic model. For *CSMSC* and child datasets, we directly employ the phonemes and their corresponding duration labeled in the dataset for training. For the Cantonese dataset, we use an open-source G2P tool [56] to obtain phonemes, and train a Montreal Forced Aligner⁵ (MFA) model to obtain the phoneme-level duration for training.

C. Model Configuration

In VQ-S3RL, we apply the MSMC-VQ-GAN with 2 stages, where the second stage has a down-sample rate of 4 along the time axis. And in each stage, a 4-head codebook is used for vector quantization, where each head is composed of 64 codewords with a dimension of 64. In the MSMC-VQ encoder, we apply the 4-layer 256-dim feedforward Transformer block with 2-head self-attention to the Transformer encoder and the down-sample block. The MLP block has two 256-dim linear layers with a Tanh activation function. The residual CNN block comprises 4 1-D residual convolutional layers with a

kernel size of 5. In the speech decoder, the frame decoder is also implemented with the 4-layer 256-dim feedforward Transformer block. The waveform generator is a Hifi-GAN-V1 [53] generator, which upsamples the input sequence 200 times to the 16kHz waveform via 4 CNN-based upsampling blocks with the upsample rates of [5, 5, 4, 2] and the kernel sizes of [11, 11, 8, 4].

In GAN training based on the UnivNet discriminator (UnivNet-c32) [52], we extract three magnitude spectrograms from the waveform using three STFT parameter sets, FFT size [256, 512, 1024], frameshift [40, 80, 160], and frame length [120, 320, 640], for multi-resolution spectrogram discriminating. And we also reshape the 1-D waveform into five 2-D sequences with the period of [2, 3, 5, 7, 11] for multi-period waveform discriminating⁶. In this work, MSMC-VQ-GAN is trained on AIShell-3 for 400k iterations using the AdamW optimizer ($\beta_1 = 0.8, \beta_2 = 0.99$) with a batch size of 16 utterances. Similar to random window discriminating [57], we also randomly select a segment with a length of 0.75 seconds from each utterance for adversarial training of the waveform generator to improve training efficiency. The learning rate of 2×10^{-4} exponentially decays with the rate $2^{-\frac{1}{200,000}}$ after 200k warm-up iterations. The weight coefficients $\lambda_{fm}, \lambda_{mel}, \lambda_{vq}, \lambda_{ms}, \lambda_{frame}$ are set to 2, 45, 10, 1, 450, respectively. To stabilize the training process, we also apply warm-up training here, i.e. no GAN training in the first 50k iterations.

The associate VQ-S3R learner compresses the MSMCR with one codebook with 64 256-dim codewords. The VQ encoder is a Transformer encoder with the same configuration as MSMC-VQ-GAN. And the global encoder is an ECAPA-TDNN [49] model with 128 channels for hidden layers. This model is trained on AIShell-3 for 200k iterations using Adam [58] optimizer ($\beta_1 = 0.9, \beta_2 = 0.98$) with the batch size of 64 utterances, and the learning rate of 2×10^{-4} exponentially decayed with the rate $2^{-\frac{1}{20,000}}$ after 20k warm-up iterations.

Finally, in TTS training, we use the supervised TTS dataset to fine-tune the speech decoder of MSMC-VQ-GAN with the same training configuration, and train the acoustic model with $\lambda_{dur} = 0.1$ based on the pre-trained multi-stage decoder in the associate VQ-S3R learner. For a standard supervised dataset with sufficient audio, i.e. 10-hour *CSMSC*, we fine-tune the vocoder for 400k iterations and the acoustic model for 200k iterations. Otherwise, for all low-resource datasets, we only fine-tune the vocoder for 100k iterations, and the acoustic model for 50k iterations to avoid over-fitting.

D. Baselines

In our experiments, we implement multiple fully-supervised and semi-supervised TTS approaches following the training configuration of QS-TTS.

1) *FastSpeech*: It is a mainstream non-autoregressive neural TTS system based on the Mel spectrogram. It comprises an acoustic model based on Transformer blocks and a Hifi-GAN

²*CSMSC* is available at <https://www.data-baker.com/data/index/source>.

³The child TTS dataset is available at <https://magichub.com/datasets/mandarin-chinese-speech-corpus-for-tts-children-speech>.

⁴The pre-trained HuBERT model is available at <https://huggingface.co/TencentGameMate/chinese-hubert-large>.

⁵The tool is available at <https://github.com/MontrealCorpusTools/Montreal-Forced-Aligner>.

⁶The implementations of Hifi-GAN generator and UnivNet discriminator are available at <https://github.com/jik876/hifi-gan> and <https://github.com/mindslab-ai/univnet>.

vocoder, having the same model hyperparameters as QS-TTS. We upgrade it to two semi-supervised versions: FastSpeech-S and FastSpeech-SS. In FastSpeech-S, the Mel-spectrogram-based Hifi-GAN vocoder is pre-trained on AI-Shell3, and then fine-tuned with the supervised dataset. In FastSpeech-SS, the Mel spectrogram is replaced with HuBERT features, and the HuBERT-based Hifi-GAN vocoder is also pre-trained on AI-Shell3.

2) *VITS* [59]: It is the SOTA end-to-end TTS system based on VAE representations. It applies a VAE-GAN model based on Hifi-GAN to learn a generative speech representation, which can be predicted from the text by a Glow-based module [60]. We implement it based on the official model configuration⁷, and apply it with the waveform generator and the discriminator with the same configuration as that of QS-TTS for a fair comparison. We also implement its semi-supervised version, VITS-SS [21], in our experiments. It first applies k-means clustering with 512 centroids on HuBERT features of AIShell-3 to create pseudo-labels for all unlabeled audio, then pre-trains VITS on this created paired dataset, and finally fine-tunes VITS using the supervised dataset.

3) *MSMC-TTS*: It can be seen as the fully-supervised version of QS-TTS, which learns the MSMCR from the Mel spectrogram via the MSMC-VQ-GAN model, and is trained with only the supervised dataset. Besides, We also implement another semi-supervised version of MSMC-TTS, MSMC-TTS-SS, which pre-trains the Mel-spectrogram-based MSMC-VQ-GAN on AIShell-3, and only fine-tunes the speech decoder in TTS training.

