

Say More with Less: Variable-Frame-Rate Speech Tokenization via Adaptive Clustering and Implicit Duration Coding

Rui-Chen Zheng¹, Wenrui Liu², Hui-Peng Du¹, Qinglin Zhang³, Chong Deng³, Qian Chen³,
Wen Wang³, Yang Ai^{1*}, Zhen-Hua Ling¹

¹ University of Science and Technology of China, Hefei, Anhui, China

² Zhejiang University, Hangzhou, Zhejiang, China

³ Independent Researcher
zhengruichen@mail.ustc.edu.cn

Abstract

Existing speech tokenizers typically assign a fixed number of tokens per second, regardless of the varying information density or temporal fluctuations in the speech signal. This uniform token allocation mismatches the intrinsic structure of speech, where information is distributed unevenly over time. To address this, we propose **VARSTok**, a **V**ariable-**f**rame-**R**ate **S**peech **T**okenizer that adapts token allocation based on local feature similarity. VARSTok introduces two key innovations: (1) a temporal-aware density peak clustering algorithm that adaptively segments speech into variable-length units, and (2) a novel implicit duration coding scheme that embeds both content and temporal span into a single token index, eliminating the need for auxiliary duration predictors. Extensive experiments show that VARSTok significantly outperforms strong fixed-rate baselines. Notably, it achieves superior reconstruction naturalness while using up to 23% fewer tokens than a 40 Hz fixed-frame-rate baseline. VARSTok further yields lower word error rates and improved naturalness in zero-shot text-to-speech synthesis. To the best of our knowledge, this is the first work to demonstrate that a fully dynamic, variable-frame-rate acoustic speech tokenizer can be seamlessly integrated into downstream speech language models.

Code — <https://zhengrachel.github.io/VARSTok>

Extended version — <https://arxiv.org/abs/2509.04685>

1 Introduction

Speech tokenization has become central to modern speech modeling, powering advances in neural audio codecs (Zeghidour et al. 2021; Wu et al. 2023; Ai et al. 2024; Huang, Meng, and Ko 2024), generative speech synthesis (Zhang et al. 2024; Chen et al. 2025b), and multimodal large language models (LLMs) (Du et al. 2024b; Chen et al. 2025a). By converting continuous speech signals into discrete token sequences, speech tokenization bridges the gap between raw waveform and token-based language modeling architectures, thus enabling the application of LLMs to speech data.

Existing speech tokenization methods generally fall into three types: semantic tokenizers (Baevski et al. 2020; Hsu

et al. 2021) that capture high-level linguistic content; acoustic tokenizers (Kumar et al. 2023; Xin et al. 2024; Liu et al. 2025) that preserve fine-grained signal fidelity; and hybrid approaches (Zhang et al. 2024; Ye et al. 2025a) that combine the merits of both. While recent acoustic tokenizers have shown impressive results, they typically operate at fixed frame rates (e.g., 75 Hz), uniformly allocating tokens across time, ignoring the underlying content or information density.

However, natural speech is inherently non-uniform over time (Keshishian, Norman-Haignere, and Mesgarani 2021). Segments with silence or stable vowels are acoustically redundant, while others with rapid articulatory transitions or expressive prosody carry dense information (Van Kuyk, Kleijn, and Hendriks 2017; Dieleman et al. 2021). Fixed-rate tokenizers fails to adapt to this variability, leading to inefficient token usage and poor alignment with temporal dynamics. This misalignment not only results in suboptimal compression but also hinders the ability of downstream speech language models (LMs) to learn natural prosody and rhythm. Addressing this challenge requires a paradigm shift from fixed-rate to content-aware, dynamic tokenization.

To this end, we propose VARSTok, a fully dynamic, variable-frame-rate speech tokenizer that adaptively allocates tokens based on local feature similarity. It introduces a temporal-aware clustering algorithm to segment speech into variable-length units. To support downstream language modeling without auxiliary duration predictors, we design an implicit duration coding scheme that embeds both content and duration into a single token index. In contrast to prior work (Zhang et al. 2025), VARSTok enables fully dynamic token allocation without the need for hierarchical fusion or predefined temporal resolutions. Crucially, the resulting duration-aware tokens can be seamlessly used in autoregressive speech LMs without modification.

We validate the effectiveness of VARSTok on speech reconstruction, semantic evaluation, and text-to-speech (TTS) language modeling. Experimental results show that despite operating at a lower average token rate of 30.95 Hz, VARSTok achieves better reconstruction quality and semantically richer representations than a 40 Hz fixed-rate baseline. In TTS, it achieves improved naturalness while maintaining comparable or lower word error rates (WER). To the best of our knowledge, this is the first work to demonstrate that a fully dynamic, variable-frame-rate acoustic speech

*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tokenizer can be seamlessly integrated into downstream speech LMs without requiring architectural modifications.

In summary, our contributions are:

1. We introduce VARSTok, the first fully dynamic, variable-frame-rate speech tokenizer that can be seamlessly integrated into downstream autoregressive speech LMs.
2. We propose a temporal-aware clustering algorithm and a novel implicit duration coding scheme that jointly produce efficient, duration-aware token sequences, eliminating the need for auxiliary duration predictors or hierarchical token streams.
3. We demonstrate that VARSTok achieves competitive performance across three tasks, consistently outperforming strong fixed-rate baselines while using fewer tokens.

2 Related Work

2.1 Speech Tokenization

Speech tokenization has gained increasing attention for its role in neural audio compression, generative speech modeling, and multimodal LLMs (Cui et al. 2024). Existing methods broadly fall into three categories.

Semantic tokenizers extract high-level linguistic units by applying clustering or vector quantization (VQ) (Van Den Oord, Vinyals et al. 2017) over features learned from large-scale self-supervised pretraining (Hsu et al. 2021; Baevski et al. 2020; Chen et al. 2022). While effective for speech understanding tasks (Yoon, Woo, and Kim 2024; Sharma 2022; Fang et al. 2025), they discard fine-grained acoustic details and do not support direct waveform reconstruction, limiting their utility in generative applications such as speech synthesis or editing.

