

Codec Does Matter: Exploring the Semantic Shortcoming of Codec for Audio Language Model

Zhen Ye¹, Peiwen Sun¹, Jiahe Lei³, Hongzhan Lin⁴,
Xu Tan², Zheqi Dai⁵, Qiuqiang Kong⁵, Jianyi Chen¹,
Jiahao Pan¹, Qifeng Liu¹, Yike Guo^{1*}, Wei Xue^{1*}

¹Hong Kong University of Science and Technology

²Microsoft

³University of Science and Technology Beijing

⁴Hong Kong Baptist University

⁵Chinese University of Hong Kong

Abstract

Recent advancements in audio generation have been significantly propelled by the capabilities of Large Language Models (LLMs). The existing research on audio LLM has primarily focused on enhancing the architecture and scale of audio language models, as well as leveraging larger datasets, and generally, acoustic codecs, such as EnCodec, are used for audio tokenization. However, these codecs were originally designed for audio compression, which may lead to suboptimal performance in the context of audio LLM. Our research aims to address the shortcomings of current audio LLM codecs, particularly their challenges in maintaining semantic integrity in generated audio. For instance, existing methods like VALL-E, which condition acoustic token generation on text transcriptions, often suffer from content inaccuracies and elevated word error rates (WER) due to semantic misinterpretations of acoustic tokens, resulting in word skipping and errors. To overcome these issues, we propose a straightforward yet effective approach called X-Codec. X-Codec incorporates semantic features from a pre-trained semantic encoder before the Residual Vector Quantization (RVQ) stage and introduces a semantic reconstruction loss after RVQ. By enhancing the semantic ability of the codec, X-Codec significantly reduces WER in speech synthesis tasks and extends these benefits to non-speech applications, including music and sound generation. Our experiments in text-to-speech, music continuation, and text-to-sound tasks demonstrate that integrating semantic information substantially improves the overall performance of language models in audio generation. The code and model have been released.

Introduction

In recent years, Large Language Models (LLMs) such as GPT (Brown et al. 2020) have demonstrated remarkable capabilities in modeling complex, high-dimensional data across various domains, including text and image generation (Zhao et al. 2023; Liu et al. 2024a). Inspired by these successes, there has been significant interest (Agostinelli et al. 2023; Borsos et al. 2023; Wang et al. 2023; Yang et al. 2023b) in exploring the application of LLMs to audio generation.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Audio codecs (Zeghidour et al. 2021) have emerged as a critical technique for audio LLMs, bridging the gap between continuous audio waveforms and token-based language models. By discretizing high-rate audio signals into a finite set of tokens, these codecs enable the application of LLM architectures to audio data, leveraging the successes of textual LLMs.

However, prior research on audio codecs has primarily focused on achieving lower compression rates and higher reconstruction quality (Kumar et al. 2024; Défossez et al. 2022; Yang et al. 2023a). Meanwhile, many efforts in audio generation have concentrated on enhancing model architecture, scaling, or leveraging larger datasets. For instance, AudioLM (Borsos et al. 2023) adopts a two-stage pipeline that models the acoustic token in an autoregressive way conditioned on the semantic token. VALL-E (Wang et al. 2023), the first TTS framework to leverage large, diverse, and multi-speaker speech data, demonstrates strong in-context learning capabilities similar to GPT-3, treating TTS as a language modeling task on audio codecs. MusicGen (Copet et al. 2024) generates music using a single-stage transformer LM alongside efficient token interleaving patterns. Similarly, UniAudio (Yang et al. 2023b) scaled up to 165K hours of audio and 1B parameters, utilizing LLM techniques to generate tokens for various types of audio, including speech, sounds, music, and singing, given different input conditions.

While these works have shown success in developing audio language models, they all rely on the acoustic codecs such as Encodec (Défossez et al. 2022) or Soundstream (Zeghidour et al. 2021) for audio tokenization and de-tokenization. However, these acoustic codecs were originally designed for audio compression rather than for audio language models. This misalignment means the design may not be optimal for audio language modeling.

