



LSCoDec: Low-Bitrate and Speaker-Decoupled Discrete Speech Codec

Yiwei Guo^{1,2}, Zhihan Li^{1,2}, Chenpeng Du^{1,2}, Hankun Wang^{1,2}, Xie Chen^{1,2}, Kai Yu^{1,2}

¹X-LANCE Lab, MoE Kay Lab of Artificial Intelligence, School of Computer Science, Shanghai Jiao Tong University, China

²Jiangsu Key Lab of Language Computing, China

yiwei.guo@sjtu.edu.cn, kai.yu@sjtu.edu.cn

Abstract

Although discrete speech tokens have exhibited strong potential for language model-based speech generation, their high bitrates and redundant timbre information restrict the development of such models. In this work, we propose LSCoDec, a discrete speech codec that has both low bitrate and speaker decoupling ability. LSCoDec adopts a multi-stage unsupervised training framework with a speaker perturbation technique. A continuous information bottleneck is first established, followed by vector quantization that produces a discrete speaker-decoupled space. A discrete token vocoder finally refines acoustic details from LSCoDec. By reconstruction evaluations, LSCoDec demonstrates superior intelligibility and audio quality with only a **single codebook** and smaller vocabulary size than baselines. Voice conversion and speaker probing experiments prove the excellent speaker disentanglement of LSCoDec, and ablation study verifies the effectiveness of the proposed training framework.

Index Terms: Discrete speech tokens, speech codecs, low-bitrate coding, speaker decoupling, speech language modeling

1. Introduction

Currently, discrete speech tokens have laid a foundation for speech generation tasks [1–3] since they can be treated in a similar manner with the prevailing language models (LMs). By their objectives, discrete speech tokens can be divided into acoustic tokens and semantic tokens [4]. Acoustic tokens, or *speech codecs* like EnCodec [5] and DAC [6], aim to reconstruct audio perfectly with a vector quantization (VQ) module. While these tokens achieve remarkable audio decompression quality, they usually require multiple VQ layers which leads to extra complex designs in speech LMs, such as coarse-to-fine modeling [2], delay pattern [7] or nested prediction [3]. Tokens with lower bitrates are favored not only for transmission, but also for simplifying the speech LM architectures to reduce cost.

Different from other aspects of information in speech, speaker timbre is a global trait nearly orthogonal to contents and prosody. In most acoustic tokens, timbre is repeatedly encoded across timesteps. We believe this leads to a considerable waste of coding efficiency. Contrarily, semantic tokens are derived from discriminative tasks, such as self-supervised learning (SSL) [8–11], that focus more on contents other than acoustics. Thus these tokens usually have fewer codebooks than acoustic ones, resulting in lower bitrate. However, semantic tokens may discard much paralinguistic information including both time-invariant timbre and time-variant prosody [12, 13]. As prosody is an important part of emotion and semantics, this makes semantic tokens not ideal for speech generation neither.

Therefore, it is highly valuable to develop a low-bitrate and speaker-disentangled discrete speech token. The bitrate of speech tokens are determined by three factors: number of quantizers Q , frame rate F and vocabulary size V for each quantizer¹. With a smaller F , the length mismatch of speech and text modalities can be mitigated. With a smaller Q and V , speech LMs will have simpler training objectives. A small V can also make more utilization of acoustic BPEs [14, 15]. Less speaker variations also reduces the modeling difficulty of speech LMs to concentrate only on time-variant features, and speaker timbre can be rendered after LM generation.

Fortunately, lowering compression bitrate and removing speaker information benefit each other mutually. It is common to involve designs that reduce global information in speech tokens. In [16–19], a speaker encoder or global-level extractor is applied to provide timbre-like features into the decoder. The natural information bottleneck in the VQ layer helps reduce time-invariant features in the encoded tokens. Particularly, [17–19] explore the scenario of only one VQ codebook. However, no explicit decoupling technique is introduced, which does not guarantee the effectiveness of the bottleneck. Single-Codec also sacrifices stream-ability due to its non-causal encoder. Explicit speaker decoupling is explored via gradient reversal with or without supervised data [20, 21], contrastive loss [22] or swapped prediction [23], but these works do not pursue low-bitrate coding. Several other single-codebook codecs have been developed [24–26] with decent reconstruction performance and low bitrates, but without speaker decoupling.