Acoustic tokenizers focus on preserving waveform fidelity by mapping raw speech to discrete tokens using residual VQ (Zeghidour et al. 2021; Défossez et al. 2023; Kumar et al. 2023). Compared to semantic tokens, acoustic tokens retain richer speech details and can be directly used in generation pipelines without a separate token-to-waveform model. Recent approaches have focused on single-codebook designs (Xin et al. 2024; Zhai et al. 2025; Zeng et al. 2025) to improve token efficiency and compatibility with downstream speech LMs (Borsos et al. 2023a,b). WavTokenizer (Ji et al. 2025) achieves strong reconstruction quality at fixed frame rates of 40 Hz or 75 Hz using a single codebook. However, they still rely on fixed-rate token assignment, failing to account for the temporal variability of speech signals.

Hybrid tokenizers aim to combine the merits of both by distilling linguistic content into acoustic embeddings (Zhang et al. 2024; Ye et al. 2025a,b). While enabling high-quality synthesis with linguistic control, they often require multiple codebooks, hierarchical token streams, or auxiliary modules, increasing system complexity and reducing generality.

VARSTok builds upon the acoustic tokenizer paradigm but departs from the fixed-rate framework by operating at a variable frame rate. Instead of assigning tokens uniformly in time, it adaptively segments speech based on local feature similarity, producing a dynamic token stream that better

reflects the temporal variation of speech. This enables improved token efficiency and greater alignment with speech structure compared to conventional fixed-rate approaches.

2.2 Adaptive Compression in Speech

Recent work has explored dynamic speech compression primarily in the context of semantic representation learning. For example, SD-HuBERT (Cho et al. 2024) uses sentence-level self-distillation to induce syllabic organization. Other approaches employ learnable temporal pooling (Dieleman et al. 2021), reinforcement learning (Cuervo et al. 2022), or syllable-aligned self-supervised representations (Baade, Peng, and Harwath 2025; Cho et al. 2025) for semantic boundary discovery. However, these methods are designed for semantic abstraction and discard fine-grained acoustic information, resulting in non-reconstructable representations that are inappropriate for waveform generation, transmission, or any task requiring high-fidelity reconstruction.

High-fidelity acoustic tokenizers predominantly, on the other hand, largely adopt a fixed-frame-rate paradigm, leaving dynamic alternatives underexplored. The most related work to ours is TFC (Zhang et al. 2025), which introduces variable temporal resolution into neural speech codec built upon multi-codebook residual VQ. TFC adaptively select among three predefined frame rates (75 Hz, 37.5 Hz, 18.75 Hz) based on entropy-based information density estimation, enabling flexible bitrate control and improved reconstruction under constrained token budgets. However, TFC relies on a fixed set of temporal granularities and constructs final representations by hierarchically fusing features across coarse-to-fine scales. Although the resulting frame rate varies, each token still adheres to one of the predefined resolutions, making the system pseudo-dynamic. Moreover, TFC ignores the modeling of token duration, limiting its utility in autoregressive modeling or alignment-sensitive tasks.

In contrast, VARSTok achieves fully dynamic token allocation through temporal-aware clustering without relying on fixed downsampling schedules. The proposed implicit duration coding further enables its use for downstream speech LMs without requiring auxiliary duration predictors. While both TFC and VARSTok aim to develop variable-frame-rate speech tokenization for high-fidelity reconstruction, our method offers a more flexible and complete solution, better aligned with the needs of downstream speech LMs.

3 Proposed Methods

3.1 Overview

We propose VARSTok, a fully dynamic, variable-frame-rate speech tokenizer that adaptively allocates tokens based on local speech feature similarity. Unlike conventional fixed-rate tokenizers, VARSTok reduces token usage in redundant regions while assigning more tokens to segments with rich variation, leading to more efficient modeling and better alignment with the intrinsic fluctuations of speech.

As illustrated in Figure 1, VARSTok comprises four major components: a speech encoder, a temporal-aware density peak clustering module, a VQ module, and a speech decoder. First, the speech encoder transforms raw waveforms into

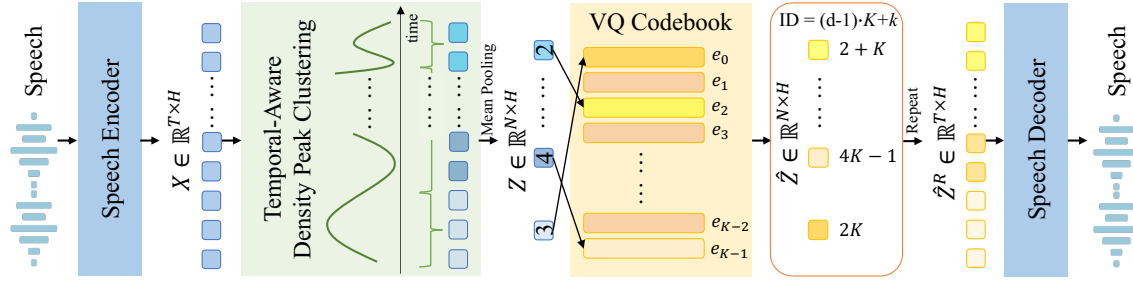


Figure 1: Overview of VARSTok. Input waveform is converted into frame-level embeddings via a speech encoder. Temporal-aware density peak clustering adaptively segments them into variable-length clusters based on similarity and temporal continuity. Each cluster is mean-pooled and quantized using a VQ codebook to produce a discrete token whose index encodes both content and duration (i.e., number of frames spanned). Each token embedding is expanded back to frame-level representations according to its duration and passed to the decoder for waveform reconstruction.

frame-level embeddings $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_T] \in \mathbb{R}^{T \times H}$, where T is the number of frames determined by the encoder’s base frame rate, and H is the embedding dimension. Adjacent frames are then adaptively grouped into N variable-length clusters $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_N\}$ based on local feature similarity, using the temporal-aware density peak clustering module. Each cluster \mathcal{C}_n is summarized into a mean-pooled embedding $\mathbf{z}_n \in \mathbb{R}^H$ and quantized using a VQ module with a single codebook $\mathcal{E} = \{e_0, \dots, e_{K-1}\}$ of size K , where e_k denotes the k -th codebook entry. The quantized embedding $\hat{\mathbf{z}}_n$ for each cluster is paired with its duration d_n , representing the number of original frames spanned by the cluster. This duration is further integrated into the token index by the proposed implicit duration coding scheme, allowing each token to jointly encode both content and temporal span, thereby preserving alignment without requiring auxiliary duration predictors. Before decoded by a speech decoder, the quantized embeddings $\hat{\mathbf{Z}} = [\hat{\mathbf{z}}_1, \dots, \hat{\mathbf{z}}_N]$ are expanded according to their durations $\mathbf{d} = [d_1, \dots, d_N]$ to form a frame-level sequence $\hat{\mathbf{Z}}^R = [\hat{\mathbf{z}}_1^R, \dots, \hat{\mathbf{z}}_T^R] \in \mathbb{R}^{T \times H}$, restoring the temporal structure for high-fidelity waveform reconstruction.

