



Single-Codec: Single-Codebook Speech Codec towards High-Performance Speech Generation

Hanzhao Li¹, Liumeng Xue², Haohan Guo³, Xinfu Zhu¹, Yuanjun Lv¹, Lei Xie^{1,*},
Yunlin Chen⁴, Hao Yin⁴, Zhifei Li⁴

¹Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, Xi'an, China

²School of Data Science, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), China

³The Chinese University of Hong Kong, Hong Kong SAR, China

⁴Shanghai Mobvoi Information Technology Co., Ltd

lihanzhao@mail.nwpu.edu.cn, lxie@nwpu.edu.cn

Abstract

The multi-codebook speech codec enables the application of large language models (LLM) in TTS but bottlenecks efficiency and robustness due to multi-sequence prediction. To avoid this obstacle, we propose Single-Codec, a single-codebook single-sequence codec, which employs a disentangled VQ-VAE to decouple speech into a time-invariant embedding and a phonetically-rich discrete sequence. Furthermore, the encoder is enhanced with 1) contextual modeling with a BLSTM module to exploit the temporal information, 2) a hybrid sampling module to alleviate distortion from upsampling and downsampling, and 3) a resampling module to encourage discrete units to carry more phonetic information. Compared with multi-codebook codecs, e.g., EnCodec and TiCodec, Single-Codec demonstrates higher reconstruction quality with a lower bandwidth of only 304bps. The effectiveness of Single-Code is further validated by LLM-TTS experiments, showing improved naturalness and intelligibility.

Index Terms: speech codec, single-codebook codec, language model, text-to-speech

1. Introduction

Large language models (LLMs) have attracted wide attention in the speech domain, particularly in text-to-speech synthesis (TTS) [1, 2]. In such LLM-based TTS systems, to operate speech synthesis as a simple next-token prediction problem, the first thing is to seek an appropriate speech codec for speech tokenization and waveform reconstruction. Multi-codebook codecs [3], as the SOTA approaches, are widely adopted in LLM-based TTS to achieve superior reconstruction quality. However, they also require the LLM to predict multiple discrete sequences, affecting efficiency and stability seriously, although various designs of codec [4, 5, 6] and LLM [7, 8, 9] are proposed to adapt this multi-sequence discrete representation better. Hence, seeking an effective approach to obtain the single-sequence discrete speech representation is critical to bypass this limitation.

However, unlike the text, it is impossible to completely represent the speech audio with abundant information in semantics and acoustics with only one discrete token sequence. Although Tortoise-TTS [10] achieves the LLM with the single-sequence discrete speech representation, A diffusion model still needs to be trained to generate Mel Spectrograms from latent

embeddings in the LLM of predicted speech units. These embeddings contain more information related to the input text and the target speaker to compensate for the compression loss. However, this operation introduces more training and inference costs. Recently, TiCodec [6] proposes introducing an additional global encoder to disentangle time-invariant information out of speech units, reducing the amount of frame-level information that needs encoding. It inspires us to re-think speech codec from the perspective of feature disentanglement.

In this study, we propose a single-codebook neural audio codec, *Single-Codec*, for high-performance speech generation. Single-Codec performs compression and reconstruction on Mel Spectrogram instead of the raw waveform, enabling efficient compression of speech information while preserving important details, as stated in Tortoise-TTS [10]. To further enhance the codec performance and applicability to speech synthesis, Single-Codec incorporates several key components:

- A global reference encoder to decouple time-invariant features. Specifically, we utilize continuous global representations rather than discrete representations and longer reference segments to capture more acoustic details, enabling embedding sufficient phonetic information into single-codebook discrete units.
- A BLSTM module for contextual modeling to help discover the correlation between adjacent frames, enhancing speech content clustering efficiency.
- A hybrid sampling module that uses both convolution and pooling to achieve downsampling, and transposed convolution and replication to achieve upsampling, alleviating upsampling and downsampling distortion.
- A resampling module to encourage the encoder to extract more phonetics-relevant information with lower short-time variance from the acoustic sequence.