In this work, we propose LSCoDec, the first effort to explicitly consider speaker decoupling that forms a single-codebook low-bitrate speech token unsupervisedly. LSCoDec adopts a three-stage learning process that firstly trains a speech variational autoencoder (VAE) with speaker perturbation to achieve reasonable speaker disentanglement and temporal compression. Then a VQ module is added to form a VQ-VAE where LSCoDec tokens are produced. Finally, a discrete token vocoder is individually trained on the LSCoDec tokens to refine audio quality. We use a simple stretching-based perturbation method to introduce explicit removal of timbre. We train two versions of LSCoDec under 50 and 25Hz frame rates, using only a single codebook and a small vocabulary size for 24kHz speech. For LSCoDec-50Hz, the vocabulary size is $V = 300$ with a resulting bitrate of 0.45kbps, while for LSCoDec-25Hz we choose $V = 1024$ with bitrate 0.25kbps. We show by experiments that LSCoDec owns superior reconstruction performance albeit having a single codebook and lower bitrate than the baselines. Voice conversion (VC) and speaker probing experiments further verify that LSCoDec achieves remarkable speaker disentanglement. Hence it paves the way for future speech LMs due to its

Kai Yu is the corresponding author.

¹Bitrates are computed by $Q \times F \times \lceil \log_2 V \rceil$ bps.

low bitrate, high compactness and speaker invariance.
Audio demos are available online².

2. LSCodec: Low-Bitrate Speaker-Decoupled Codec

The architecture and training diagram of LSCodec are illustrated in Fig.1. LSCodec adopts a three-stage learning process where speaker perturbation is applied on the first two stages. We elaborate on the perturbation method and three training stages in this section.

2.1. Speaker Perturbation

To explicitly disentangle discrete speech tokens and speaker information, appropriate modification on speaker timbre is necessary before tokenization, while keeping content and prosodic variations retained. Several techniques have been explored in previous study, such as formant & F0 scaling, random equalizers [22], and vocal tract length perturbation [27]. In this work, we use a simple time stretching approach that produces time-aligned perturbed speech with only pitch and timbre changes. Given a coefficient β , our perturbation process starts with a rate-based speed-up effect to scale the total duration of an utterance by β times. This operation changes pitch and formant positions, hereby altering the timbre. Then, a pitch-preserving tempo effect is applied to re-stretch the utterance to its original duration, via the WSOLA algorithm [28]. This perturbation³ only changes the global pitch position and timbre features while retaining the content and pitch variations. Therefore, to reconstruct speech, timbre must be learned from additional inputs instead of the perturbed segments. With a strong information bottleneck in the reconstruction process, speaker can be explicitly disentangled from the speech tokens. In training stage, β is independently sampled within an interval centered at 1 for every utterance to provide randomized variations.

2.2. Training Stage 1: Speech VAE

The first training stage of LSCodec is a speech VAE in a continuous space. Given an utterance waveform, we randomly cut it into a reference prompt and content segment. The content segment x is passed to the perturbation algorithm in Sec.2.1 with a random β . Then, a convolutional neural network (CNN) encoder compresses the signal in time domain and outputs isotropic Gaussian posteriors $\mathcal{N}(\mu_t, \sigma_t^2)$ for each frame t . The frame rate is controlled by the CNN strides. The sampled posteriors are fed to a Conformer [29]-based decoder. To form an information bottleneck before the decoder like [27, 30], we provide the decoder with sufficient timbre information. We use the position-agnostic cross attention [31, 32] to feed timbre from prompt segments into the decoder, as it exhibits superior timbre controllability. To provide more discriminative timbre features, we use the WavLM [11] SSL model to extract hidden embeddings from reference prompt, inspired by its widely-verified advantage on speaker verification [33, 34]. As the reference and content segments are different and the cross attention mechanism is position-agnostic, the reference prompt can almost only provide timbre information. With sufficient timbre provided in the decoder and a perturbed input, the information bottleneck formed by VAE naturally disentangles the posterior and timbre to a degree. The WavLM embeddings are fed to a CNN prompt prenet before entering the decoder.

²<https://cantabile-kwok.github.io/LSCodec/>

³The perturbation process can be efficiently implemented via SoX.