The overall design enables VARSTok to produce compact and temporally-aligned token sequences that can be directly used in downstream speech LMs, achieving high reconstruction quality with significantly improved token efficiency. We describe each component in detail in the following sections.

3.2 Model Architecture

VARSTok adopts an encoder-VQ-decoder architecture following the design of WavTokenizer (Ji et al. 2025), which achieves advanced performance with a single codebook.

Speech Encoder and Decoder The speech encoder begins with a 1D convolutional layer followed by four convolutional blocks, each containing a residual unit and a strided convolution for temporal downsampling while doubling the number of channels at each stage. The downsampled features are further processed by a two-layer bidirectional LSTM and a projection layer, producing frame-level embeddings \mathbf{X} . ELU activations are used throughout.

The speech decoder adopts an improved architecture combining attention layers, ConvNextV2 (Woo et al. 2023)

Algorithm 1: Temporal-Aware Density Peak Clustering

Require: Embeddings $\mathbf{X} \in \mathbb{R}^{T \times D}$, neighbors m , threshold τ , penalty β , max span S_{\max}
Ensure: Cluster embeddings $\mathbf{Z} = \{\mathbf{z}_1, \dots, \mathbf{z}_N\}$ and durations $\mathbf{d} = \{d_1, \dots, d_N\}$

- 1: Compute similarity $\phi(\mathbf{x}_i, \mathbf{x}_j)$ for all pairs
- 2: Compute local density ρ_i and peak distance δ_i
- 3: Compute peak score $s_i = \rho_i \cdot \delta_i$
- 4: Initialize assigned = False for all frames
- 5: **while** some frames are unassigned **do**
- 6: Select seed $i^* = \arg \max s_i$ among unassigned
- 7: Initialize cluster $\mathcal{C}_n = \{i^*\}$
- 8: **for** offset $t = \pm 1$ to $\pm S_{\max}$ **do**
- 9: Check similarity and score: $\phi(\mathbf{x}_{i^*}, \mathbf{x}_t) - \beta \cdot s_t > \tau$
- 10: If valid and unassigned, add t to \mathcal{C}_n , else break
- 11: **end for**
- 12: Sort \mathcal{C}_n ; compute mean $\mathbf{z}_n = \frac{1}{|\mathcal{C}_n|} \sum_{t \in \mathcal{C}_n} \mathbf{x}_t$, duration $d_n = |\mathcal{C}_n|$
- 13: **end while**
- 14: **return** $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_N]$, $\mathbf{d} = [d_1, \dots, d_N]$

blocks, and an inverse Fourier transform-based upsampling module. In VARSTok, the decoder receives an expanded sequence of quantized cluster embeddings $\hat{\mathbf{Z}}^R$ obtained by repeating each quantized cluster embedding $\hat{\mathbf{z}}_n$ according to their durations d_n , and finally reconstructs the original waveform.

VQ Module We adopt the VQ module from WavTokenizer, which employs a single codebook with $K = 4096$ entries. Each cluster embedding $\mathbf{z}_n \in \mathbb{R}^H$ is mapped to its nearest codebook entry $e_{k_n} \in \mathcal{E}$ based on L2 distance after factorization. The codebook is updated using exponential moving averages to encourage stability and high codebook usage. To further improve token diversity and prevent codebook collapse, random awakening is applied during training.

3.3 Temporal-Aware Density Peak Clustering

To adaptively segment the encoder output \mathbf{X} into variable-length units, we propose a temporal-aware density peak

clustering algorithm. Unlike prior clustering-based methods (Wu et al. 2025) designed for unordered inputs like image patches, speech features are inherently sequential and clustering must respect temporal continuity. Our proposed algorithm enforces strict time-order constraints to produce segments that are both semantically meaningful and temporally aligned. The full algorithm is detailed in Algorithm 1.

We first identify potential cluster centers among all the frames by calculating the local density ρ_i and peak distance δ_i of each frame. The local density ρ_i measures how closely a frame is surrounded by similar neighbors in the embedding space:

$$\rho_i = \exp\left(\frac{1}{m} \sum_{j \in \text{KNN}(i)} \phi(\mathbf{x}_i, \mathbf{x}_j)\right), \quad (1)$$

where $\text{KNN}(i)$ are the indices of the m -most-similar frames of x_i and $\phi(\cdot, \cdot)$ denotes the normalized cosine similarity between two frames:

$$\phi(\mathbf{x}_i, \mathbf{x}_j) = \frac{1 + \langle \mathbf{x}_i, \mathbf{x}_j \rangle}{2}. \quad (2)$$

Intuitively, a higher ρ_i indicates \mathbf{x}_i is in a dense region of the embedding space. The peak distance δ_i measures a frame’s separation from regions of higher density:

$$\delta_i = \begin{cases} \min_{j: \rho_j > \rho_i} 1 - \phi(\mathbf{x}_i, \mathbf{x}_j), & \text{if such } j \text{ exists,} \\ \max_j 1 - \phi(\mathbf{x}_i, \mathbf{x}_j), & \text{otherwise.} \end{cases} \quad (3)$$

A larger δ_i implies that \mathbf{x}_i is an isolated peak in the density landscape. The final peak score is defined as:

$$s_i = \rho_i \cdot \delta_i. \quad (4)$$

Frames with high s_i are both locally dense and relatively isolated, making them ideal candidates to seed new clusters.