To design a better audio codec for Audio LLMs, we drew inspiration from the initial purpose of LLMs such as GPT, which were designed to process text. These models focus on understanding and generating natural language, which is inherently rich in semantics. Motivated by this, we assume that a better audio tokenizer should encapsulate rich semantic information to facilitate an easy understanding of audio

content, thus reducing the language model’s burden in interpreting tokens. However, most audio codecs focus on acoustic reconstruction which ignores the semantic information. As a result, LLM essentially tries to predict the local fluctuations of the audio signal, which is difficult, and methods like VALL-E, which condition acoustic token generation on text transcriptions, frequently result in content inaccuracies causing elevated word error rates (WER), stemming from the semantic misinterpretations of acoustic tokens, leading to word skipping and errors.

To address this issue, approaches like SpeechTokenizer (Zhang et al. 2023) have attempted to disentangle speech into separate tokens for content and timbre and perform distillation-based semantic and acoustic integration. However, this method may not integrate smoothly with all audio LLMs, especially those requiring uniform token treatment across different layers, such as utilizing flattened codec tokens (Yang et al. 2023b; Copet et al. 2024).

In this paper, We propose a straightforward yet effective method termed “X-codec”, which integrates both semantic and acoustic features into a unified tokenization framework. The X-Codec architecture employs a distinctive “X-shaped” structure, characterized by two inputs and two outputs, unifying semantic and acoustic information within a single Residual Vector Quantizer (RVQ) structure. This design enables simultaneous embedding learning of semantic richness and acoustic fidelity for every token, resulting in better performance for audio LLM.

We have conducted comprehensive evaluations of X-Codec across various applications, including text-to-speech, music continuation, and text-to-sound synthesis. The results consistently demonstrate the effectiveness of the proposed method. Furthermore, our comparative evaluation on VALL-E based TTS demonstrates that X-Codec outperforms existing disentanglement techniques, thereby highlighting its efficacy and versatility in advancing audio LLM technologies.

Related Works

Audio Language Model

The success of Large Language Models (LLMs) has sparked a significant trend in leveraging language foundation models for audio generation tasks (Rubenstein et al. 2023; Zhang et al. 2024; Wu et al. 2023b,a; Yang et al. 2023b; Chen et al. 2023). Audio, much like language, consists of variable-length sequences, making it well-suited for modeling with language foundation models. One pioneering method, AudioLM (Borsos et al. 2023), employs a multi-stage strategy to harness the predictive capabilities of foundation models for generating tokens unconditionally. This approach involves predicting semantic tokens from various conditions (e.g., phonemes, text descriptions, MIDI) in the initial stage, followed by transforming them into acoustic tokens through coarse-to-fine modeling, ultimately generating the waveform. Representative systems such as SPEAR-TTS (Kharitonov et al. 2023) for speech synthesis and MusicLM (Agostinelli et al. 2023) for music generation have also been proposed. However, the two-stage process can lead to complexity in training and suboptimal performance due to the

separate development of semantic and acoustic tokens, leading to error accumulation.

Conversely, recent advancements have shown that methods employing a single-stage language model outperform two-stage approaches. For example, VALL-E (Wang et al. 2023) utilizes an autoregressive (AR) model to predict the first token and a non-autoregressive (NAR) model to estimate the residual tokens, demonstrating superior performance compared to AudioLM. Similarly, MusicGen (Copet et al. 2024) employs a single-stage transformer language model and incorporates a delay pattern strategy for efficient token interleaving, achieving better results than MusicLM. Other notable works include CLAM-TTS (Kim et al. 2024), VoiceCraft (Peng et al. 2024), and UniAudio (Yang et al. 2023b).

Despite recent advancements, directly modeling the intricate low-level acoustic fluctuations with an LLM poses challenges. LLMs are primarily designed for processing natural language, which is inherently rich in semantics. In order to overcome this limitation, we propose X-Codec, a novel enhancement that aims to enrich semantic processing within acoustic codecs. By doing so, we aim to improve the overall performance of audio LLMs.

Audio Codec

Recent advancements have seen a surge in deep learning methodologies employing vector quantization (Van Den Oord, Vinyals et al. 2017) to reconstruct continuous signals into discrete representations for AR generation. Notably, audio codecs based on the VQ-GAN framework (Esser, Rombach, and Ommer 2021) have gained prominence. For example, SoundStream (Zeghidour et al. 2021) introduces a versatile codec adaptable to various audio types, integrating Residual Vector Quantization (RVQ) and Generative Adversarial Network (GAN) to refine quantization and reconstruction. Similarly, Encodec (Défossez et al. 2022) enhances compression through a multi-scale discriminator and a loss-balancing strategy alongside a language model. HiFi-Codec (Yang et al. 2023a) employs Group-Residual Vector Quantization (GRVQ) to minimize the need for extensive codebooks while maintaining high reconstruction fidelity. DAC (Kumar et al. 2024) addresses codebook collapse, where some codes remain unused, by applying improved codebook learning to achieve higher compression rates.