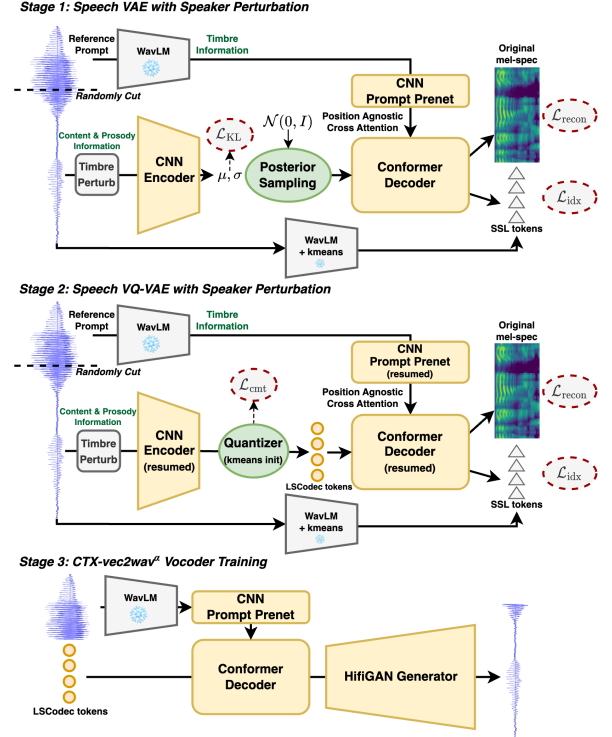


Figure 1: Architecture and Training Diagram of LSCodec.

The decoder predicts two features: mel-spectrogram and SSL semantic tokens from the original un-perturbed content segment. In other words, the VAE model performs multi-task learning that simultaneously regresses towards acoustic mel features and classifies towards SSL semantic tokens in a fixed vocabulary size. The ground truth SSL tokens are extracted using WavLM [11] and k-means clustering. The prediction of mel-spectrograms ensures reconstruction ability of the bottleneck features. As semantic tokens retain rich content information with a compact discrete space, the SSL token prediction task is crucial for guiding the bottleneck features to encode sufficient contents. Meanwhile, prosody information is largely damaged in the discrete WavLM tokens [12], so keeping the mel prediction task ensures prosody to be encoded in the bottleneck. Note that SSL models do not require labels to train, so the whole training pipeline is still unsupervised unlike FACodec [21].

Therefore in this stage, the training loss contains a KL loss $\mathcal{L}_{KL} = \frac{1}{T} \sum_{t=1}^T D_{KL}(\mathcal{N}(\mu_t, \sigma_t^2) || \mathcal{N}(\mathbf{0}, \mathbf{1}))$, an L1 loss of predicted mel features \mathcal{L}_{recon} and cross entropy loss of predicted SSL indexes \mathcal{L}_{idx} . Three losses are weighted with γ_{KL} , γ_{recon} and γ_{idx} , respectively. Due to the information bottleneck in this VAE, speaker timbre can already be partially removed in a continuous space.

2.3. Training Stage 2: Speech VQ-VAE

Next, we inject a VQ layer with codebook size V into the trained VAE to obtain the desired LSCodec tokens. We first extract the Gaussian means μ from the VAE using a portion of training data, and perform a V -centroid k-means clustering for initializing the codebook. The architecture of VQ-VAE remains almost identical to the VAE in stage 1, and we resume the encoder, decoder and prompt prenet parameters before training. The VQ layer quantizes the encoder outputs to the codebook entry with the smallest Euclidean distance in V candi-

dates. Straight-through estimator [35] is applied for gradient back-propagation. The quantized tokens are mapped to code-vectors before entering the decoder. The codebook in the VQ layer is updated via exponential moving average (EMA), following [5]. The training criterion also inherits that from stage 1, except for replacing the KL loss with the commitment loss [35] \mathcal{L}_{cmt} in the VQ layer. Note that the dimensions for previous Gaussian variances are disabled in this stage. The loss \mathcal{L}_{cmt} is weighted via γ_{cmt} .

With a VQ layer and resumed parameters, this VQ-VAE constructs a discrete space based on the insufficiently speaker-decoupled continuous space from stage 1. The information bottleneck thus further restricts timbre from being encoded. The indexes after the VQ layer are the desired LSCoDec tokens.