With peak scores $\{s_i\}_{i=1}^T$ computed, we form clusters greedily. In each step, we select the unassigned frame i^* with the highest peak score s_{i^*} to initialize a new cluster $\mathcal{C}_n = \{i^*\}$. The cluster is then expanded bidirectionally from the seed i^* . A candidate frame t is added to the current cluster \mathcal{C}_n only if it meets two conditions. First it must satisfy a similarity criterion:

$$\phi(\mathbf{x}_{i^*}, \mathbf{x}_t) - \beta \cdot s_t > \tau, \quad (5)$$

where β is a penalty factor and τ is a manually defined similarity threshold. The term $-\beta \cdot s_t$ penalizes adding frames that are themselves strong cluster seeds. Second and crucially for speech, the process is temporal-aware: t -th frame considered for inclusion only if its immediate temporal neighbor (i.e., $t - 1$ for forward expansion and $t + 1$ for backward expansion) has already been assigned to the current cluster \mathcal{C}_n . This constraint ensures that all clusters are temporally contiguous segments. The expansion stops if either condition fails or if the cluster span reaches a pre-defined maximum S_{\max} . Once the expansion for cluster \mathcal{C}_n is complete, we compute its mean-pooled embedding $\mathbf{z}_n = \frac{1}{|\mathcal{C}_n|} \sum_{t \in \mathcal{C}_n} \mathbf{x}_t$ and record its span $d_n = |\mathcal{C}_n|$.

This process repeats until all frames are assigned, yielding a variable-length sequence of cluster embeddings $[\mathbf{z}_1, \dots, \mathbf{z}_N]$.

3.4 Implicit Duration Coding via Extended Index

A key challenge in variable-frame-rate speech tokenization is managing the duration of each token. Accurate duration modeling is critical for faithfully reconstructing the temporal structure of speech and ensuring proper alignment with other modalities in downstream speech LMs. A natural solution is to incorporate a separate FastSpeech-style duration predictor (Ren et al. 2019, 2020). However, this increases architectural complexity and breaks the simplicity of an end-to-end tokenizer. More importantly, we found empirically that jointly training a learned duration predictor within the tokenizer pipeline leads to severe optimization instability and poor convergence. To address this, we propose a simple yet effective implicit duration coding scheme that embeds both content identity and temporal span into a single token index. This design preserves temporal alignment without requiring auxiliary predictors, while maintaining the stability and modularity of the overall tokenizer architecture.

The core idea is to expand the VQ codebook index space by a factor of S_{\max} , the maximum allowable cluster duration, resulting in a conceptual vocabulary of $K \cdot S_{\max}$ unique token IDs. In practice, only the original single codebook \mathcal{E} of size K is instantiated and trained. Specifically, a quantized cluster embedding $\hat{\mathbf{z}}_n$ with VQ codebook index $k_n \in \{0, \dots, K - 1\}$ and duration $d_n \in \{1, \dots, S_{\max}\}$ is mapped to a single, unified token ID:

$$\text{ID}_n = (d_n - 1) \cdot K + k_n. \quad (6)$$

During the decoding stage, this process is trivially reversed to recover both the content and duration from the token ID. The duration is recovered via integer division:

$$d_n = \left\lfloor \frac{\text{ID}_n}{K} \right\rfloor + 1, \quad (7)$$

and the original VQ index is recovered via

$$k_n = \text{ID}_n \bmod K. \quad (8)$$

The corresponding quantized cluster embeddings $\hat{\mathbf{z}}_n = \mathbf{e}_{k_n}$ is then repeated d_n times to form the input for the speech decoder, restoring the correct temporal resolution for waveform reconstruction.

When applied to downstream speech LMs, the implicit duration coding scheme offers substantial advantages. By embedding both content and duration within a single token, the modeling process is significantly streamlined: the speech LM no longer requires an auxiliary duration predictor to infer temporal spans. The model operates directly on a highly compact sequence of extended token IDs without frame-level repetition or upsampling. This reformulation reduces speech modeling to a standard autoregressive prediction task, where each step selects the next token from an expanded vocabulary of size $K \times S_{\max}$. This design preserves the LM’s end-to-end simplicity while remaining lightweight, flexible, and readily scalable with different codebook sizes and maximum-allowed durations.

3.5 Training Objective

We follow the same training objective as WavTokenizer, jointly optimizing a mel-spectrogram reconstruction loss, a

Model	τ	S_{\max}	Frame Rate/Hz \downarrow	Bitrate/kbps \downarrow	UTMOS \uparrow	PESQ \uparrow	STOI \uparrow	V/UV F1 \uparrow
GT	/	/	/	/	4.1185	/	/	/
DAC	/	/	100.00	1.00	1.4940	1.2464	0.7706	0.7941
WavTokenizer	/	/	75.00	0.90	4.0247	2.4543	0.9188	0.9339
XCodec2.0	/	/	50.00	0.80	3.4727	1.8659	0.8635	0.9136
VARSTok	0.7	2	46.50	0.60	4.0379	2.0694	0.8935	0.9209
SQCodec	/	/	44.44	0.75	3.9601	1.8898	/	0.9197
BigCodec	/	/	40.00	0.52	3.9802	1.8796	0.8653	0.9133
WavTokenizer	/	/	40.00	0.48	3.6107	1.7075	0.8652	0.9095
VARSTok	0.8	4	36.81	0.52	4.0000	1.8887	0.8814	0.9186
VARSTok	0.7	4	30.95	0.43	3.8949	1.7095	0.8601	0.9047
VARSTok	0.6	4	26.29	0.37	3.8304	1.5855	0.8411	0.8985
VARSTok	0.7	8	22.38	0.34	3.6466	1.4532	0.8203	0.8860

Table 1: Speech reconstruction performance of VARSTok compared with single-codebook baselines. Results are reported on the LibriTTS test-clean set. UTMOS, PESQ, and V/UV F1 scores are computed by downsampling the speech to 16 kHz. The STOI for SQCodec is marked as ‘/’ since STOI values for the other systems are calculated at 24 kHz. The optimal results below 40Hz frame rate are marked in bold.

vector quantization loss, an adversarial loss and a feature matching loss as follows:

$$\mathcal{L} = \lambda_{mel}\mathcal{L}_{mel} + \lambda_q\mathcal{L}_q + \lambda_{adv}\mathcal{L}_{adv} + \lambda_{feat}\mathcal{L}_{feat}. \quad (9)$$

More details are provided in the Appendix B.