To the best of our knowledge, Single-Codec is the first single-codebook codec dedicatedly designed for LLM-based speech generation. We comprehensively compare Single-Codec with SOTA multi-codebook codecs, including EnCodec and TiCodec, by conducting both objective and subjective tests in analysis-synthesis and TTS. The results show that Single-Codec with the lower bandwidth has better speech reconstruction quality, and significantly improves intelligibility, naturalness, and speaker similarity of synthesized speech in zero-shot LLM-TTS. Audio samples are available at <https://kkksuper.github.io/Single-Codec>.

* Corresponding author.

2. Method

2.1. Architecture of Single-Codec

The architecture of Single-Codec is shown in Figure 1. It is built on Vector Quantised-Variational AutoEncoder (VQVAE) [11] with Mel Spectrogram input and reconstruction, similar to Tortoise TTS [10]. We adopt a Conformer-based encoder to encode a Mel Spectrogram segment seg_2 into a latent content representation c , which is then passed to the Vector Quantizer (VQ) for vector quantization. The convolution-based decoder reconstructs the Mel Spectrogram $s\tilde{e}g_2$ from the quantized content representation c . Additionally, we apply a discriminator to improve generation quality [12]. Finally, we use a neural vocoder BigVGAN [13] to reconstruct waveform from codec output, i.e., Mel Spectrogram.

To achieve a high-quality single-codebook codec, we improve the codec architecture with four modules. Specifically, we add a reference encoder to decouple time-invariant information in speech from a Mel Spectrogram segment seg_1 , yielding a global representation g . A hybrid sampling module is adopted to alleviate sampling loss. Moreover, we introduce a BLSTM [14] module and resampling module in both codec encoder and decoder to enhance contextual information and phonetics-relevant information, respectively.

2.2. Reference Encoder

Speech contains multiple aspects of information, such as time-variant content, time-invariant timbre, and acoustic environment. Multiple codebook in codec makes it easy to encode these various information. However, for a single-codebook codec, it is challenging to compress all information into a limited number of discrete units. To solve this problem, we decouple global information (such as timbre and acoustic environment) that is almost invariable in all frames of the utterance and discretize speech content into code.

We introduce a reference encoder to derive global representation g that is mainly related to timbre and acoustic environment. The input of the reference encoder is a segment seg_1 randomly selected from the input utterance. We set the length of the segment seg_1 for reference input to 600 frames while the input segment seg_2 for codec encoder to 200 frames, where the short segment seg_2 can reduce the amount of calculation and memory overhead, while the longer reference segment seg_1 can help to obtain more robust global features. The output g of the reference encoder is fed to the codec encoder and decoder after passing through different linear layers, where it subtracts with output of the encoder blocks and adds to the input of the decoder blocks.

2.3. BLSTM Module

Codecs are generally trained on large-scale speech data to ensure good generalization. The diversity of speech content creates challenges for single-codebook codecs with appropriate sizes. Unlike EnCodec [3], which introduces the sequence modelling with LSTM [15] and finds that it can improve the Scale-Invariant Signal-to-Noise Ratio (SI-SNR) [16], we add BLSTM modules before and after the quantizer to enhance contextual information. We found this can improve the efficiency of speech content modelling and make it easier to form stable clustering centers.