2.4. Training Stage 3: Vocoder CTX-vec2wav $^\alpha$

As synthesizing waveforms from predicted mel-spectrograms often yields degraded audio quality, we additionally train a specialized vocoder from the obtained discrete tokens like Vocos [36]. Since timbre is decently removed in the LSCoDec tokens, we resort to CTX-vec2wav [31] that exhibits strong timbre controllability [32]. To be consistent with the previous two stages, we develop an improved version CTX-vec2wav $^\alpha$ that receives the same WavLM outputs as timbre features instead of mel-spectrograms from reference prompts. Other parts of the vocoder follow the original CTX-vec2wav. With this vocoder, acoustic details can be refined to produce high-fidelity waveforms. This vocoder is trained from scratch. The final model can also serve as a high-quality VC model, where timbre can be well altered by providing different reference prompts.

3. Experimental Results

3.1. Data and Model Setup

We use LibriTTS [37] for model training, which is an English corpus with 585 hours of 24kHz speech data. To cover enough timbre variations, we use all of the train splits that contain about 2500 speakers. We discard all utterances shorter than 6s, remaining around 360 hours of training data. The prompts are randomly cut with lengths sampled between one third and one half of the original duration.

We train two versions of LSCoDec in 50Hz and 25Hz frame rates. The CNN encoders contain 11 and 12 residual blocks respectively and 512 hidden dimensions. In stage 1, the output of encoder is 128 dimensional before being evenly split and projected into μ, σ^2 . In stage 2, only the first half dimensions of encoder output are used for quantizing and decoding. The code vectors is then 64-dimensional. Codebook is updated with EMA weight $\gamma = 0.99$. Code expiration is enabled only after 5000 update steps. The Conformer decoder contains 2 Conformer blocks each with 2 heads and 184 attention dimensions. The prompt prenet has four CNN blocks with scaled residual connections, where the hidden dimensions are 128, 256 and 512 before being fed to cross attentions.

We use the output of the sixth layer from a pretrained WavLM-Large model⁴ to provide timbre features. In the SSL token prediction task, we use the same pretrained WavLM and perform 2048-centroid k-means clustering on around a 83-hour subset of our training data. As this SSL feature resides in 50Hz, in the 25Hz version of LSCoDec, we repeat each token by 2 times along the temporal axis before the Conformer decoder. The 80-bin cepstral-normalized mel-spectrograms for

⁴<https://github.com/microsoft/unilm/tree/master/wavlm>

Table 1: *Reconstruction results. Bitrates are presented in kbps. Parentheses denote model versions, and “-” means not comparable because of sampling rate.*

Discrete Speech Token	Bitrate \downarrow	WER \downarrow	GPE \downarrow	SECS \uparrow	MOS \uparrow
Ground truth	-	1.20	0.00	1.000	4.71 \pm 0.06
<i>Semantic (SSL) tokens with CTX-vec2wav$^\alpha$ vocoder</i>					
wav2vec2.0 (Large)	0.83	3.24	2.92	0.947	4.45 \pm 0.08
WavLM (Large)	0.55	1.67	17.94	0.934	4.36 \pm 0.07
<i>Acoustic tokens with low bitrates at 16kHz</i>					
SemantiCodec	0.63	4.16	2.15	0.926	-
LLM-Codec	0.85	6.25	1.86	0.919	-
Stable-Codec (2 posthoc)	0.70	5.06	1.73	0.889	-
<i>Acoustic tokens with low bitrates at 24kHz</i>					
DAC (First VQ only)	0.75	25.76	8.20	0.746	unintelligible
TiCodec (1VQ)	0.75	10.44	2.24	0.857	3.56 \pm 0.06
WavTokenizer (Small, 40Hz)	0.48	7.86	1.89	0.907	4.14 \pm 0.07
LSCoDec (50Hz, V=300)	0.45	3.33	2.42	0.954	4.49\pm0.07
LSCoDec (25Hz, V=1024)	0.25	5.46	2.80	0.945	4.41 \pm 0.07

reconstruction also have a frame rate of 50Hz.

Each stage of LSCoDec is trained up to 200 epochs, with loss weights $\gamma_{\text{recon}} = \gamma_{\text{KL}} = 60$, $\gamma_{\text{idx}} = 2$, $\gamma_{\text{cmt}} = 1$. The perturbation coefficient β is sampled uniformly in [0.8, 1.2]. Other hyper-parameters and the training of CTX-vec2wav $^\alpha$ follow that in [31].