4 Experiments

4.1 Experimental Setup

Datasets and Tasks We trained VARSTok on the 585-hour LibriTTS corpus (Zen et al. 2019) at a 24 kHz sampling rate for fair comparison with prior work. Our experiments cover three distinct tasks for comprehensive evaluation:

- **Speech Reconstruction:** We evaluated the speech reconstruction performance of VARSTok using the LibriTTS test-clean subset. Additional results on the more challenging test-other subset are provided in the Appendix G to assess its robustness and generalizability.
- **Semantic Representation Quality:** To evaluate the semantic quality of the learned representations, we adopted the speech portion of the ARCH benchmark (La Quatra et al. 2024), including SLURP (Bastianelli et al. 2020) (intent classification), EMOVO (Costantini et al. 2014) and RAVDESS (Livingstone and Russo 2018) (emotion classification), and AudioMNIST (Becker et al. 2024) (spoken digit recognition). Additional details about the ARCH benchmark are provided in the Appendix E.
- **Downstream TTS Language Modeling:** We evaluated VARSTok’s compatibility with downstream speech LMs by training an TTS model based on a decoder-only autoregressive architecture using LibriTTS dataset. Training details are provided in the Appendix F. Its zero-shot TTS performance was assessed on the test set from UniCATS (Du et al. 2024a).

Baselines We compared VARSTok against a set of strong, single-codebook acoustic tokenizers for a fair and comprehensive evaluation. Our primary baselines are the 40 Hz and

75 Hz variants of WavTokenizer, as VARSTok is built upon its architecture. We evaluated both models using the publicly released small-version checkpoints provided by the authors. This setup allows for a direct and controlled comparison of our variable-rate mechanism against its fixed-rate counterpart across all three evaluation tasks.

To further situate our model’s reconstruction performance within the broader field, we benchmarked it against several other representative tokenizers specifically on the speech reconstruction task. These include the single-codebook DAC (Kumar et al. 2023) with results quoted from the WavTokenizer paper, alongside BigCodec (Xin et al. 2024) and XCodec2.0 (Ye et al. 2025b) which we trained from scratch on 24 kHz LibriTTS corpus. We also included results from SQCodec as optimistic references. It is important to note that SQCodec and XCodec2.0 benefit from either being trained and evaluated on 16kHz speech using substantially larger datasets or leveraging additional semantic features during inference, making their results not directly comparable. For reproducibility, all detailed configurations for baseline models are provided in the Appendix D. We do not include earlier neural codecs such as EnCodec or SoundStream as baselines, as they rely on multi-codebook or hierarchical token streams that are not directly compatible with downstream speech language modeling. For experimental efficiency, we limited semantic evaluation and TTS comparisons to VARSTok and the 40 Hz variant of WavTokenizer.

Evaluation Metrics We employed a comprehensive set of standard metrics to evaluate VARSTok’s performance across three tasks. For speech reconstruction, we report UTMOS (Saeki et al. 2022) (a neural estimator of Mean-Opinion Score), PESQ (perceptual evaluation of speech quality), STOI (short-time objective intelligibility) and V/UV F1 Score (voiced/unvoiced frame classification accuracy) to assess the naturalness, quality, intelligibility, and prosodic consistency of the reconstructed speech, respectively. For VARSTok, the reported ‘‘Frame Rate’’ is the average token rate per second calculated across the entire test set. To en-

Model	Frame Rate/Hz	EMOVO		RAVDESS		AUDIO MNIST		SLUPR	
		ACC	F1	ACC	F1	ACC	F1	ACC	F1
WavTokenizer	40.00	0.2194	0.1676	0.2847	0.2319	0.4597	0.4509	0.0658	0.0055
VARSTok($\tau = 0.8$)	36.81	0.2500	0.1900	0.2639	0.2241	0.5958	0.5930	0.0755	0.0109
VARSTok($\tau = 0.7$)	30.95	0.2347	0.1763	0.2951	0.2508	0.6111	0.6078	0.0746	0.0109
VARSTok($\tau = 0.6$)	26.29	0.2364	0.1682	0.2674	0.2348	0.6207	0.6175	0.0729	0.0098

Table 2: Classification accuracy (ACC) and macro-averaged F1 score across four datasets in the speech portion of ARCH benchmark. The optimal results are marked in bold.

sure fair comparison, we calculate the bitrate for VARSTok based on the expanded token space as:

$$\text{Bitrate} = \text{Frame Rate} \times \log_2(K \cdot S_{\max}). \quad (10)$$

This formulation accounts for the effective token vocabulary size after expansion due to the inclusion of duration. For baseline models the bitrate is calculated as:

$$\text{Bitrate} = \text{Frame Rate} \times \log_2(K). \quad (11)$$

For semantic evaluation, we use classification accuracy and macro-averaged F1 score to quantify the discriminability of learned representations across the three ARCH tasks mentioned above. For TTS language modeling, we evaluated word error rate (WER) using the Whisper-large-v3 (Radford et al. 2023) model, speaker similarity via cosine similarity of WavLM (Chen et al. 2022) embeddings, and UTMOS to assess perceptual naturalness of the synthesized speech. Subjective evaluations were also conducted. The setup is described in Appendix I. We report mean opinion scores (MOS) for naturalness and similarity MOS (SMOS) to assess perceived speaker similarity relative to the reference.

Implementation Details We initialized VARSTok by loading the encoder, decoder, and VQ codebook from the pretrained small-version 75 Hz WavTokenizer. This initialization provided a solid foundation and facilitated fast convergence when adapting to the proposed variable-frame-rate tokenization scheme. No additional losses or auxiliary predictors are used beyond those inherited from WavTokenizer. For the clustering algorithm, we set the penalty coefficient to $\beta = 0.2$ and nearest neighbors $m = 5$. To investigate the impact of clustering granularity, we varied the similarity threshold $\tau \in \{0.6, 0.7, 0.8\}$. We set $S_{\max} = 4$ by default. For the speech reconstruction task, we additionally explored different values of $S_{\max} \in \{2, 4, 8\}$ to study its influence on temporal compression. A detailed description of our training configuration is provided in Appendix C.

4.2 Speech Reconstruction

Main Results The speech reconstruction results in Table 1 confirm VARSTok’s superior efficiency over fixed-rate methods. For instance, VARSTok configured with $\tau = 0.7$, $S_{\max} = 4$ achieves a UTMOS score of 3.8949, surpassing the 40Hz WavTokenizer despite a 23% reduction in token frame rate. More strikingly, VARSTok configured with $\tau = 0.8$, $S_{\max} = 4$ achieves a UTMOS of 4.0000, nearly on par with the 75 Hz WavTokenizer while using less than

half the tokens. These highlight VARSTok’s ability to discard redundancy without sacrificing essential acoustic details. Furthermore, among all models operating below 40 Hz in Tabel 1, VARSTok consistently delivers leading performance across all metrics, underscoring its competitiveness in the low-frame-rate regime.