These codecs primarily focus on acoustic reconstruction and higher compression rates, often overlooking their potential as tokenizers for audio LLMs. Some attempts have been made to develop more suitable tokenizers for audio LLMs. For example, SpeechTokenizer (Zhang et al. 2023) utilizes HuBERT to separate speech into distinct VQ components for content and timbre/acoustic details. This separation improves the modeling of content in the AR stage of VALL-E, while the NAR stage enriches the acoustic details. However, a distillation framework is exploited, this makes SpeechTokenizer may not be compatible with all LLM architectures, especially those that require uniform treatment of tokens, such as methods using flattened codec tokens (Yang et al. 2023b; Copet et al. 2024). Another attempt is pre-

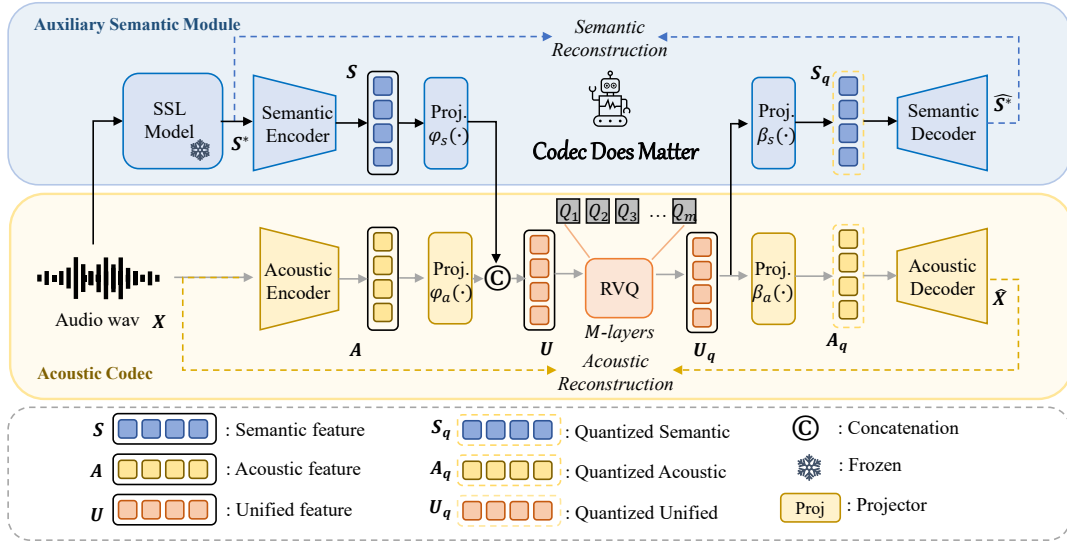


Figure 1: The pipeline of X-codec.

sented by SemantiCodec (Liu et al. 2024b), which employs a pre-trained AudioMAE (Huang et al. 2022) to generate distinct semantic and acoustic tokens from mel-spectrograms. However, this method inherits the issues of SpeechTokenizer and introduces additional complexity in token modeling. Moreover, since the audioMAE is performed on 2D time-frequency mel-spectrograms, LLMs must effectively handle dual scales (time and frequency), which may require significant modifications to existing LLM structures.

In contrast, our proposed X-Codec provides a uniform and comprehensive enhancement of semantic information for all tokens, resulting in significant performance improvements for existing audio LLMs without requiring any structural modifications.

Methods

In this section, we propose X-codec, a straightforward yet effective method to overcome the semantic shortcomings of the current acoustic codecs.