3.2. Speech Reconstruction

To evaluate the reconstruction ability of LSCoDec against other low-bitrate speech tokens, we use the testset-B in [31] that contains 500 utterances from unseen speakers in LibriTTS test-clean split. Since low-bitrate speech coding cannot fully restore every detail in the original waveform, we emphasize on comparing the preservation of key information: content, prosody, and speaker timbre. We use word error rate (WER, %) to measure intelligibility and gross pitch error (GPE, %) to measure prosody preservation. WERs are computed using NeMo-ASR⁵ with the ground truth transcriptions. GPE stands for the ratio of pitch values larger than 20% relative error from the ground truth. We also compute speaker similarity cosine distance (SECS) from Resemblyzer⁶ between reconstructed and original utterances. We conduct MOS tests as a subjective evaluation of audio quality.

The baselines include semantic SSL tokens with CTX-vec2wav $^\alpha$ vocoders, where wav2vec 2.0 [10] tokens are obtained from its inner quantizer, and WavLM tokens are 2048-centroid cluster indexes from the last layer⁷. To enable a fair comparison among the low-bitrate acoustic speech codecs, we compare those with bitrate lower than 1kbps. The baselines then include DAC [6] (first VQ only), TiCodec [17], WavTokenizer-small [24] as single-codebook baselines, and SemantiCodec [38], LLM-Codec [39], Stable-Codec [25] as multi-codebook baselines. Official checkpoints are used for these baselines. We also report the bitrate of these baselines to indicate model versions. For SemantiCodec, we use the 50Hz token rate version with semantic vocabulary size 4096. For Stable-Codec, we use 2 posthoc bottlenecks. MOS tests are conducted among tokens at 24kHz sampling rate and decent reconstruction quality.

The results are presented in Table 1, which reveal that LSCoDec-50Hz exhibits significantly better WER, SECS and subjective audio quality than other acoustic token baselines although having even a lower bitrate. Although LSCoDec has a slightly higher GPE than TiCodec and WavTokenizer, the im-

⁵https://huggingface.co/nvidia/sttLen_fastConformer_transducer_large

⁶<https://github.com/resemble-ai/Resemblyzer>

⁷Same as that used in SSL token prediction task.

Table 2: VC results. Bitrates are presented in kbps.

Discrete Speech Token	Bitrate↓	WER↓	SECS↑	P.Corr↑
<i>Semantic (SSL) tokens with CTX-vec2wav^α vocoder</i>				
wav2vec2.0 (50Hz)	0.83	4.40	0.814	0.759
WavLM (50Hz)	0.55	1.92	0.872	0.374
<i>Acoustic tokens with VC ability</i>				
FACodec (80Hz, w/ detail)	4.80	1.57	0.774	0.583
FACodec (80Hz, w/o detail)	1.60	2.01	0.822	0.558
TiCodec (75Hz, 1VQ)	0.75	9.40	0.714	0.785
TiCodec (75Hz, 2VQ)	1.50	2.62	0.642	0.886
LSCodec (50Hz, V=300)	0.45	4.04	0.852	0.697
LSCodec (25Hz, V=1024)	0.25	6.32	0.853	0.753

provement on WER is still worthwhile, since a 1% difference in GPE is hardly noticeable in human perception, but WER matters more in downstream tasks. Compared to semantic tokens, LSCodec-50Hz has comparable WER with wav2vec 2.0 tokens but better GPE and SECS, and both LSCodec versions have a much better GPE than HuBERT tokens. Note that wav2vec 2.0 has much worse speaker decoupling than LSCodec as will be shown in later experiments. The 25Hz version of LSCodec has a minor degradation than the 50Hz version, since its bitrate is only 55% of the latter. LSCodec-25Hz still has competitive WER and SECS among baselines. These findings suggest that LSCodec owns remarkable reconstruction performance using a very low bitrate.

3.3. Any-to-Any Voice Conversion

While reconstruction evaluations reflect the preservation of content and prosody information, speaker disentanglement can be evaluated by VC. In this section, we randomly assign a different target speaker for every source utterance in testset-B for VC task. We measure SECS between the target and converted utterances. Higher SECS means better timbre controllability of the tokens. In addition, we also compute the correlation coefficient of pitch contours (P.Corr) between the source and converted speech. This allows us to have a better picture of a VC model both from timbre and prosody similarity. Both SECS and P.Corr range from -1 to 1. Note that P.Corr becomes meaningless when SECS is low. We include the same SSL tokens with CTX-vec2wav^α for comparison as in Table 1. As only few codec models have the ability of VC, we include TiCodec (both 1 VQ and 4 VQ variants) and FACodec [21] to comparison. The *detail* tokens in FACodec can either be discarded or preserved, which will lead to different VC outcomes.