Analysis of Hyperparameters We conducted an analysis to investigate the impact of clustering hyperparameters on the rate-quality trade-off in Table 1. The similarity threshold τ controls the selectivity of cluster expansion from seed frames. A higher τ enforces stricter similarity, resulting in shorter clusters and thus a higher token rate.

For instance, increasing τ from 0.6 to 0.8 boosts the UTMOS score from 3.83 to 4.00, but at the cost of an increased frame rate. The maximum span S_{\max} constrains the maximum duration of each token, thereby limiting the extent of temporal compression. Larger S_{\max} values allow more aggressive clustering and thus lower token rates, but may underfit fine-grained acoustic detail. We observed that increasing S_{\max} from 2 to 8 reduces the average token rate from 46.50 Hz to 22.38 Hz, but causes severe degradation in UTMOS and PESQ. Overall, we found $\tau = 0.7$ and $S_{\max} = 4$ to strikes the best balance, achieving a token rate of 30.95 Hz with strong performance. This analysis validates VARSTok’s flexibility, offering granular control over the rate-quality spectrum via an intuitive hyperparameters, which is a capability that fixed-rate models lack.

Visualization of Token Boundaries To better understand how VARSTok allocates tokens over time, we visualize the token boundaries for two speech segments in Figure 2. As illustrated, the 75 Hz and 40 Hz WavTokenizer assigns tokens uniformly across time regardless of the acoustic structure, leading to potential over- or under-tokenization in regions of varying complexity. In contrast, VARSTok dynamically adjusts token duration in response to local feature similarity: shorter tokens are assigned to acoustically rich regions (e.g., rapid formant transitions at 0.15-0.20s in (a)), while longer tokens are assigned to redundant regions (e.g., steady vowels at 0.35-0.40s of (a) or silence at 0.35-0.50s of (b)). This adaptive allocation enables improved token efficiency while preserving critical acoustic details, as reflected in higher UTMOS and PESQ scores under comparable or lower frame rates. These qualitative findings support our design motivation: variable token durations aligns more closely with the temporal dynamics of speech and yields a more semantically faithful and resource-efficient representation.

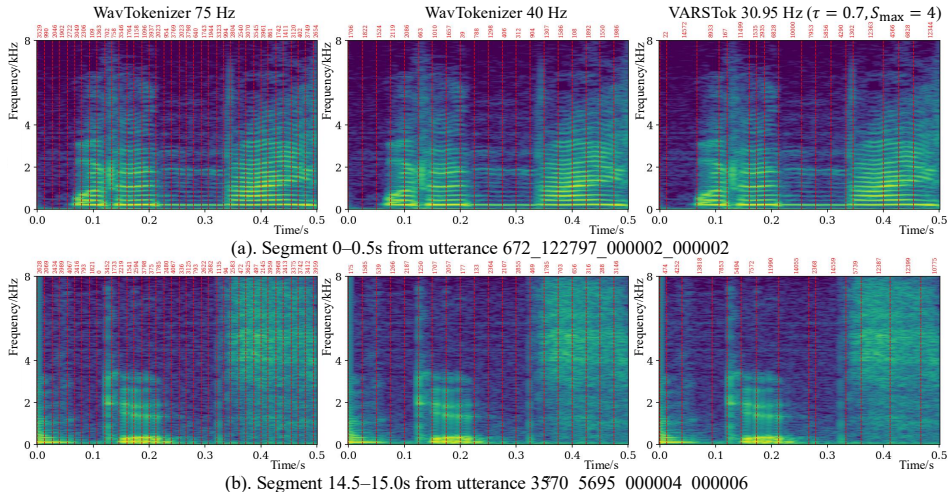


Figure 2: Token boundary visualization on two speech segments for three tokenizers : WavTokenizer (75 Hz), WavTokenizer (40 Hz), and VARSTok (30.95 Hz). Vertical red lines indicate token boundaries and token IDs are annotated above.

Tokenizer	Frame Rate/Hz↓	WER/%↓	SIM↑	UTMOS↑	MOS↑	SMOS↑
WavTokenizer	40.00	7.481	0.918	3.920	3.983 ± 0.065	3.918 ± 0.063
VARSTok($\tau = 0.8$)	36.81	6.787	<u>0.899</u>	4.246	4.053 ± 0.063	3.946 ± 0.062
VARSTok($\tau = 0.7$)	30.95	<u>7.294</u>	0.895	<u>4.199</u>	<u>4.036 ± 0.065</u>	<u>3.941 ± 0.065</u>
VARSTok($\tau = 0.6$)	26.29	9.393	0.880	4.083	3.986 ± 0.064	3.913 ± 0.066

Table 3: Comparison of WER, speaker similarity (SIM), UTMOS between fixed-rate WavTokenizer at a 40Hz frame rate and the proposed VARSTok. The optimal and suboptimal results are marked in bold and underlined, respectively.

4.3 Semantic Evaluation

We adopted ARCH benchmark to assess the semantic quality of the learned speech representations. We report classification accuracy and macro-averaged F1 score in Table 2. Despite operating at a lower average token rate, VARSTok achieves superior performance compared to the 40 Hz WavTokenizer across all tasks and all settings. These results demonstrate that dynamic token allocation leads to more semantically expressive representations.

4.4 TTS Language Modeling

To evaluate the compatibility of VARSTok with downstream speech LMs, we trained an TTS model based on a decoder-only autoregressive architecture. Similar to VALL-E (Chen et al. 2025b), given an input text and a 3-second speech prompt from the target speaker, the model autoregressively generates speech tokens, which are then directly decoded into waveform using the decoder of the pretrained speech tokenizer. As shown in Table 3, VARSTok with $\tau = 0.8$ and $S_{max} = 4$ achieves the best overall performance despite using fewer tokens on average, improving naturalness while reducing WER compared to the 40 Hz WavTokenizer baseline. Even at a lower token rate of 30.95 Hz, our model achieves comparable WER and higher UTMOS, demonstrating that VARSTok enhances both efficiency and

generation quality. We note a slight decrease in the objective speaker similarity score as the token rate drops. However, such metrics are known to sometimes diverge from human perception. Subjective results from human evaluators confirm that VARSTok matches the baseline in perceived speaker similarity and naturalness, indicating the minor drop in the objective SIM score is not perceptually significant. These results reinforces the overall superiority of our dynamic tokenization approach for enabling token-efficient and high-fidelity speech synthesis.