Acoustic Audio codec

As illustrated in Figure 1, our model builds upon the framework established by existing acoustic codecs such as Encodec (Défossez et al. 2022) and DAC (Kumar et al. 2024). An acoustic audio codec is composed of three main components: an acoustic encoder, a quantizer, and an acoustic decoder. The input of the codec is the raw waveform $\mathbf{X} \in \mathbb{R}^n$, where n represents the number of waveform samples. This waveform is fed into the acoustic encoder, which consists of several convolutional layers and employs temporal downscaling to extract frame-level latent acoustic features $\mathbf{A} \in \mathbb{R}^{H_a \times T}$, where H_a denotes the hidden size of the acoustic features and T is the number of frames. These continuous features are then transformed into a series of discrete tokens $\mathbf{Q} \in \mathbb{R}^{M \times T}$ using a Residual Vector Quantizer (RVQ) with M quantizer layers. During training, a

specific codebook for the quantizer is learned, enabling the conversion of discrete tokens back to continuous features $\mathbf{A}_q \in \mathbb{R}^{H_a \times T}$. The acoustic decoder then reconstructs the waveform $\hat{\mathbf{X}}$ from \mathbf{A}_q using several convolutional layers and temporal upsampling. The training process is supervised using various losses, including mel loss, STFT loss, and GAN loss, to ensure high-quality acoustic reconstruction.

Analysing Semantic Shortcoming

In this section, we investigate the impact of acoustic codecs on the performance of audio LLMs, focusing specifically on VALL-E, a pioneering model that leverages language model principles for text-to-speech. Our analysis reveals that training VALL-E using Encodec results in high word error rates (WER) and frequent inaccuracies in content generation. For example, when the input text “he passed through Henley Saint Albans and came so near to London as Harrow on the Hill” is synthesized, it is erroneously produced as “he passed through henley saint albeans and camsel knew to lunglan as herold the lor”. This misinterpretation, which is beyond simply improving the audio quality, suggests a fundamental limitation in Encodec’s ability to differentiate phonemes, possibly due to its inadequate semantic processing capabilities.

To substantiate the above hypothesis, we conducted Phonetic Discriminability ABX Tests to evaluate the phonetic discriminability of Encodec’s representations. The details are provided in the experiment section. Our findings reveal that Encodec’s representations exhibit poor phonetic discriminability, which confirms the presence of semantic inadequacies in the codec. Based on these results, we assert that these semantic shortcomings are a significant contributing factor to the observed inaccuracies of language model based audio generation.

To effectively address these semantic limitations, we introduce a novel approach that integrates more comprehen-

sive semantic features into the codec’s architecture. This enhancement is designed to enrich the codec’s understanding of audio content, thereby alleviating the interpreting load on the language model. Detailed elaboration of this method is provided in the subsequent section.

Designing Auxiliary Semantic Module

Our approach employs a straightforward method that enhances audio codecs by directly concatenating semantic and acoustic features. Initially, we extract the semantic feature vector $\mathbf{S}^* \in \mathbb{R}^{H_s \times T}$ from the audio waveform \mathbf{x} . This extraction utilizes a self-supervised, pre-trained model such as HuBERT (Hsu et al. 2021) or wav2vec 2.0 (Baevski et al. 2020). The extracted features are then processed through multiple convolutional layers within a semantic encoder to yield the refined semantic feature vector \mathbf{S} . Concurrently, the acoustic branch produces the feature \mathbf{A} . These outputs, \mathbf{S} and \mathbf{A} , are subsequently concatenated using a linear projection ϕ , formulated as:

$$\mathbf{U} = \text{concat}(\phi_s(\mathbf{S}), \phi_a(\mathbf{A})), \quad (1)$$

where the concatenated feature $\mathbf{U} \in \mathbb{R}^{H_u \times T}$ is designed to maximize information preservation from both semantic and acoustic sources. This combined feature is then subject to RVQ using an M -layer quantizer, resulting in tokens that encapsulate a rich mixture of semantic and acoustic information.

The quantized feature \mathbf{U}_q is designed to meet the decoder’s objectives through two projectors, β_s and β_a , which enable the decoders to reconstruct the original semantic feature $\hat{\mathbf{S}}^*$ and the audio waveform $\hat{\mathbf{x}}$. We adhere to established acoustic reconstruction methods from previous works while introducing a Mean Squared Error (MSE) loss specifically for the reconstruction of semantic features. Furthermore, a constant weight γ is applied to the semantic loss to ensure that its scale is aligned with other losses, thus promoting a balanced training objective.