4) *Back Translation*: We also implement this semi-supervised training approach to pre-train the acoustic model. It first employs the supervised TTS dataset to train an ASR system, and then transcribe unlabeled speech audio of AIShell-3 to create a low-precision but large-scale paired dataset. The acoustic model is pre-trained on this dataset, and finally fine-tuned with the supervised TTS dataset. This approach highly relies on the quality of ASR to ensure precise transcription, so we train a CTC-ASR based on the pre-trained Mandarin HuBERT model to convert the audio into phonemes and their corresponding durations extracted from CTC alignments. Besides, we also implement the k-means-based approach proposed in VITS-SS to avoid training ASR models. It directly quantize HuBERT features of AIShell-3 into 512 codewords by k-means to obtain pseudo-labels and corresponding durations, which are then applied to pre-train the acoustic model.

E. Evaluation

In this work, we conduct objective and subjective tests to evaluate the proposed TTS approach comprehensively.

1) *Objective Metrics*: Except for Mel Cepstral Distortion (MCD) [61], computing the perceptual difference between two fully-aligned audio files, we also propose evaluating the synthesis quality using Frèchet distances [62] in various embedding spaces. Frèchet distance can measure the distance

between two sets of samples by calculating the difference between them in distributions as follows:

$$\mathbb{F}(\mathcal{N}_s, \mathcal{N}_t) = \|\mu_s - \mu_t\|^2 + \text{tr}(\Sigma_s + \Sigma_t - 2\sqrt{\Sigma_s \Sigma_t}) \quad (11)$$

where \mathcal{N}_s and \mathcal{N}_t denote the normal distribution of the synthesized audio set, and the target ground-truth audio set in the embedding space. And we can use different audio classification models to extract embeddings and calculate their mean vectors (μ_s, μ_t) and covariance matrices (Σ_s, Σ_t) . This work employs three embedding spaces: acoustics, speaker, and phonetics. The Frèchet distance in acoustic space, i.e. Frèchet Audio Distance (FAD) [63], has been well applied in evaluating the synthesis quality of neural vocoder, which extracts the embedding for each 4-second audio from an audio classification model. Similarly, we extract utterance-level speaker embeddings using an ECAPA-TDNN-based speaker verification model, and extract utterance-level phonetic embeddings using a Transformer-based ASR model by averaging its encoder output sequence into one vector. These three distances are denoted as FD-AC, FD-SV, and FD-ASR in our following work⁸. Notably, the Frèchet distance in speaker space is multiplied by 10 to align with other distances.

We also evaluate intelligibility, the most crucial factor in evaluating a TTS system, by calculating the Character Error Rate (CER) and Phoneme Error Rate (PER) by transcribing the synthesized audio using ASR tools⁹. And we use G2P tools¹⁰ to convert the transcribed and ground-truth text into phonemes to calculate PER, which focuses more on the pronunciation accuracy of phonemes.

2) *Subjective Metrics*: We conduct MOS (mean opinion score) tests to subjectively evaluate TTS systems on synthesis quality. In each MOS test, 10 native speakers are hired to rate each audio sample in 20 test cases, where each test case contains multiple audio samples synthesized by different TTS approaches but from the same text. The rating ranges from 1 to 5 with an increment of 0.5, where the higher score indicates better quality. Finally, we statistic the scores of each method to obtain the MOS with a 95% confidence interval.

V. RESULTS

A. TTS comparison: Semi-Supervised v.s. Supervised

First, we investigate the impact of unlabeled speech data on TTS by comparing the proposed semi-supervised TTS system with supervised TTS systems: FastSpeech, VITS, and MSMC-TTS.

Table I shows the MOS test result of standard single-speaker TTS systems on the 10-hour CSMSC. FastSpeech, the modular TTS system based on the Mel Spectrogram, shows the worst

⁸The pre-trained audio classification, speaker verification, and ASR models are available at <https://github.com/harrityaylor/torchvggish>, <https://huggingface.co/speechbrain/spkrec-ecapa-voxceleb>, and <https://github.com/openai/whisper/tree/main> (multi-lingual base version).

⁹The ASR tools for Mandarin and Cantonese are available at <https://github.com/wenet-e2e/wenet/tree/main/runtime/binding/python> and <https://huggingface.co/Scrya/whisper-large-v2-cantonese>.

¹⁰The G2P tools for Mandarin and Cantonese are available at <https://github.com/mozillazg/python-pinyin> and <https://github.com/mozillazg/python-pinyin>.

⁷The official implementation of VITS is available at <https://github.com/jaywalnut310/vits>

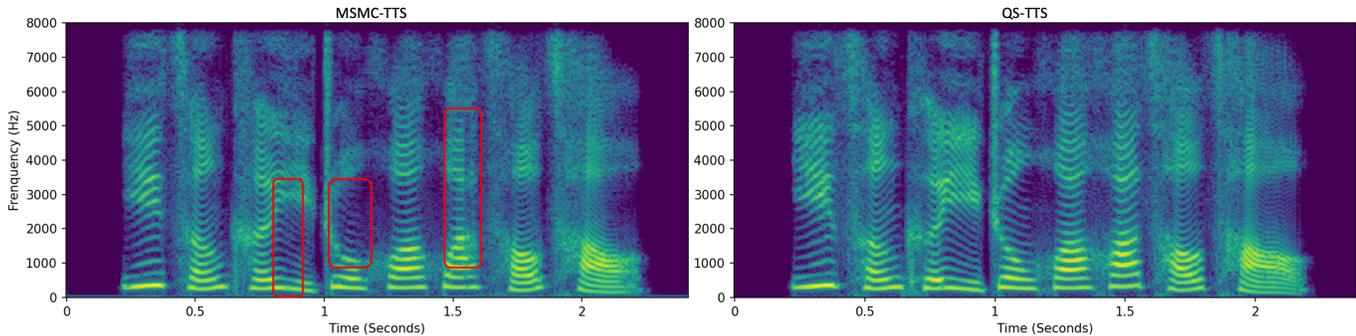


Fig. 5. The magnitude spectrograms of audio samples synthesized from the same input text by MSMC-TTS (left) and QS-TTS (right) trained with 15-minute CSMSC. Red boxes highlight areas with significant differences.