5 Conclusion

We introduce VARSTok, a variable-frame-rate speech tokenizer that produces compact token sequences by dynamically segmenting speech using a temporal-aware clustering algorithm. To support seamless integration with downstream speech LMs, we incorporate an implicit duration coding scheme that encodes both content and temporal span into a single token index. Comprehensive experiments confirmed that VARSTok enhances token efficiency while maintaining or surpassing the performance of fixed-rate baselines, establishing it as a principled and effective framework for dynamic speech tokenization. Future work could extend this dynamic paradigm to other audio domains such as music.

Acknowledgments

This work was partially funded by the National Nature Science Foundation of China under Grant 62301521 and the Anhui Province Major Science and Technology Research Project under Grant S2023Z20004.

References

- Ai, Y.; Jiang, X.-H.; Lu, Y.-X.; Du, H.-P.; and Ling, Z.-H. 2024. APCodec: A neural audio codec with parallel amplitude and phase spectrum encoding and decoding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32: 3256–3269.
- Baade, A.; Peng, P.; and Harwath, D. 2025. SyllableLM: Learning Coarse Semantic Units for Speech Language Models. In *Proc. ICLR 2025*.
- Baevski, A.; Zhou, Y.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. In *Proc. NeurIPS 2020*, volume 33, 12449–12460.
- Bastianelli, E.; Vanzo, A.; Swietojanski, P.; and Rieser, V. 2020. SLURP: A Spoken Language Understanding Resource Package. In *Proc. EMNLP 2020*, 7252–7262.
- Becker, S.; Vielhaben, J.; Ackermann, M.; Müller, K.-R.; Lapuschkin, S.; and Samek, W. 2024. AudioMNIST: exploring explainable artificial intelligence for audio analysis on a simple benchmark. *Journal of the Franklin Institute*, 361(1): 418–428.
- Borsos, Z.; Marinier, R.; Vincent, D.; Kharitonov, E.; Pietquin, O.; Sharifi, M.; Roblek, D.; Teboul, O.; Grangier, D.; Tagliasacchi, M.; et al. 2023a. Audioldm: a language modeling approach to audio generation. *IEEE/ACM transactions on audio, speech, and language processing*, 31: 2523–2533.
- Borsos, Z.; Sharifi, M.; Vincent, D.; Kharitonov, E.; Zeghidour, N.; and Tagliasacchi, M. 2023b. Soundstorm: Efficient parallel audio generation. *arXiv preprint arXiv:2305.09636*.
- Chen, Q.; Chen, Y.; Chen, Y.; Chen, M.; Chen, Y.; Deng, C.; Du, Z.; Gao, R.; Gao, C.; Gao, Z.; et al. 2025a. Minmo: A multimodal large language model for seamless voice interaction. *arXiv preprint arXiv:2501.06282*.
- Chen, S.; Wang, C.; Chen, Z.; Wu, Y.; Liu, S.; Chen, Z.; Li, J.; Kanda, N.; Yoshioka, T.; Xiao, X.; et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6): 1505–1518.
- Chen, S.; Wang, C.; Wu, Y.; Zhang, Z.; Zhou, L.; Liu, S.; Chen, Z.; Liu, Y.; Wang, H.; Li, J.; et al. 2025b. Neural codec language models are zero-shot text to speech synthesizers. *IEEE Transactions on Audio, Speech and Language Processing*, 33: 705–718.
- Cho, C. J.; Lee, N.; Gupta, A.; Agarwal, D.; Chen, E.; Black, A.; and Anumanchipalli, G. 2025. Sylber: Syllabic Embedding Representation of Speech from Raw Audio. In *Proc. ICLR 2025*.
- Cho, C. J.; Mohamed, A.; Li, S.-W.; Black, A. W.; and Anumanchipalli, G. K. 2024. Sd-hubert: Sentence-level self-distillation induces syllabic organization in hubert. In *Proc. ICASSP 2024*, 12076–12080. IEEE.
- Costantini, G.; Iaderola, I.; Paoloni, A.; Todisco, M.; et al. 2014. EMOVO corpus: an Italian emotional speech database. In *Proc. LREC 2014*, 3501–3504. European Language Resources Association (ELRA).
- Cuervo, S.; Lancucki, A.; Marxer, R.; Rychlikowski, P.; and Chorowski, J. K. 2022. Variable-rate hierarchical CPC leads to acoustic unit discovery in speech. In *Proc. NeurIPS 2022*, volume 35, 34995–35006.
- Cui, W.; Yu, D.; Jiao, X.; Meng, Z.; Zhang, G.; Wang, Q.; Guo, Y.; and King, I. 2024. Recent advances in speech language models: A survey. *arXiv preprint arXiv:2410.03751*.
- Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2023. High Fidelity Neural Audio Compression. *Transactions on Machine Learning Research*.
- Dieleman, S.; Nash, C.; Engel, J.; and Simonyan, K. 2021. Variable-rate discrete representation learning. *arXiv preprint arXiv:2103.06089*.
- Du, C.; Guo, Y.; Shen, F.; Liu, Z.; Liang, Z.; Chen, X.; Wang, S.; Zhang, H.; and Yu, K. 2024a. Unicats: A unified context-aware text-to-speech framework with contextual vq-diffusion and vocoding. In *Proc. AAAI 2024*, volume 38, 17924–17932.
- Du, Z.; Chen, Q.; Zhang, S.; Hu, K.; Lu, H.; Yang, Y.; Hu, H.; Zheng, S.; Gu, Y.; Ma, Z.; et al. 2024b. Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens. *arXiv preprint arXiv:2407.05407*.
- Fang, Q.; Guo, S.; Zhou, Y.; Ma, Z.; Zhang, S.; and Feng, Y. 2025. Llama-omni: Seamless speech interaction with large language models. In *Proc. ICLR 2025*.
- Hsu, W.-N.; Bolte, B.; Tsai, Y.-H. H.; Lakhotia, K.; Salakhutdinov, R.; and Mohamed, A. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM transactions on audio, speech, and language processing*, 29: 3451–3460.
- Huang, Z.; Meng, C.; and Ko, T. 2024. RepCodec: A Speech Representation Codec for Speech Tokenization. In *Proc. ACL 2024*, 5777–5790.
- Ji, S.; Jiang, Z.; Wang, W.; Chen, Y.; Fang, M.; Zuo, J.; Yang, Q.; Cheng, X.; Wang, Z.; Li, R.; et al. 2025. Wavtokenizer: an efficient acoustic discrete codec tokenizer for audio language modeling. In *Proc. ICLR 2025*.
- Keshishian, M.; Norman-Haignere, S.; and Mesgarani, N. 2021. Understanding adaptive, multiscale temporal integration in deep speech recognition systems. In *Proc. NeurIPS 2021*, volume 34, 24455–24467.
- Kumar, R.; Seetharaman, P.; Luebs, A.; Kumar, I.; and Kumar, K. 2023. High-fidelity audio compression with improved rvqgan. *Proc. NeurIPS 2024*, 36: 27980–27993.
- La Quatra, M.; Koudounas, A.; Vaiani, L.; Baralis, E.; Cagliero, L.; Garza, P.; and Siniscalchi, S. M. 2024. Benchmarking representations for speech, music, and acoustic events. In *Proc. ICASSPW 2024*, 505–509. IEEE.