Experiments

Given that established audio codecs such as Encodec, Spechtokenizer, and DAC are trained on diverse datasets with varying configurations, we meticulously design experiments to rigorously evaluate the efficacy of our proposed solution, X-Codec. To ensure a fair and unbiased comparison, each experiment employs a baseline acoustic codec that is precisely aligned with our X-Codec in terms of training data, training steps, and other hyperparameters. The primary distinction between the baseline codec and X-Codec lies in the exclusion of the auxiliary semantic module in the baseline configuration. This controlled experimental design enables us to isolate and evaluate the specific contributions of our semantic enhancements to the overall performance of the audio LLMs.

Text-to-Speech

In this subsection, we critically evaluate the performance of various audio codecs in training the VALL-E model for zero-

shot Text-to-Speech (TTS) tasks. Our investigation is guided by two primary objectives:

- To determine whether the X-Codec can enhance the performance of audio LLMs in TTS applications.
- To evaluate the comparative advantages of X-Codec over the disentanglement strategy employed by SpeechTokenizer, specifically within the context of the VALL-E model.

Baselines For a comprehensive comparison, we employ several state-of-the-art neural audio codecs as baselines:

- **EnCodec**¹: The open-source EnCodec model (Défossez et al. 2022), trained on a diverse range of 24kHz audio data, can compress audio to bitrates between 1.5 and 24.0 kbps while maintaining high fidelity.
- **DAC**²: The open-source DAC model (Kumar et al. 2024) utilizes enhanced VQ techniques. For our experiments, we employ the official 16kHz version.
- **SpeechTokenizer**³: This model (Zhang et al. 2023) is a unified speech tokenizer that leverages distinct VQ layers to separate speech into content and timbre components. We utilize their official checkpoints in our evaluations.

Training Details of X-Codec Given our objective to assess the efficacy of X-Codec in leveraging semantic information, we meticulously align our experimental setup with that used for SpeechTokenizer. Both models are trained on the same dataset, LibriSpeech, and utilize the same pre-trained self-supervised representations from HuBERT-base-ls960⁴. To ensure comparability, we also adopt the strategy of employing the average representation across various layers of HuBERT as our semantic training objective.

Training Details of VALL-E For reproduction of the VALL-E, we utilize the resources specified in the provided repository⁵. The training data is the LibriTTS, retaining the default settings as specified in the repository, except for the learning rate during the AR stage, which is adjusted to 0.01 to enhance model stability. The training process span 100 epochs for the AR stage and 200 epochs for the non-autoregressive (NAR) stage, same for all audio codecs for a fair comparison.

Evaluation Metrics To assess the performances of zero-shot TTS systems, we employ the following metrics:

- **WER (Word Error Rate)**: We utilize an Automatic Speech Recognition (ASR) model to transcribe the generated audio (Wang et al. 2023). The discrepancies between these transcriptions and the original texts are quantified using WER, providing a critical measure of audio intelligibility.
- **Sim-O (Similarity Objective)**: This metric assesses the objective similarity between synthesized speech and the

¹https://huggingface.co/facebook/encodec_24khz

²<https://github.com/descriptinc/descript-audio-codec>

³<https://github.com/ZhangXInFD/SpeechTokenizer>

⁴<https://huggingface.co/facebook/hubert-base-ls960>

⁵<https://github.com/lifeiteng/vall-e>

original reference speech. Sim-O uses feature embeddings extracted from a pre-trained speaker verification model to measure this similarity (Hsu et al. 2021; Kim et al. 2024)⁶, reflecting the codec’s ability to preserve speaker characteristics.

- **UTMOS:** We evaluate the audio quality using UTMOS, a Speech MOS (Mean Opinion Score) predictor (Saeki et al. 2022)⁷ that automatically measures the naturalness of speech. This metric provides insights into the overall auditory quality of the synthesized speech.

Zero-shot TTS Results We use librispeech-test-clean (Panayotov et al. 2015) for zero-shot TTS evaluation following VALL-E-continual-setting (Wang et al. 2023). The results in Table 1 demonstrate the following key findings:

- When comparing both X-Codec and SpeechTokenizer against the baseline and other acoustic codecs like DAC and Encodec, we observe improvements in WER. This supports our hypothesis that integrating semantic information helps audio LLMs better understand content.
- Comparing the baseline acoustic codec and SpeechTokenizer, SpeechTokenizer exhibited lower Sim-O scores. We attribute this reduction to its initial disentanglement phase, which exclusively focuses on content prediction. This specialization potentially hampers the NAR phase’s ability to accurately reconstruct speaker timbre when conditioned solely on tokens derived from the primary content-focused stage, resulting in poor speaker similarity.
- X-Codec not only shows better WER but also higher Sim-O and UTMOS scores compared to SpeechTokenizer. This confirms the effectiveness of our approach, indicating that our codec handles the integration of semantic and acoustic information more proficiently.