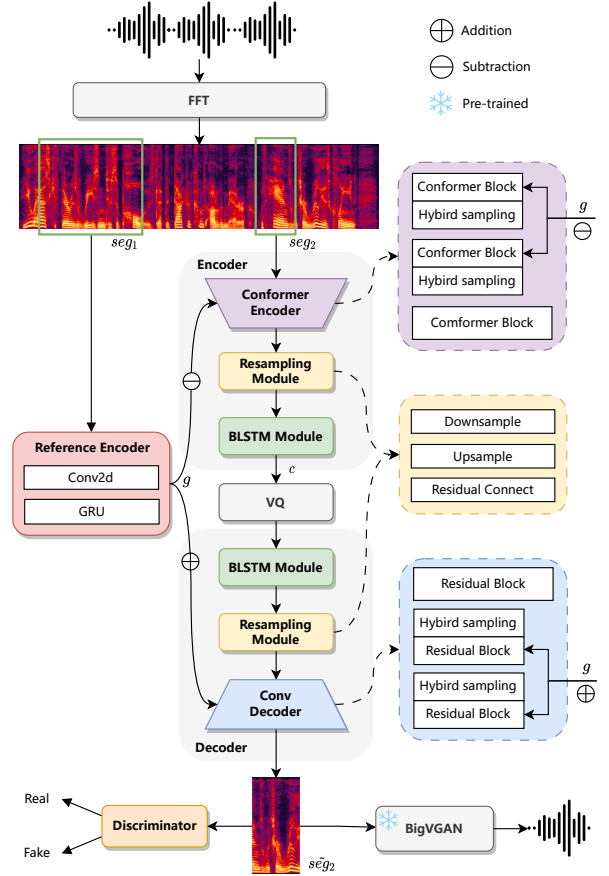


Figure 1: The architecture of Single-Codec.

2.4. Hybrid Sampling Module

Neural codecs usually employ a sampling module to reduce the sequence length of the discrete representation. Currently, the up-sampling and down-sampling operations in codecs are usually implemented by convolution, transposed convolution, or pooling and repeat. The sampling process inevitably produces sampling loss, resulting in reduced encoding and decoding capabilities. Inspired by MR-HuBERT [17], we introduce an improved hybrid sampling module that uses both convolution and pooling to achieve downsampling and transposed convolution and replication to achieve upsampling. The combination of different sampling methods can alleviate sampling distortion.

2.5. Resampling Module

The main goal of a single-codebook speech codec is to extract short-term invariant speech units from acoustic representations. The diversity of acoustic representations brings challenges to the learning of codebook vectors. To solve this problem, we propose a novel resampling module, which first downsamples the input feature for local modelling and then residual connect after upsampling. This bottlenecking operation along the time axis encourages the encoder to extract more phonetics-relevant information with lower short-time variance from the acoustic sequence.

Table 1: Objective metrics scores of various codecs, where the bold numbers highlight best results in 304bps.

Experiment	Model	Bandwidth (bps)	STOI \uparrow	PESQ \uparrow	MCD \downarrow	UTMOS \uparrow	SPK \uparrow
Ablation	VQVAE		0.805	1.591	4.489	2.847	0.727
	Ref-short		0.827	1.797	4.184	2.884	0.790
	Ref-long		0.833	1.819	4.122	2.983	0.809
	Ref-BLSTM	304	0.837	1.879	4.054	2.933	0.814
	Ref-HybSam		0.837	1.899	4.077	2.961	0.800
	Ref-BLSTM-HybSam		0.841	1.912	4.042	3.081	0.809
Comparison	Ref-BLSTM-HybSam-Conf		0.839	1.901	4.091	3.029	0.811
	EnCodec-1VQ	750	0.765	1.319	4.515	2.521	0.599
	TiCodec-1VQ	750	0.808	1.694	4.486	2.919	0.684
	TiCodec-2VQ	1500	0.866	2.140	3.865	3.079	0.762
	Single-Codec	304	0.842	1.933	4.017	3.031	0.817

3. Experiments

3.1. Dataset

We train speech codecs and a zero-shot TTS system, VALL-E [1], using five open-source datasets, including LibriTTS[18], Hi-Fi TTS[19], VCTK[20], AISHELL-1[21], and AISHELL-3[22]. A total of 1165.3 hours of English and Chinese speech is used.

3.2. Comparison Models

We adopt EnCodec [3] with one codebook (EnCodec-1VQ) and TiCodec [6] with one codebook (TiCodec-1VQ) and two codebooks (TiCodec-2VQ) as the baselines to compare with our proposed Single-Codec. For VALL-E, we use EnCodec with one, four, and eight codebooks, representing EnCodec-1VQ, EnCodec-4VQ, and EnCodec-8VQ, TiCodec with one codebook (TiCodec-1VQ) as the baselines to evaluate the performance of codecs on speech synthesis.