- Liu, W.; Guo, Z.; Xu, J.; Lv, Y.; Chu, Y.; Liu, Z.; and Lin, J. 2025. Analyzing and Mitigating Inconsistency in Discrete Speech Tokens for Neural Codec Language Models. In *Proc. ACL 2025*, 31035–31046.
- Livingstone, S. R.; and Russo, F. A. 2018. The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in North American English. *PLoS one*, 13(5): e0196391.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2023. Robust speech recognition via large-scale weak supervision. In *Proc. ICML 2023*, 28492–28518. PMLR.
- Ren, Y.; Hu, C.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T.-Y. 2020. FastSpeech 2: Fast and High-Quality End-to-End Text to Speech. In *Proc. ICLR 2020*.
- Ren, Y.; Ruan, Y.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T.-Y. 2019. FastSpeech: Fast, robust and controllable text to speech. In *Proc. NeurIPS 2019*, volume 32.
- Saeki, T.; Xin, D.; Nakata, W.; Koriyama, T.; Takamichi, S.; and Saruwatari, H. 2022. UTMOS: UTokyo-SaruLab System for VoiceMOS Challenge 2022. In *Proc. Interspeech 2022*, volume 2022, 4521–4525.
- Sharma, M. 2022. Multi-lingual multi-task speech emotion recognition using wav2vec 2.0. In *Proc. ICASSP 2022*, 6907–6911. IEEE.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. In *Proc. NeurIPS 2017*, volume 30.
- Van Kuyk, S.; Kleijn, W. B.; and Hendriks, R. C. 2017. On the information rate of speech communication. In *Proc. ICASSP 2017*, 5625–5629. IEEE.
- Woo, S.; Debnath, S.; Hu, R.; Chen, X.; Liu, Z.; Kweon, I. S.; and Xie, S. 2023. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In *Proc. CVPR 2023*, 16133–16142.
- Wu, S.; Fei, H.; Li, X.; Ji, J.; Zhang, H.; Chua, T.-S.; and Yan, S. 2025. Towards semantic equivalence of tokenization in multimodal llm. In *Proc. ICLR 2025*.
- Wu, Y.-C.; Gebru, I. D.; Marković, D.; and Richard, A. 2023. Audiodec: An open-source streaming high-fidelity neural audio codec. In *Proc. ICASSP 2023*, 1–5. IEEE.
- Xin, D.; Tan, X.; Takamichi, S.; and Saruwatari, H. 2024. Bigcodec: Pushing the limits of low-bitrate neural speech codec. *arXiv preprint arXiv:2409.05377*.
- Ye, Z.; Sun, P.; Lei, J.; Lin, H.; Tan, X.; Dai, Z.; Kong, Q.; Chen, J.; Pan, J.; Liu, Q.; et al. 2025a. Codec does matter: Exploring the semantic shortcoming of codec for audio language model. In *Proc. AAAI 2025*, volume 39, 25697–25705.
- Ye, Z.; Zhu, X.; Chan, C.-M.; Wang, X.; Tan, X.; Lei, J.; Peng, Y.; Liu, H.; Jin, Y.; DAI, Z.; et al. 2025b. Llasa: Scaling Train-Time and Inference-Time Compute for Llama-based Speech Synthesis. *arXiv preprint arXiv:2502.04128*.
- Yoon, J. W.; Woo, B. J.; and Kim, N. S. 2024. HuBERT-EE: Early Exiting HuBERT for Efficient Speech Recognition. In *Proc. Interspeech 2024*, 2400–2404.
- Zeghidour, N.; Luebs, A.; Omran, A.; Skoglund, J.; and Tagliasacchi, M. 2021. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30: 495–507.
- Zen, H.; Dang, V.; Clark, R.; Zhang, Y.; Weiss, R. J.; Jia, Y.; Chen, Z.; and Wu, Y. 2019. LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech. In *Proc. Interspeech 2019*, 1526–1530.
- Zeng, A.; Du, Z.; Liu, M.; Zhang, L.; Dong, Y.; Tang, J.; et al. 2025. Scaling Speech-Text Pre-training with Synthetic Interleaved Data. In *Proc. ICLR 2025*.
- Zhai, L.; Ding, H.; Zhao, C.; Wang, G.; Zhi, W.; Xi, W.; et al. 2025. One Quantizer is Enough: Toward a Lightweight Audio Codec. *arXiv preprint arXiv:2504.04949*.
- Zhang, H.; Guo, Y.; Li, Z.; Hao, X.; Chen, X.; and Yu, K. 2025. Unlocking Temporal Flexibility: Neural Speech Codec with Variable Frame Rate. *arXiv preprint arXiv:2505.16845*.
- Zhang, X.; Zhang, D.; Li, S.; Zhou, Y.; and Qiu, X. 2024. SpeechTokenizer: Unified Speech Tokenizer for Speech Language Models. In *Proc. ICLR 2024*.