synthesis quality among all TTS approaches, obtaining the lowest MOS of 3.55 and the highest FD-AC and FD-SV of 0.79 and 1.19. But it has a satisfying performance in pronunciation accuracy, showing the lowest PER of 0.70%. VITS, the end-to-end approach based on VAE, significantly improves audio quality by a much lower FD-AC and FD-SV, and shows a higher MOS of 3.74. However, intelligibility also seriously degrades with the increase in CER and PER. Instead, MSMC-TTS performs well in both intelligibility and synthesis quality, obtaining much-decreased metrics over FastSpeech, and a higher MOS of 3.99, which indicates the effectiveness of this approach based on VQ speech representations. QS-TTS inherits this approach, and is enhanced via the proposed VQ-S3RL using more unlabeled speech data. It further enhances TTS but with limited improvement since sufficient supervised data is given in training.

TABLE I
MOS TEST: STANDARD SINGLE-SPEAKER MANDARIN TTS

Systems	FDs			ERs (%)		MOS (\pm 95%CI)
	AC	SV	ASR	CER	PER	
Recording	-	-	-	4.70	0.67	4.33 \pm 0.10
FastSpeech	0.79	1.19	1.52	6.28	0.70	3.55 \pm 0.09
VITS	0.22	0.96	1.85	7.96	1.15	3.74 \pm 0.09
MSMC-TTS	0.34	0.84	1.51	6.11	0.80	3.99 \pm 0.09
QS-TTS	0.19	0.81	1.47	6.05	0.80	4.00\pm0.09

TABLE II
MOS TEST: LOW-RESOURCE SINGLE-SPEAKER MANDARIN TTS

Systems	FDs			ERs (%)		MOS (\pm 95%CI)
	AC	SV	ASR	CER	PER	
Recording	-	-	-	4.70	0.67	4.51 \pm 0.10
FastSpeech	4.17	1.79	4.10	12.28	3.23	2.58 \pm 0.11
VITS	0.92	1.52	3.19	19.56	6.31	2.86 \pm 0.11
MSMC-TTS	0.54	1.40	4.44	16.80	4.99	3.15 \pm 0.11
QS-TTS	0.25	1.00	2.20	8.98	1.46	3.75\pm0.10

Then, we conduct this comparison in a more challenging TTS task, a low-resource scenario with only 15 minutes of supervised data. We randomly select 257 utterances from the full training set of CSMSC to form a 15-minute training set. As shown in Table II, all supervised approaches degrade seriously in audio quality and intelligibility under this scenario, achieving very low subjective scores. VITS and MSMC-TTS still achieve higher audio quality than FastSpeech, showing

lower FD-AC and FD-SV. However, due to insufficient data to learn effective speech representations, their performance on intelligibility has more severe degradation, where the end-to-end approach, VITS, shows the highest CER and PER. QS-TTS addresses this problem by conducting VQ-S3RL on more unlabeled speech data, significantly improving audio quality and intelligibility. It shows significantly improved audio quality by an FD-AC of 0.25 and an FD-SV of 1.00, even lower than FastSpeech trained with 10 hours of supervised data. The intelligibility is also enhanced with the CER of 8.98 and PER of 1.46. Finally, its MOS of 3.75 further validates that QS-TTS surpasses all supervised methods greatly in this low-resource scenario. It demonstrates that unlabeled speech data plays a more critical role in low-resource TTS.

As shown in Fig. 5, we also visualize the magnitude spectrograms of samples synthesized by the MSMC-TTS and QS-TTS to investigate their differences further. The high-quality audio usually presents clear and smooth harmonics in the middle- and low-frequency parts to produce accurate pronunciation perceptually. However, MSMC-TTS trained with only 15 minutes of supervised data cannot synthesize the expected harmonics stably. It often presents a jittered pitch, and fuzzy middle-frequency harmonics, as shown in the red boxes, which leads to a degradation of audio quality and intelligibility. In QS-TTS, these issues are alleviated significantly. The synthesized audio shows clearer smoother harmonics in low- and middle-frequency parts, and has a higher variance in the high-frequency part, finally leading to a higher MOS in the subjective test.

B. TTS Comparison: Semi-Supervised Approaches

In the experiment, we aim to demonstrate the effectiveness of the proposed semi-supervised approach for TTS, by comparing it with semi-supervised TTS systems, FastSpeech-SS and VITS-SS. First, we compare different semi-supervised TTS systems under an intra-lingual cross-style low-resource scenario using the 10-minute child speech dataset. It shares the same languages as AIShell-3, but has a unique timbre of a five-year-old girl unseen in AIShell-3. As shown in Table III, in this more challenging task with fewer supervised data, all approaches show higher Fréchet distances and error rates. First, the SOTA-supervised approach, MSMC-TTS, shows the worst intelligibility with the high CER and

PER of 27.53% and 13.02%, and also performs poorly on audio quality with the highest FD-AC of 1.76. The semi-supervised version of FastSpeech, FastSpeech-SS, is significantly enhanced over its fully-supervised version, achieving comparable performance to MSMC-TTS. And VITS-SS enhanced by semi-supervised training outperforms MSMC-TTS significantly, achieving much lower Fr chet distances and error rates. Finally, although QS-TTS has a higher FD-AC than VITS-SS, it obtains the best performance on all other metrics, especially with a CER of 14.61% of a PER of 5.35% which are twice as low as those of MSMC-TTS. It indicates the universal effectiveness of QS-TTS in extreme scenarios.