To verify the effectiveness of our designed modules in Single-Codec, we conduct ablation studies on the following models.

- **VQVAE**: A basic VQVAE codec with a discriminator for perceptual loss, the structure and configuration of the VQVAE codec is similar to that in Tortoise TTS [10].
- **Ref-short**: VQVAE with a reference encoder that consumes a short segment with 200 frames as input.
- **Ref-long**: VQVAE with a reference encoder that consumes a long segment with 600 frames as input.
- **Ref-BLSTM**: Ref-long with the BLSTM module to verify the effectiveness of the BLSTM module.
- **Ref-HybSam**: Ref-long with the hybrid sampling module to verify the effectiveness of the hybrid sampling module.
- **Ref-BLSTM-HybSam**: Ref-long with the BLSTM and hybrid sampling modules to verify the effectiveness of the combination of BLSTM and hybrid sampling modules.
- **Ref-BLSTM-HybSam-Conf**: Ref-BLSTM-HybSam with the Conformer-based encoder, excluding the resampling module.

3.3. Model Parameters and Training Details

The audio sample rate is 24kHz, and the hop length and window length of the Mel Spectrogram are 256 and 1024, respectively. The downsampling rate is 4, resulting in a total downsampling of 1024 times (about 23 discrete tokens per second). The codebook size is 8192. The model size of the codec is 256. The sizes of the intermediate hidden states in convolution blocks are 256, 512, and 1024, while the hidden size of the Conformer [23]

block is 1024. The reference encoder consists of 6 layers of 2D convolution with a kernel size of 3 and a GRU [24] layer. The residual block [25] consists of two residual units. Each residual unit includes 2 one-dimensional convolutions with kernel sizes of 3 and 1 respectively. Discriminator consists of 4 layers of 2D convolution with a kernel size of 5 and 2 layers of 2D convolution with a kernel size of 3. The BLSTM module contains two LSTM layers with a hidden size of 128.

During training, we conduct 300k iterations using a batch size of 1024 on a single V100 GPU for Single-Codec. The baseline model Encodec utilizes with the code reimplemented in HifiCodec¹ [26], and is trained for 25 epochs. TiCodec² is trained for 300k steps with a batch size of 40 on two V100 GPUs. For TTS, we employ VALL-E, reimplemented in Amphon³ [27], with dynamic batch sizing and a maximum token limit of 4000 per batch. The single-codebook codec only utilizes the AR stage, while the multiple-codebook codec trains both the AR and NAR stages simultaneously. Eight A800 GPUs and 70 epochs are employed for training VALL-E.

3.4. Ablation Studies

We calculate STOI [28], PESQ [29], Mel cepstral distortion (MCD) [30], UTMOS⁴ [31] and speaker cosine similarity (SPK)⁵ to objectively evaluate the quality of speech reconstruction. The test set is composed of 100 randomly selected sentences from unseen speakers. The objective result is shown in Table 1.

Compared with VQVAE, either Ref-short or Ref-long obtains better performance on all metrics. It indicates that it is effective to decouple global information from speech for the single-codebook codec. Moreover, Ref-long outperforms Ref-short in both reconstruction and speaker similarity, suggesting that longer reference segments help capture more accurate time-invariant information and enhance content modelling. Ref-BLSTM, Ref-HybSam, and Ref-BLSTM-HybSam get higher reconstruction quality, showing the effectiveness of the BLSTM and hybrid sampling modules. Moreover, Ref-BLSTM-HybSam-Con yields on-pair performance with Ref-BLSTM-HybSam but gets further improvement after adding the resampling module, i.e., our proposed Single-Code, achieving the best results.