The results are shown in Table 2. Concretely, LSCodec has better speaker similarity in VC than other acoustic tokens, especially FACodec with supervised decoupling of speakers. We find that TiCodec has unsatisfactory performance on speaker similarity, which leads to inordinately high P.Corr since converted speech is very similar to the source one. This demonstrates the difficulty of disentangling speakers only via implicit bottleneck methods. In VC scenarios, LSCodec still exhibits lower WERs than TiCodec. Compared with semantic tokens, LSCodec owns notably higher SECS than wav2vec 2.0 tokens and P.Corr than HuBERT, achieving a better balance between the two metrics. These findings reflect that speaker timbre in LSCodec tokens have been removed well with little harm on prosody, which results in competitive VC ability.

3.4. Speaker Information Probing

We also conduct speaker probing experiments as a more straight proof of the disentanglement ability. We do so by speaker classification on LibriTTS with the Xvector [40] network as the




Figure 2: Accuracy curve in speaker probing experiments.

Table 3: Ablation study on training techniques

Ablation	Configuration	S1 (VAE)		S2 (VQ-VAE)	
		WER↓	SECS↑	WER↓	SECS↑
Original LSCodec-50Hz	With \mathcal{L}_{adv} , $\beta \in [0.8, 1.2]$	4.96	0.811	3.39	0.817
On perturbation	No Perturb	5.48	0.799	3.97	0.806
	$\beta \in [0.9, 1.1]$	4.78	0.804	3.78	0.807
	$\beta \in [0.7, 1.3]$	5.40	0.809	3.62	0.809
On SSL token prediction	No \mathcal{L}_{adv}	11.22	0.811	8.93	0.816
On multi-stage training	No stage 1 training	-	-	3.84	0.800

speaker classifier. We use code-vectors associated with each token to feed the classifier. We leave out 10% utterances from the LibriTTS training set for validation. Lower classification accuracy means less speaker information present in the quantized space. The accuracy curves on the validation set along with training epochs are visualized in Fig.2. Among acoustic tokens in Table.2, LSCodec achieves the lowest probing accuracy, demonstrating its success on speaker timbre removal. For semantic tokens, LSCodec still outperforms wav2vec2.0 tokens. Since LSCodec retains prosody information, it is reasonable that it gets higher accuracy than WavLM tokens.

3.5. Ablation Study

We conduct another ablation study to verify the effectiveness of speaker perturbation, SSL token prediction task and the multi-stage training framework. We only include the first two training stages where perturbation and SSL token prediction take place. The results are shown in Table 3. Compared with the LSCodec-50Hz with $\beta \in [0.8, 1.2]$, canceling or restricting speaker perturbation leads to worse WER and SECS in both stages. This shows that the model is sacrificing local contents for more global information to decrease the loss values. Similar patterns can be found when β ranges are enlarged, since this brings too much damage to data. Canceling the SSL token prediction loss has a significant impact on WERs in both stages. Directly training from the VQ-VAE stage also has a degraded performance in both metrics. Hence, we conclude that the proposed training framework effectively produces the desired properties of LSCodec.

4. Conclusion

We propose LSCodec, a low-bitrate and speaker-decoupled speech codec with a highly compact representation space. LSCodec constructs a multi-stage training framework with speaker perturbation. A single-codebook VQ layer is applied after a VAE that disentangles speaker in a continuous space. A token vocoder is then trained upon the quantized codes. Experiments prove that LSCodec achieves strong reconstruction and VC ability given a bitrate as low as 0.45 and 0.25kbps. Probing experiments and ablation studies validate the proposed methods. Meanwhile, stronger perturbation methods, better content preservation, and scaling ability can be explored as future work.

5. Acknowledgements

This work is supported by China NSFC Project (No. 92370206), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), and Key Research and Development Program of Jiangsu Province, China (No.BE2022059).