Then, we further validate its effectiveness in a cross-lingual scenario, which builds a TTS system for a low-resource language, Cantonese, using a 15-minute supervised dataset, where the language is also unseen in the unlabeled speech dataset. As shown in Table IV, this dataset has a relatively lower audio fidelity and lower expressiveness on prosody, hence achieving the MOS of 3.86 only. FastSpeech-SS shows a significant improvement over MSMC-TTS in this task, obtaining a higher MOS of 3.58, and lower metrics in audio quality and intelligibility. But VITS-SS fails to keep a stable performance, showing seriously degraded intelligibility with much higher CER and PER of 42.35% and 19.92%. The pseudo-labels extracted from the Mandarin dataset using k-means cannot adapt to Cantonese well, providing a seriously biased pre-trained model. It leads to the lowest MOS of 3.34, showing the limitation of this semi-supervised approach in cross-lingual scenarios. However, QS-TTS still performs best, with the highest MOS of 3.66. It still has the lowest CER and PER, and keeps high audio quality with comparable or lower Fr chet distances than other approaches.

In conclusion, through these challenging low-resource experiments, baseline systems show either lower performance or worse generalization, but QS-TTS is validated as a more effective and stable approach in improving both audio quality and intelligibility of TTS over other supervised and semi-supervised approaches.

TABLE III
MOS TEST: DIFFERENT SEMI-SUPERVISED TTS APPROACHES ON THE 10-MINUTE CHILD TTS DATASET

Systems	FDs			ERs (%)		MOS (\pm 95%CI)
	AC	SV	ASR	CER	PER	
Recording	-	-	-	5.06	2.03	4.17 \pm 0.10
MSMC-TTS	1.76	3.37	7.49	27.53	13.02	2.93 \pm 0.13
FastSpeech-SS	1.72	3.31	7.73	25.28	15.48	2.80 \pm 0.12
VITS-SS	1.07	2.90	6.84	20.79	10.85	3.20 \pm 0.13
QS-TTS	1.48	2.81	6.32	14.61	5.35	3.23\pm0.14

C. The Principal VQ-S3R Learner

We conduct experiments for two VQ-S3R learners respectively to investigate the impact of the proposed VQ-S3RL on TTS. The principal learner combines HuBERT-based contrastive S3RL and MSMC-VQ-GAN-based generative S3RL to benefit TTS maximally. To validate the effectiveness of these two components, we implement QS-TTS-P, i.e. QS-TTS with only the principal learner, and compare it with the following

TABLE IV
MOS TEST: DIFFERENT SEMI-SUPERVISED TTS APPROACHES ON THE 15-MINUTE CANTONESE TTS DATASET

Systems	FDs			ERs (%)		MOS (\pm 95%CI)
	AC	SV	ASR	CER	PER	
Recording	-	-	-	8.53	-	3.86 \pm 0.11
MSMC-TTS	1.57	2.16	5.89	25.47	11.61	3.37 \pm 0.11
FastSpeech-SS	0.49	1.63	4.03	20.00	6.76	3.58 \pm 0.10
VITS-SS	0.57	1.67	5.20	42.35	19.92	3.34 \pm 0.11
QS-TTS	0.51	1.53	3.86	18.76	6.26	3.66\pm0.10

semi-supervised TTS systems: FastSpeech-S, FastSpeech-SS, and MSMC-TTS-SS. All of these systems use AISHell-3 for pre-training and the 15-minute supervised data of CSMSC for TTS training.

As shown in Table V, first, The Mel-spectrogram-based FastSpeech-S, which does not use HuBERT and MSMC-VQ, performs best in analysis-synthesis, since the Mel spectrogram contains sufficient acoustic information for complete speech reconstruction. However, this feature with abundant information is hard to predict by the acoustic model without sufficient supervised data for training, leading to the worst TTS performance on audio quality. In MSMC-TTS-SS, MSMC-VQ compressed the Mel spectrogram into a more compact representation with less information. This lossy compression degrades the analysis-synthesis quality, but makes the feature easier to predict from the model trained with less supervised training data. It leads to a smaller gap between the ground-truth and predicted features, making TTS synthesis closer to analysis-synthesis in audio quality. However, the information loss also degrades intelligibility, causing a trade-off between audio quality and intelligibility in TTS synthesis. Hence, it is not advisable to overly compress features to enhance TTS unless we can keep sufficient phonetic information in compression. In FastSpeech-SS, the HuBERT, a contrastive speech representation learned from massive speech audio, discards more acoustic information, showing a lower audio quality in analysis-synthesis, but keeps richer phonetic information, which has a lower PER of 0.58% than that of the Mel spectrogram, validates its completeness in representing speech. This feature also significantly improves TTS in both audio quality and intelligibility. Finally, QS-TTS-P applies MSMC-VQ with HuBERT to learn the compact representation with rich phonetic information, and achieves the best performance in TTS, strongly verifying the effectiveness of the principal VQ-S3R learner in TTS.

D. The Associate VQ-S3R Learner

The associate VQ-S3R learner aims to provide a practical pre-trained model to enhance the acoustic model to predict the MSMCR better. To validate its effectiveness on TTS, as shown in Table. VI, we compare it with another pre-training approach, back translation, under low-resource scenarios using 15 minutes of supervised data in Mandarin or Cantonese.

First, the baseline approach, ASR-based back-translation, does not achieve consistent performance in these two scenarios. In Mandarin TTS, the ASR enhanced by Mandarin-HuBERT and trained with the Mandarin TTS dataset can tran-

TABLE V
THE PERFORMANCE OF SEMI-SUPERVISED APPROACHES BASED ON DIFFERENT SPEECH REPRESENTATIONS IN ANALYSIS-SYNTHESIS AND TTS SYNTHESIS USING THE 15-MINUTE MANDARIN TTS DATASET.