¹<https://github.com/yangdongchao/AcademiCodec>

²<https://github.com/y-ren16/TiCodec>

³<https://github.com/open-mmlab/Amphon>

⁴<https://github.com/tarepan/SpeechMOS>

⁵We adopt WeSpeaker [32] to extract speaker embedding for cosine similarity computation

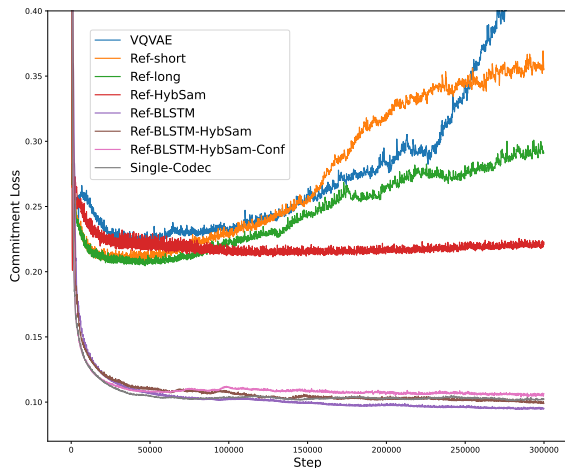


Figure 2: The commitment loss of different codec while training.

3.5. Commitment Loss Analysis

We further analyze the commitment loss in the training to explore the impact of different designed modules on the single-codebook codec. Commitment loss is the difference between representations before and after quantization. The degree of convergence of the commitment loss can reflect the relationship between the encoder output and the cluster center in the codebook. As shown in Figure 2, the commitment loss of VQVAE tends to diverge after model training, indicating that the entanglement of time-invariant global information and time-variant content information hinders forming a limited variety of content-related speech units. After considering time-invariant decoupled modelling, the loss of Ref-short increases slowly, indicating the effectiveness of global information disentanglement for speech unit learning. Ref-long further verifies this result, illustrating the effectiveness of a longer reference segment. The loss curve of Ref-HybSam is flat, indicating that the hybrid sampling module effectively improves codec performance. Moreover, the losses of the models with context modelling via the BLSTM module are all converged. It demonstrates that the models have learned stable phonetic units before quantization, indicating the effectiveness of context modelling in codecs.

Furthermore, considering the results presented in Table 1, we observe that the commitment loss is not strictly inversely related to reconstruction quality. However, the convergence status of the commitment loss (divergence, flat, convergence) is indeed associated with reconstruction quality. Specifically, the converged codec surpasses the codec which is not converged. This result further highlights the significance of achieving a stable clustering center in the single codebook codec, which directly impacts the overall reconstruction quality.

3.6. Speech Reconstruction Evaluation

We compare the performance in speech reconstruction of the proposed Single-Codec with other codecs. The results, as presented in Table 1, demonstrate that despite lower bandwidth, the proposed Single-Codec surpasses other codecs with 1 codebook and is on par with the TiCodec with 2 codebooks in terms of reconstruction quality and speaker similarity. VQVAE performs better than EnCodec with 1 codebook, demonstrating the high quantization efficiency of codecs operated on Mel Spectrogram. Compared to TiCodec, which also quantized the decoupled time-invariant information, Single-Codec achieves higher speaker similarity and reconstruction quality, indicating the effectiveness of continuous time-invariant representations and longer reference length.

Table 2: The objective and subjective metrics scores of VALL-E.

Codec Model	Subjective		Objective	
	N-MOS \uparrow	S-MOS \uparrow	WER \downarrow	S-Cosine \uparrow
VQVAE	2.81 \pm 0.06	3.24 \pm 0.07	30.9	0.649
EnCodec-1VQ	3.36 \pm 0.06	3.68 \pm 0.08	28.7	0.669
EnCodec-4VQ	3.70 \pm 0.10	3.86 \pm 0.08	15.9	0.725
EnCodec-8VQ	3.77 \pm 0.12	4.02 \pm 0.11	11.4	0.738
TiCodec-1VQ	3.68 \pm 0.08	3.97 \pm 0.07	19.4	0.683
Single-Codec	4.02 \pm 0.07	4.13 \pm 0.09	13.2	0.792

3.7. Zero-shot TTS Evaluation

To evaluate the performance of codecs applied in speech synthesis tasks, we train VALL-E [1] using discrete tokens extracted from EnCodec in the number of 1,4,8 codebooks, TiCodec in 1 codebook, and Single-Codec. We conduct naturalness Mean Opinion Score (N-MOS) and speaker similarity MOS (S-MOS) for subjective evaluation of the synthesized speech. The test set consists of 30 sentences, including Chinese and English speech. The 20 listeners who are Chinese Mandarin native speakers and familiar to English are invited to participate in each MOS test. Meanwhile, we calculate the word error rate (WER) using an ASR model⁶ [33] to measure speech intelligibility. We also use WeSpeaker [32] to extract speaker embedding to calculate the speaker embedding cosine similarity.