Analysing the Effect of Codec To further analyse the above results caused by different audio codecs, we evaluate phonetic discriminability using the ABX error rate (Schatz et al. 2013). This metric assesses how well different codecs can distinguish between similar phonetic sounds within and across various contexts. We specifically examine the continuous representations for VQ as indicated by the results in the following table 2. We compare the performance of various models in terms of within and across phonetic discriminability:

Key insights include:

- Both SpeechTokenizer and X-Codec significantly outperform pure acoustic codecs like Encodec and DAC in phonetic discriminability, which supports our claim that enhancing semantic understanding in codecs helps modelling content such as phonetic details.
- The X-Codec demonstrates a notable trend of improved phonetic discriminability with an increase in the number of quantizations (nq). Specifically, as nq increases

from 1 to 8, the ABX error rates consistently decrease, thereby highlighting effectiveness of the X-Codec’s design in enhancing semantic integration across multiple quantization layers.

- In contrast, the SpeechTokenizer, while exhibiting commendable performance at a lower quantization level (nq = 1), fails to show significant improvement as nq is increased. This suggests a design limitation; the codec’s reliance on the initial quantization to carry semantic information restricts its ability to process a broader spectrum of semantic information. Notably, the performance of X-Codec at nq = 8 significantly exceeds that of SpeechTokenizer.

These results underline the effectiveness of our method in facilitating enhanced semantic integration, leading to better phonetic discriminability and audio LLMs. In addition, these results also show that our simple concatenate methods surpass disentangle methods such as speechookenizer.

Music and Sound Generation

To the best of our knowledge, this is the first exploration into the potential benefits of incorporating semantic information into audio codecs for enhancing music and general sound generation through audio LLMs. Conventional methods for general audio representation learning, aiming at capturing the semantic discriminability of audios, are generally based on 2D mel-spectrogram, such as AudioMAE (Huang et al. 2022) and Beats (Chen et al. 2022). These methods are in stark contrast to traditional codecs that process audio sequentially, frame-by-frame. This difference poses challenges for direct integration into existing audio generation frameworks.

To bridge this gap, we have developed a variant of HuBERT, specifically adapted for general audio, which we refer to as HuBERT-General-Audio. This HuBERT-General-Audio is trained on an expansive internal dataset of approximately 200,000 hours, with a similar distribution as AudioSet. Additionally, our proposed X-Codec is also trained using these data for 400,000 steps until convergence, incorporating the HuBERT-General-Audio model within its semantic module. For a fair comparison, we train a baseline acoustic codec under identical settings but excluding semantic information.

Training Details of Self-Supervised General Audio Representation HuBERT-General-Audio is trained using 8 NVIDIA H800 GPUs on 2.6 million tokens across 325,000 iterations. For training stability, we adopt an inverse square root learning schedule, a modification from the polynomial decay schedule originally utilized in (Hsu et al. 2021). The learning rate is set at 0.0003 with warmup steps of 32,000. Unlike the original HuBERT, which utilizes MFCCs as the training target unit designed specifically for speech, our model leverages the first VQ layer of Encodec as the training target for acoustic unit discovery in the general audio. This choice eliminates the need for the K-means discretization step, saving significant time and computational resources.

⁶https://github.com/microsoft/UniSpeech/tree/main/downstreams/speaker_verification

⁷<https://github.com/tarepan/SpeechMOS>

Codec	Training Data of Audio Codec	VALL-E AR stage			VALL-E AR+NAR stages		
		WER ↓	SIM-O ↑	UTMOS ↑	WER ↓	SIM-O ↑	UTMOS ↑
GT	-	2.23	0.67	4.10	2.23	0.67	4.10
Encodec (Défossez et al. 2022)	C+DNS+AS+FSD+J	47.17	0.09	1.24	6.37	0.33	3.02
DAC (Kumar et al. 2024)	C+DNS+V+AS+J+M	85.55	0.03	1.24	6.81	0.34	3.31
Spechttokenizer (Zhang et al. 2023)	LibriSpeech	7.53	0.10	1.26	5.24	0.36	3.84
Baseline Acoustic Codec	LibriSpeech	22.32	0.16	3.15	7.70	0.41	3.89
X-Codec-hubert	LibriSpeech	5.27	0.22	3.85	4.07	0.42	4.16
X-Codec-wavlm-base-plus	MLS English	4.83	0.24	4.02	3.26	0.41	4.22