TTS System	Speech Representations		Analysis-Synthesis						TTS Synthesis				
	HuBERT	MSMC-VQ	MCD (dB)	FDs			ERs (%)		FDs			ERs (%)	
				AC	SR	ASR	CER	PER	AC	SR	ASR	CER	PER
FastSpeech-S	×	×	2.23	0.12	0.17	0.15	4.61	0.66	2.29	1.51	3.52	10.54	2.47
MSMC-TTS-SS	×	✓	3.43	0.23	0.66	0.52	5.61	0.99	0.74	1.95	3.07	11.98	2.83
FastSpeech-SS	✓	×	3.09	0.12	0.45	0.30	4.79	0.58	0.43	1.10	2.73	10.87	2.19
QS-TTS-P	✓	✓	3.32	0.15	0.59	0.41	5.11	0.79	0.34	1.03	2.34	9.31	1.72

TABLE VI
THE TTS PERFORMANCE OF ACOUSTIC MODELS BASED ON DIFFERENT PRE-TRAINING METHODS IN MANDARIN AND CANTONESE TTS WITH 15-MINUTE SUPERVISED TRAINING DATA. “-” DENOTES THAT NO PRE-TRAINING IS APPLIED TO THE ACOUSTIC MODEL.

Methods	Mandarin					Cantonese				
	FDs			ERs (%)		FDs			ERs (%)	
	AC	SR	ASR	CER	PER	AC	SR	ASR	CER	PER
-	0.34	1.03	2.34	9.31	1.72	0.58	1.62	4.07	20.39	7.23
Back Translation (ASR)	0.25	0.99	2.11	7.58	1.09	0.62	1.58	4.00	21.63	7.89
Back Translation (k-means)	0.25	1.03	2.37	10.77	2.30	0.52	1.55	4.10	23.97	9.11
Proposed	0.25	1.00	2.20	8.98	1.46	0.51	1.53	3.86	18.76	6.26

scribe the Mandarin unlabeled speech dataset well, providing a good pre-training paired dataset with sufficient transcription precision to support model pre-training, achieving the best TTS quality among all methods. However, in Cantonese TTS, the ASR trained with the Cantonese TTS dataset cannot transcribe Mandarin speech audio well into Cantonese phonemes, leading to the low-quality paired dataset. The pre-trained acoustic model on this dataset contaminates the following finetuning with Cantonese supervised data instead, showing higher CER and PER. It indicates the limitation of the ASR-based back-translation under the cross-lingual application. Although the k-means-based back-translation avoids ASR training, and shows consistent performance in both scenarios, it only improves the audio quality while degrading intelligibility. Unlike these back-translation-based approaches, the associate VQ-S3R learner can enhance TTS in audio quality and intelligibility under intra-lingual and cross-lingual scenarios, hence validated as an effective and general pre-training approach.

with a large gap to highlight the difference between them. The results in these two scenarios show opposite conclusions, that the Mandarin TTS system sharing the same language as the pre-training set prefers a smaller codebook size, while the Cantonese TTS system in a different language from the pre-training set prefers a larger codebook size. The highly-compact VQ sequence can abstract phonetic information in the Mandarin training set well, but also lacks generalization to represent cross-lingual speech. Hence, the larger codebook size benefits Cantonese TTS instead. In practice, we suggest training multiple learners with different codebook sizes, and selecting the suitable learner in TTS training regarding the difference in language between the pre-training set and the supervised dataset.

E. Requirements for Supervised Data

Finally, we investigate the supervised data requirements of QS-TTS by comparing MSMC-TTS and QS-TTS trained with different sizes of CSMSC, and drawing the line charts of FD-ASR and CER, as shown in Fig. 7. First, the fully-supervised MSMC-TTS and semi-supervised QS-TTS achieve similar and good TTS performance with sufficient supervised data. Under the situation with 10-hour supervised data, the gap between these two systems is just 0.04 on FD-ASR and 0.06% on CER. However, as data size decreases, the synthesis quality of MSMC-TTS degrades significantly, showing the rapidly increased FD-ASR and CER. In the low-resource scenario with only 15 minutes of supervised data, it achieves the FD-ASR of 4.44 and CER of 16.80%, which are 193% and 175% higher than those in the system trained with 10-hour supervised data. Instead, QS-TTS shows a much-lower audio quality and intelligibility decay as supervised data size decreases. The QS-TTS trained with 15-minute supervised data achieves the FD-ASR of 2.2 and CER of 8.98%, which are only 50% and 48% higher than QS-TTS trained with 10-hour supervised data. And the gap between these two approaches is also widened to 2.22 on FD-ASR and 7.82% on CER, which is nearly a hundred

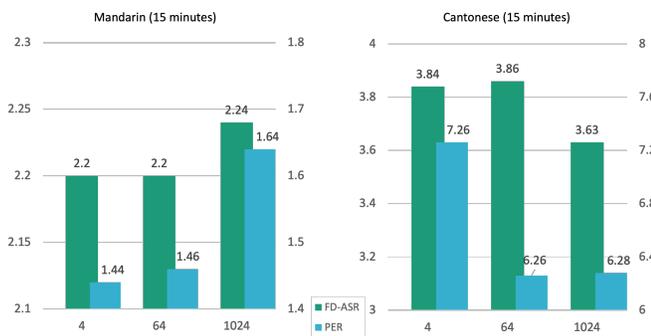


Fig. 6. The impact of different codebook sizes to associate VQ-S3R learner in Mandarin and Cantonese low-resource TTS

We also investigate the impact of the codebook size of the associate VQ-S3R learner on TTS. We train VQ-VAE models with different codebook sizes, and then apply them for acoustic model training in Mandarin and Cantonese. Fig. 6 shows the results based on 4, 64, and 1024 codewords