Table 2 shows the subjective and objective results. Single-Codec outperforms other models in terms of naturalness and speaker similarity. In single-codebook scenes, TiCodec-1VQ and Single-Codec are significantly better than other codec models in speaker similarity, naturalness, and stability. This is because decoupling global information makes the frame-level codebook pay more attention to content modelling and enables more global information transmission. Meanwhile, Single-Codec performs better than TiCodec, indicating the effectiveness of continuous global representation and additional content modelling. In addition, Single-Codec exceeds multiple-codebook codecs regarding speaker similarity and naturalness while WER is slightly higher than EnCodec-8VQ. This is mainly because the higher bandwidth brings higher-resolution speech unit perception.

4. Conclusion

In this paper, we propose Single-Codec, the first single-codebook codec dedicatedly designed for LLM-based speech generation. Single-Codec employs a disentangled VQ-VAE on Mel Spectrograms to decouple speech into time-invariant global embedding and one phonetically-rich discrete sequence quantized by one codebook. Furthermore, the encoder is enhanced with a BLSTM module for contextual modelling, a hybrid sampling module to alleviate distortion from upsampling and downsampling, and a resampling module to encourage discrete units to carry more phonetic-relevant information with lower short-time variance. In experiments, compared with multi-codebook codecs, e.g. EnCodec and TiCodec, Single-Codec demonstrates higher speech reconstruction quality with lower bandwidth of only 304bps, and enables a higher-quality LLM-TTS with better naturalness and intelligibility. In the future, we will focus on developing a more efficient single-codebook codec for speech reconstruction and speech synthesis.