6. References

- [1] C. Du, Y. Guo, X. Chen, and K. Yu, "VQTTs: High-Fidelity Text-to-Speech Synthesis with Self-Supervised VQ Acoustic Feature," in *Proc. ISCA Interspeech*, 2022, pp. 1596–1600.
- [2] S. Chen, C. Wang, Y. Wu *et al.*, "Neural codec language models are zero-shot text to speech synthesizers," *IEEE/ACM Trans. ASLP*, vol. 33, pp. 705–718, 2025.
- [3] D. Yang, J. Tian, X. Tan *et al.*, "UniAudio: Towards Universal Audio Generation with Large Language Models," in *Proc. ICML*, 2024.
- [4] Y. Yang, F. Shen, C. Du, Z. Ma, K. Yu, D. Povey, and X. Chen, "Towards universal speech discrete tokens: A case study for ASR and TTS," in *Proc. IEEE ICASSP*, 2024, pp. 10 401–10 405.
- [5] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High Fidelity Neural Audio Compression," *Transactions on Machine Learning Research*, 2023.
- [6] R. Kumar, P. Seetharaman, A. Luebs, I. Kumar, and K. Kumar, "High-fidelity audio compression with improved RVQGAN," *Proc. NeurIPS*, vol. 36, 2024.
- [7] P. Peng, P.-Y. Huang, S.-W. Li, A. Mohamed, and D. Harwath, "VoiceCraft: Zero-Shot Speech Editing and Text-to-Speech in the Wild," in *Proc. ACL*, Aug. 2024, pp. 12 442–12 462.
- [8] A. Baevski, S. Schneider, and M. Auli, "vq-wav2vec: Self-Supervised Learning of Discrete Speech Representations," in *Proc. ICLR*, 2020.
- [9] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units," *IEEE/ACM Trans. ASLP*, vol. 29, pp. 3451–3460, 2021.
- [10] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations," *Proc. NeurIPS*, vol. 33, pp. 12 449–12 460, 2020.
- [11] S. Chen, C. Wang, Z. Chen *et al.*, "WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [12] P. Mousavi, L. Della Libera, J. Duret, A. Ploujnikov, C. Subakan, and M. Ravanelli, "DASB–Discrete Audio and Speech Benchmark," *arXiv preprint arXiv:2406.14294*, 2024.
- [13] J. Shi, X. Ma, H. Inaguma, A. Sun, and S. Watanabe, "MMM: Multi-Layer Multi-Residual Multi-Stream Discrete Speech Representation from Self-supervised Learning Model," in *Proc. ISCA Interspeech*, 2024, pp. 2569–2573.
- [14] F. Shen, Y. Guo, C. Du, X. Chen, and K. Yu, "Acoustic BPE for speech generation with discrete tokens," in *Proc. IEEE ICASSP*. IEEE, 2024, pp. 11 746–11 750.
- [15] Y. Guo, Z. Li, H. Wang, B. Li, C. Shao, H. Zhang, C. Du, X. Chen, S. Liu, and K. Yu, "Recent advances in discrete speech tokens: A review," *arXiv preprint arXiv:2502.06490*, 2025.
- [16] X. Jiang, X. Peng, Y. Zhang, and Y. Lu, "Disentangled feature learning for real-time neural speech coding," in *Proc. IEEE ICASSP*. IEEE, 2023, pp. 1–5.
- [17] Y. Ren, T. Wang, J. Yi, L. Xu, J. Tao, C. Y. Zhang, and J. Zhou, "Fewer-token neural speech codec with time-invariant codes," in *Proc. IEEE ICASSP*. IEEE, 2024, pp. 12 737–12 741.
- [18] H. Li, L. Xue, H. Guo, X. Zhu, Y. Lv, L. Xie, Y. Chen, H. Yin, and Z. Li, "Single-Codec: Single-Codebook Speech Codec towards High-Performance Speech Generation," in *Proc. ISCA Interspeech*, 2024, pp. 3390–3394.
- [19] Y. Zheng, W. Tu, Y. Kang, J. Chen, Y. Zhang, L. Xiao, Y. Yang, and L. Ma, "FreeCodec: A disentangled neural speech codec with fewer tokens," *arXiv preprint arXiv:2412.01053*, 2024.
- [20] Á. Martín-Cortinas, D. Sáez-Trigueros, I. Vallés-Pérez *et al.*, "Enhancing the Stability of LLM-based Speech Generation Systems through Self-Supervised Representations," *arXiv preprint arXiv:2402.03407*, 2024.
- [21] Z. Ju, Y. Wang, K. Shen *et al.*, "NaturalSpeech 3: Zero-Shot Speech Synthesis with Factorized Codec and Diffusion Models," in *Proc. ICML*, 2024.