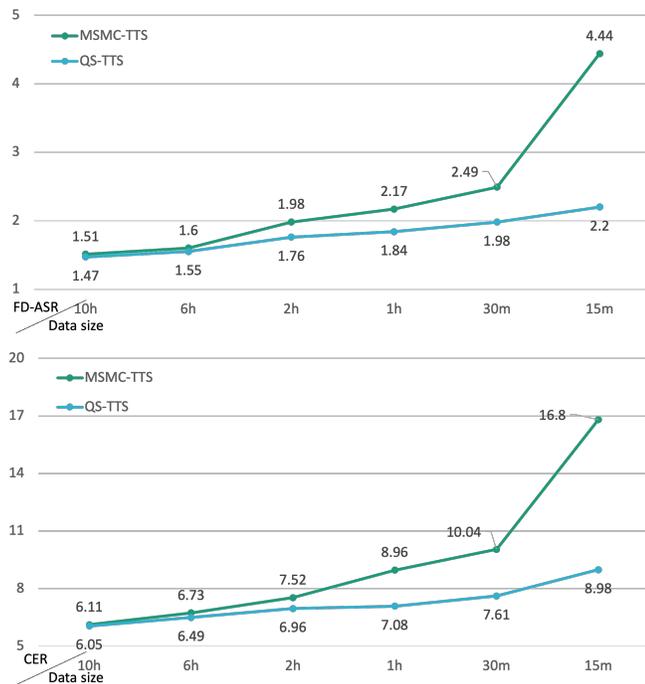


Fig. 7. The FD-ASRs and PERs of MSMC-TTS and QS-TTS trained with different data sizes of CSMSC.

times larger. Hence, under the low-resource scenario, MSMC-TTS has produced serious pronunciation issues, but QS-TTS can still keep good performance. This result strongly validates that the proposed semi-supervised TTS framework, QS-TTS, has a lower requirement for supervised data, and indicates its great potential in low-resource scenarios.

VI. CONCLUSION

This paper proposes QS-TTS, a novel semi-supervised TTS framework based on VQ-S3RL to effectively utilize more unlabeled speech audio to improve TTS quality while reducing its requirements for supervised data. The VQ-S3RL is conducted through two learners: The principal learner combines Multi-Stage Multi-Codebook (MSMC) VQ-GAN with contrastive S3RL to learn high-quality generative MSMC VQ-S3R, while decoding it to the high-quality audio; the associate learner further compresses the MSMCR into a highly-compact VQ representation via a VQ-VAE-based model. Then, TTS is implemented based on the MSMCR, and applied with VQ-S3R learners via transfer learning to achieve higher synthesis quality with lower supervised data requirements. This proposed framework can synthesize high-quality speech with lower supervised data requirements, significantly outperforming mainstream supervised and semi-supervised TTS approaches, especially in low-resource scenarios. Besides, the proposed VQ-S3RL also shows its effectiveness in providing better speech representations and pre-trained models for TTS by comparing with TTS systems with different speech representations and transfer learning methods. Finally, the slowly decayed performance of QS-TTS as supervised data decreases further validates its lower requirement for supervised data, and indicates its great potential in low-resource scenarios.

REFERENCES

- [1] Lehlohonolo Mohasi and Daniel Mashao. Text-to-speech technology in human-computer interaction. In *5th conference on human computer interaction in Southern Africa, South Africa (CHISA 2006, ACM SIGHI)*, pages 79–84, 2006.
- [2] Haohan Guo, Shaofei Zhang, Frank K Soong, Lei He, and Lei Xie. Conversational end-to-end TTS for voice agents. In *Proc. SLT*, pages 403–409. IEEE, 2021.
- [3] Danni Liu, Changhan Wang, Hongyu Gong, Xutai Ma, Yun Tang, and Juan Pino. Incremental speech synthesis for speech-to-speech translation. *arXiv preprint arXiv:2110.08214*, 2021.
- [4] Katsuhito Sudoh, Takatomo Kano, Sashi Novitasari, Tomoya Yanagita, Sakriani Sakti, and Satoshi Nakamura. Simultaneous speech-to-speech translation system with neural incremental ASR, MT, and TTS. *arXiv preprint arXiv:2011.04845*, 2020.
- [5] Chenshuang Zhang, Chaoning Zhang, Sheng Zheng, Mengchun Zhang, Maryam Qamar, Sung-Ho Bae, and In So Kweon. A survey on audio diffusion models: Text to speech synthesis and enhancement in generative ai. *arXiv preprint arXiv:2303.13336*, 2, 2023.
- [6] Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharonov, Kushal Lakhotia, Wei-Ning Hsu, Abdelrahman Mohamed, and Emmanuel Dupoux. Speech resynthesis from discrete disentangled self-supervised representations. *Proc. INTERSPEECH*, pages 3615–3619, 2021.
- [7] Haohan Guo, Fenglong Xie, Frank K Soong, Xixin Wu, and Helen Meng. A multi-stage multi-codebook vq-vae approach to high-performance neural tts. In *Proc. INTERSPEECH*, 2022.
- [8] Aäron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Proc. NeurIPS*, 2017.
- [9] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460, 2021.
- [10] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al. Tacotron: Towards end-to-end speech synthesis. In *Proc. Interspeech*, 2017.
- [11] Jonathan Shen, Ruoming Pang, Ron J Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, Rj Skerry-Ryan, et al. Natural TTS synthesis by conditioning wavenet on Mel spectrogram predictions. In *Proc. ICASSP*, pages 4779–4783. IEEE, 2018.
- [12] Yu-An Chung, Yuxuan Wang, Wei-Ning Hsu, Yu Zhang, and RJ Skerry-Ryan. Semi-supervised training for improving data efficiency in end-to-end speech synthesis. In *Proc. ICASSP*, pages 6940–6944. IEEE, 2019.
- [13] Takato Fujimoto, Shinji Takaki, Kei Hashimoto, Keiichiro Oura, Yoshihiko Nankaku, and Keiichi Tokuda. Semi-supervised learning based on hierarchical generative models for end-to-end speech synthesis. In *Proc. ICASSP*, pages 7644–7648. IEEE, 2020.
- [14] Ye Jia, Heiga Zen, Jonathan Shen, Yu Zhang, and Yonghui Wu. Png bert: Augmented bert on phonemes and graphemes for neural tts. In *Interspeech*, 2021.
- [15] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [16] Andros Tjandra, Sakriani Sakti, and Satoshi Nakamura. Listening while speaking: Speech chain by deep learning. In *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 301–308. IEEE, 2017.
- [17] Isaac Caswell, Ciprian Chelba, and David Grangier. Tagged back-translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, 2019.
- [18] Katsuki Inoue, Sunao Hara, Masanobu Abe, Tomoki Hayashi, Ryuichi Yamamoto, and Shinji Watanabe. Semi-supervised speaker adaptation for end-to-end speech synthesis with pretrained models. In *Proc. ICASSP*, pages 7634–7638. IEEE, 2020.
- [19] Eugene Kharonov, Damien Vincent, Zalán Borsos, Raphaël Marinier, Sertan Girgin, Olivier Pietquin, Matt Sharifi, Marco Tagliasacchi, and Neil Zeghidour. Speak, read and prompt: High-fidelity text-to-speech with minimal supervision. *arXiv preprint arXiv:2302.03540*, 2023.
- [20] Tao Tu, Yuan-Jui Chen, Alexander H Liu, and Hung-yi Lee. Semi-supervised learning for multi-speaker text-to-speech synthesis using discrete speech representation. *Proc. Interspeech 2020*, pages 3191–3195, 2020.