⁶<https://huggingface.co/openai/whisper-large>

5. References

- [1] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Neural codec language models are zero-shot text to speech synthesizers," *CoRR*, vol. abs/2301.02111, 2023.
- [2] E. Kharitonov, D. Vincent, Z. Borsos, R. Marinier, S. Girgin, O. Pietquin, M. Sharifi, M. Tagliasacchi, and N. Zeghidour, "Speak, read and prompt: High-fidelity text-to-speech with minimal supervision," *CoRR*, vol. abs/2302.03540, 2023.
- [3] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *CoRR*, vol. abs/2210.13438, 2022.
- [4] X. Zhang, D. Zhang, S. Li, Y. Zhou, and X. Qiu, "Speechtokenizer: Unified speech tokenizer for speech large language models," *CoRR*, vol. abs/2308.16692, 2023.
- [5] S. Ji, M. Fang, Z. Jiang, R. Huang, J. Zuo, S. Wang, and Z. Zhao, "Language-codec: Reducing the gaps between discrete codec representation and speech language models," *CoRR*, vol. abs/2402.12208, 2024.
- [6] Y. Ren, T. Wang, J. Yi, L. Xu, J. Tao, C. Zhang, and J. Zhou, "Fewer-token neural speech codec with time-invariant codes," *CoRR*, vol. abs/2310.00014, 2023.
- [7] Z. Borsos, M. Sharifi, D. Vincent, E. Kharitonov, N. Zeghidour, and M. Tagliasacchi, "Soundstorm: Efficient parallel audio generation," *CoRR*, vol. abs/2305.09636, 2023.
- [8] D. Yang, J. Tian, X. Tan, R. Huang, S. Liu, X. Chang, J. Shi, S. Zhao, J. Bian, X. Wu, Z. Zhao, S. Watanabe, and H. Meng, "Uniaudio: An audio foundation model toward universal audio generation," *CoRR*, vol. abs/2310.00704, 2023.
- [9] A. Ziv, I. Gat, G. L. Lan, T. Remez, F. Kreuk, A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "Masked audio generation using a single non-autoregressive transformer," *CoRR*, vol. abs/2401.04577, 2024.
- [10] J. Betker, "Better speech synthesis through scaling," *CoRR*, vol. abs/2305.07243, 2023.
- [11] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," *ArXiv*, vol. abs/1711.00937, 2017.
- [12] P. Esser, R. Rombach, and B. Ommer, "Taming transformers for high-resolution image synthesis," in *CVPR*. Computer Vision Foundation / IEEE, 2021, pp. 12 873–12 883.
- [13] S. Lee, W. Ping, B. Ginsburg, B. Catanzaro, and S. Yoon, "Bigvgan: A universal neural vocoder with large-scale training," in *ICLR*. OpenReview.net, 2023.
- [14] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Trans. Signal Process.*, vol. 45, pp. 2673–2681, 1997.
- [15] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [16] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Trans. Speech Audio Process.*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [17] J. Shi, H. Inaguma, X. Ma, I. Kulikov, and A. Y. Sun, "Multi-resolution hubert: Multi-resolution speech self-supervised learning with masked unit prediction," *CoRR*, vol. abs/2310.02720, 2023.
- [18] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "Libritts: A corpus derived from librispeech for text-to-speech," in *INTERSPEECH*. ISCA, 2019, pp. 1526–1530.
- [19] E. Bakhturina, V. Lavrukhin, B. Ginsburg, and Y. Zhang, "Hi-fi multi-speaker english TTS dataset," in *Interspeech*. ISCA, 2021, pp. 2776–2780.
- [20] Z. Liu and B. K.-W. Mak, "Cross-lingual multi-speaker text-to-speech synthesis for voice cloning without using parallel corpus for unseen speakers," *arXiv: Audio and Speech Processing*, 2019.
- [21] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHELL-1: an open-source mandarin speech corpus and a speech recognition baseline," in *O-COCOSDA*. IEEE, 2017, pp. 1–5.
- [22] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "AISHELL-3: A multi-speaker mandarin TTS corpus and the baselines," *CoRR*, vol. abs/2010.11567, 2020.
- [23] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *INTERSPEECH*. ISCA, 2020, pp. 5036–5040.
- [24] J. Chung, Çağlar Gülçehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," *ArXiv*, vol. abs/1412.3555, 2014.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*. IEEE Computer Society, 2016, pp. 770–778.
- [26] D. Yang, S. Liu, R. Huang, J. Tian, C. Weng, and Y. Zou, "Hifi-codec: Group-residual vector quantization for high fidelity audio codec," *CoRR*, vol. abs/2305.02765, 2023.
- [27] X. Zhang, L. Xue, Y. Wang, Y. Gu, X. Chen, Z. Fang, H. Chen, L. Zou, C. Wang, J. Han, K. Chen, H. Li, and Z. Wu, "Amphion: An open-source audio, music and speech generation toolkit," *CoRR*, vol. abs/2312.09911, 2023.
- [28] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "A short-time objective intelligibility measure for time-frequency weighted noisy speech," in *ICASSP*. IEEE, 2010, pp. 4214–4217.
- [29] I.-T. Recommendation, "Perceptual evaluation of speech quality (pesq): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs," *Rec. ITU-T P. 862*, 2001.
- [30] R. F. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," *Proceedings of IEEE Pacific Rim Conference on Communications Computers and Signal Processing*, vol. 1, pp. 125–128 vol.1, 1993.
- [31] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, "UTMOS: utokyo-sarulab system for voicemos challenge 2022," in *INTERSPEECH*. ISCA, 2022, pp. 4521–4525.
- [32] H. Wang, C. Liang, S. Wang, Z. Chen, B. Zhang, X. Xiang, Y. Deng, and Y. Qian, "Wespeaker: A research and production oriented speaker embedding learning toolkit," in *ICASSP*. IEEE, 2023, pp. 1–5.
- [33] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*, ser. Proceedings of Machine Learning Research, vol. 202. PMLR, 2023, pp. 28 492–28 518.