- [22] K. Qian, Y. Zhang, H. Gao, J. Ni, C.-I. Lai, D. Cox, M. Hasegawa-Johnson, and S. Chang, "Contentvec: An improved self-supervised speech representation by disentangling speakers," in *Proc. ICML*. PMLR, 2022, pp. 18 003–18 017.
- [23] H.-J. Chang, A. H. Liu, and J. Glass, "Self-supervised fine-tuning for improved content representations by speaker-invariant clustering," in *Proc. ISCA Interspeech*, 2023, pp. 2983–2987.
- [24] S. Ji, Z. Jiang, W. Wang *et al.*, "WavTokenizer: An Efficient Acoustic Discrete Codec Tokenizer for Audio Language Modeling," in *Proc. ICLR*, 2025.
- [25] J. D. Parker, A. Smirnov, J. Pons, C. Carr, Z. Zukowski, Z. Evans, and X. Liu, "Scaling transformers for low-bitrate high-quality speech coding," in *Proc. ICLR*, 2025.
- [26] H. Wu, N. Kanda, S. E. Eskimez, and J. Li, "TS3-Codec: Transformer-based simple streaming single codec," *arXiv preprint arXiv:2411.18803*, 2024.
- [27] C. H. Chan, K. Qian, Y. Zhang, and M. Hasegawa-Johnson, "SpeechSplit2.0: Unsupervised Speech Disentanglement for Voice Conversion without Tuning Autoencoder Bottlenecks," in *Proc. IEEE ICASSP*, 2022, pp. 6332–6336.
- [28] W. Verhelst and M. Roelands, "An overlap-add technique based on waveform similarity (WSOLA) for high quality time-scale modification of speech," in *Proc. IEEE ICASSP*, vol. 2. IEEE, 1993, pp. 554–557.
- [29] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proc. ISCA Interspeech*, 2020, pp. 5036–5040.
- [30] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "AutoVC: Zero-Shot Voice Style Transfer with Only Autoencoder Loss," in *Proc. ICML*. PMLR, 2019, pp. 5210–5219.
- [31] C. Du, Y. Guo, F. Shen, Z. Liu, Z. Liang, X. Chen, S. Wang, H. Zhang, and K. Yu, "UniCATS: A Unified Context-Aware Text-to-Speech Framework with Contextual VQ-Diffusion and Vocoding," in *Proc. AAAI*, vol. 38, no. 16, 2024, pp. 17 924–17 932.
- [32] J. Li, Y. Guo, X. Chen, and K. Yu, "SEF-VC: Speaker Embedding Free Zero-Shot Voice Conversion with Cross Attention," in *Proc. IEEE ICASSP*, 2024, pp. 12 296–12 300.
- [33] J. weon Jung, W. Zhang, J. Shi *et al.*, "ESPnet-SPK: full pipeline speaker embedding toolkit with reproducible recipes, self-supervised front-ends, and off-the-shelf models," in *Proc. ISCA Interspeech*, 2024, pp. 4278–4282.
- [34] S. wen Yang, P.-H. Chi, Y.-S. Chuang *et al.*, "SUPERB: Speech Processing Universal PERFORMANCE Benchmark," in *Proc. ISCA Interspeech*, 2021, pp. 1194–1198.
- [35] A. Van Den Oord, O. Vinyals *et al.*, "Neural discrete representation learning," *Proc. NeurIPS*, vol. 30, 2017.
- [36] H. Siuzdak, "Vocos: Closing the gap between time-domain and Fourier-based neural vocoders for high-quality audio synthesis," in *Proc. ICLR*, 2024.
- [37] H. Zen, V. Dang, R. Clark *et al.*, "LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech," in *Proc. ISCA Interspeech*, 2019, pp. 1526–1530.
- [38] H. Liu, X. Xu, Y. Yuan, M. Wu, W. Wang, and M. D. Plumbley, "SemantiCodec: An ultra low bitrate semantic audio codec for general sound," *IEEE Journal of Selected Topics in Signal Processing*, vol. 18, no. 8, pp. 1448–1461, 2024.
- [39] D. Yang, H. Guo, Y. Wang, R. Huang, X. Li, X. Tan, X. Wu, and H. M. Meng, "UniAudio 1.5: Large language model-driven audio codec is a few-shot audio task learner," in *Proc. NeurIPS*, 2024.
- [40] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. IEEE ICASSP*, 2018, pp. 5329–5333.