- [21] Minchan Kim, Myeonghun Jeong, Byoung Jin Choi, Sunghwan Ahn, Joun Yeop Lee, and Nam Soo Kim. Transfer learning framework for low-resource text-to-speech using a large-scale unlabeled speech corpus. In *Proc. INTERSPEECH*, pages 788–792, 2022.
- [22] Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaløe, et al. Self-supervised speech representation learning: A review. *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [23] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering*, 35(1):857–876, 2021.
- [24] Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. wav2vec: Unsupervised pre-training for speech recognition. *Proc. Interspeech 2019*, pages 3465–3469, 2019.
- [25] Alexei Baevski, Steffen Schneider, and Michael Auli. vq-wav2vec: Self-supervised learning of discrete speech representations. In *Proc. ICLR*, 2020.
- [26] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Proc. NeurIPS*, 33:12449–12460, 2020.
- [27] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518, 2022.
- [28] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020.
- [29] Tzu-hsun Feng, Annie Dong, Ching-Feng Yeh, Shu-wen Yang, Tzu-Quan Lin, Jiatong Shi, Kai-Wei Chang, Zili Huang, Haibin Wu, Xuankai Chang, et al. Superb@ slt 2022: Challenge on generalization and efficiency of self-supervised speech representation learning. In *2022 IEEE Spoken Language Technology Workshop (SLT)*, pages 1096–1103. IEEE, 2023.
- [30] Hung-yi Lee, Abdelrahman Mohamed, Shinji Watanabe, Tara Sainath, Karen Livescu, Shang-Wen Li, Shu-wen Yang, and Katrin Kirchhoff. Self-supervised representation learning for speech processing. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Tutorial Abstracts*, pages 8–13, 2022.
- [31] Songjun Cao, Yueteng Kang, Yanzhe Fu, Xiaoshuo Xu, Sining Sun, Yike Zhang, and Long Ma. Improving streaming transformer based asr under a framework of self-supervised learning. In *Proc. INTERSPEECH*, pages 706–710, 2021.
- [32] Zhengyang Chen, Sanyuan Chen, Yu Wu, Yao Qian, Chengyi Wang, Shujie Liu, Yanmin Qian, and Michael Zeng. Large-scale self-supervised speech representation learning for automatic speaker verification. In *Proc. ICASSP*, pages 6147–6151. IEEE, 2022.
- [33] Qiu-Shi Zhu, Jie Zhang, Zi-Qiang Zhang, and Li-Rong Dai. Joint training of speech enhancement and self-supervised model for noise-robust asr. *arXiv preprint arXiv:2205.13293*, 2022.
- [34] Jing Zhao and Wei-Qiang Zhang. Improving automatic speech recognition performance for low-resource languages with self-supervised models. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1227–1241, 2022.
- [35] Yohan Lim, Namhyeong Kim, Seung Yun, Sanghun Kim, and Seung-Ik Lee. A preliminary study on wav2vec 2.0 embeddings for text-to-speech. In *2021 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 343–347. IEEE, 2021.
- [36] Siyang Wang, Gustav Eje Henter, Joakim Gustafson, and Éva Székely. A comparative study of self-supervised speech representations in read and spontaneous tts. *Proc. ICASSP*, 2023.
- [37] Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in beta-vae. *arXiv preprint arXiv:1804.03599*, 2018.
- [38] Hui Lu, Zhiyong Wu, Xixin Wu, Xu Li, Shiyin Kang, Xunying Liu, and Helen Meng. Vaenar-tts: Variational auto-encoder based non-autoregressive text-to-speech synthesis. *Proc. Interspeech*, 2021.
- [39] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [40] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proc. CVPR*, pages 10696–10706, 2022.
- [41] Cristina Gârbacea, Aäron van den Oord, Yazhe Li, Felicia SC Lim, Alejandro Luebs, Oriol Vinyals, and Thomas C Walters. Low bit-rate speech coding with VQ-VAE and a WaveNet decoder. In *Proc. ICASSP*, pages 735–739. IEEE, 2019.
- [42] Yi Chen, Shan Yang, Na Hu, Lei Xie, and Dan Su. TeNC: Low bit-rate speech coding with VQ-VAE and GAN. In *ICMI*, pages 126–130, 2021.
- [43] Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. High fidelity neural audio compression. *Transactions on Machine Learning Research*, 2023.
- [44] Da-Yi Wu and Hung-yi Lee. One-shot voice conversion by vector quantization. In *Proc. ICASSP*, pages 7734–7738. IEEE, 2020.
- [45] Tomoki Hayashi and Shinji Watanabe. Discretalk: Text-to-speech as a machine translation problem. *arXiv preprint arXiv:2005.05525*, 2020.
- [46] Chengyi Wang, Sanyuan Chen, Yu Wu, Ziqiang Zhang, Long Zhou, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, et al. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*, 2023.
- [47] Chenpeng Du, Yiwei Guo, Xie Chen, and K. Yu. VQTTs: High-fidelity text-to-speech synthesis with self-supervised VQ acoustic feature. In *Proc. Interspeech*, 2022.
- [48] Herve Jegou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.
- [49] Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck. Ecapadnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification. *Proc. INTERSPEECH*, pages 3830–3834, 2020.
- [50] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Proc. NeurIPS*, pages 14866–14876, 2019.
- [51] Rewon Child. Very deep VAEs generalize autoregressive models and can outperform them on images. In *Proc. ICLR*, 2021.
- [52] Won Jang, Dan Lim, Jaesam Yoon, Bongwan Kim, and Juntae Kim. UnivNet: A neural vocoder with multi-resolution spectrogram discriminators for high-fidelity waveform generation. In *Proc. Interspeech*, 2021.
- [53] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis. In *Proc. NeurIPS*, 2020.
- [54] Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. FastSpeech: Fast, robust and controllable text to speech. In *Proc. NeurIPS*, pages 3165–3174, 2019.
- [55] Binbin Zhang, Hang Lv, Pengcheng Guo, Qijie Shao, Chao Yang, Lei Xie, Xin Xu, Hui Bu, Xiaoyu Chen, Chenchen Zeng, et al. Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition. In *Proc. ICASSP*, pages 6182–6186. IEEE, 2022.
- [56] Jackson Lee, Litong Chen, Charles Lam, Chaak Ming Lau, and Tsz-Him Tsui. PyCantonese: Cantonese linguistics and NLP in python. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6607–6611, Marseille, France, June 2022. European Language Resources Association.
- [57] Jeff Donahue, Sander Dieleman, Mikolaj Binkowski, Erich Elsen, and Karen Simonyan. End-to-end adversarial text-to-speech. In *Proc. ICLR*, 2021.
- [58] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *Proc. ICLR*, 2015.
- [59] Jaehyeon Kim, Jungil Kong, and Juhee Son. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *Proc. ICML*, pages 5530–5540. PMLR, 2021.
- [60] Jaehyeon Kim, Sungwon Kim, Jungil Kong, and Sungroh Yoon. Glow-TTS: A Generative flow for text-to-speech via monotonic alignment search. In *Proc. NeurIPS*, 2020.
- [61] Robert Kubichek. Mel-cepstral distance measure for objective speech quality assessment. In *Proceedings of IEEE pacific rim conference on communications computers and signal processing*, volume 1, pages 125–128. IEEE, 1993.
- [62] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Proc. NeurIPS*, 30, 2017.
- [63] Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi. Fr`echet audio distance: A metric for evaluating music enhancement algorithms. *Proc. INTERSPEECH*, 2019.



Haohan Guo (Student Member, IEEE) received the B.S. degree and the M.S. degree from the Northwestern Polytechnical University (NWPUP), Xi'an, Shannxi Province, China, in 2017 and 2020, respectively. He has been working toward a Ph.D. degree at the Chinese University of Hong Kong (CUHK) since August 2021. Before that, he also worked at Sogou Inc., China, as an associate researcher. His research interests include speech synthesis, voice conversion, speech and language processing, and machine learning.



Xixin Wu (Member, IEEE) received the B.S. degree from Beihang University, Beijing, China, the M.S. degree from Tsinghua University, Beijing, China, and the Ph.D. degree from The Chinese University of Hong Kong, Hong Kong. He had been a Research Associate with the Machine Intelligence Laboratory, Cambridge University Engineering Department, and from 2021 has been a Research Assistant Professor with the Stanley Ho Big Data Decision Analytics Research Centre, The Chinese University of Hong Kong. His research interests include speech synthesis and recognition, speaker verification, and neural network uncertainty. He is a Member of ISCA.



Fenglong Xie (Member, IEEE) received the B.S. and M.S. degrees in Computer Science and Technology from Harbin Institute of Technology in 2012, 2014, respectively. He received the Ph.D. degree in a joint Ph.D. program between Harbin Institute of Technology and Microsoft Research Asia in 2019. Currently He is the audio intelligence team leader of Xiaohongshu where he engages in research and product development in speech, music and audio intelligence. From 2018 to 2021, he was a senior researcher responsible for transferring speech synthesis technology from

research to several products at WeChat, Tencent. Prior to that he worked at Microsoft Research Asia from 2012 to 2018 conducting fundamental research on speech synthesis and recognition. His research interests include speech synthesis, music information retrieval, speech signal processing and speech recognition.



Jiawen Kang (IEEE Student Member) received the B.S. degree from Jilin University, Changchun, Jilin Province, China, in 2019. Since 2021, he has been pursuing his Ph.D. degree at the Chinese University of Hong Kong (CUHK). From 2019 to 2021, he worked as a research assistant at the Tsinghua University and CUHK, respectively. His research interests include spoken language processing in healthcare, speech recognition, and speaker recognition.



Helen Meng (Fellow, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA. In 1998, she joined the Chinese University of Hong Kong, Hong Kong, where she is currently a Patrick Huen Wing Ming Professor with the Department of Systems Engineering & Engineering Management. She was the former Department Chairman and the Associate Dean of Research with the faculty of Engineering. Her research interests include human-computer interaction via multimodal and multilingual spoken language systems, spoken dialog systems, computer-aided pronunciation training, speech processing in assistive technologies, health-related applications, and big data decision analytics. She was the Editor-in-Chief of the IEEE TRANSACTIONS ON AUDIO, SPEECH AND LANGUAGE PROCESSING between 2009 and 2011. She was the recipient of the IEEE Signal Processing Society Leo L. Beranek Meritorious Service Award in 2019. She was also an Elected Board Member of the International Speech Communication Association (ISCA) and an ISCA International Advisory Board Member. She is a Fellow of ISCA, HKCS, HKIE and IEEE.



Yujia Xiao (Student Member, IEEE) received the B.S. degree and the M.S. degree from the South China University of Technology (SCUT), Guangzhou, Guangdong Province, China, in 2015 and 2018, respectively. She has been working toward a Ph.D. degree at the Chinese University of Hong Kong (CUHK) since August 2022. Before that, she also worked at Microsoft China, as an applied scientist. Her research interests include speech synthesis, computer-aided language learning, speech and language processing, and machine learning.