Table 1: Objective performance comparison on *continuation* zero-shot speech synthesis tasks using VALL-E trained on LibriTTS with different audio codec. Abbreviation: C (Common Voice), DNS (DNS Challenge 4 speech), AS (AudioSet), FSD (FSD50K), J (Jamendo), V (VCTK), M(MUSDB)

Model	n_q	within ↓	across ↓
hubert-ls-960 (Hsu et al. 2021)	-	3.3	4.1
Encodec (Défossez et al. 2022)	1	21.5	28.3
Encodec (Défossez et al. 2022)	8	17.5	27.0
DAC (Kumar et al. 2024)	1	26.3	32.7
DAC (Kumar et al. 2024)	12	21.7	33.2
Spechttokenizer (Zhang et al. 2023)	1	3.5	4.3
Spechttokenizer(Zhang et al. 2023)	8	3.6	4.5
Baseline Acoustic Codec	1	26.4	31.2
Baseline Acoustic Codec	8	20.1	28.3
X-Codec	1	3.9	4.9
X-Codec	8	3.3	4.3

Table 2: Comparison of Phonetic Discriminability within and across ABX error rate for various models, with different n_q values. Lower values indicate better performance.

Model	FD ↓	FAD ↓	FD-MERT-layer-9 ↓
Acoustic codec	16.17	1.43	2.88
X-Codec	12.66	1.37	2.62

Table 3: Comparison between baseline acoustic codec and our X-Codec on music continue.

Music Continuation Training Details: Acquiring high-quality text-music pair data is challenging; therefore, we gathered approximately 100,000 hours of music-only data, including about one million songs for the music continuation task. We deployed nanoGPT⁸ to implement a GPT-2-medium (approximately 300M parameters) (Radford et al. 2019) as our generative model. This model utilizes the first VQ from our codec to construct the training sequences, with additional experiments involving multiple VQs detailed in the appendix. We set the block size of sequence modelling to 4096, corresponding to roughly 82 seconds of audio, and adjust the vocabulary size from 50,257 to 1024, matching our codec’s codebook size. Other training hyperparameters are consistent with previous GPT-2-medium configurations. We train 300,000 steps on 8 NVIDIA H800 GPUs. The batch size is set to 20, with a learning rate of 3e-4 and a warmup phase of 2000 steps.

Experiments: For music continuation, we randomly crop 600 samples with each 40 seconds in duration from the MUSDB18 dataset (Rafii et al. 2017). The initial 10 seconds of each sample are used as prompts for the audio LLM,

⁸<https://github.com/karpathy/nanoGPT>

Model	FD ↓	IS ↑	FAD ↓	CLAP ↑
Acoustic codec	59.03	3.89	6.19	0.417
X-Codec	46.31	5.29	4.10	0.483

Table 4: Comparison between baseline acoustic codec and X-Codec on text-to-sound tasks.

while the subsequent 30 seconds are generated by the model. These generated segments are then compared against the corresponding ground truth (GT) segments. To ensure that the assessment is independent of the codec’s reconstruction fidelity, both the generated and GT audio are reconstructed using the first VQ layer of the codec, ensuring performance differences attributed solely to the generative models themselves.

The evaluation metrics of the generated music include: Frechet Distance (FD) computed using features from Pre-trained Audio Neural Networks (PANNs) (Kong et al. 2020), Frechet Audio Distance (FAD), and FD-MERT Layer 9 (Li et al. 2023). The results, as summarized in Table 3, reveal that the X-Codec significantly outperforms the baseline acoustic codec across all metrics. This superior performance indicates the X-Codec has a better understanding and enabling more effective reproduction of complex musical structures.

Text-to-Sound Training Details: Still, GPT-2-medium (approximately 300M parameters) are adopted for conditional text-to-sound tasks, where the condition embedding is extracted from text captions using LAION-CLAP (Wu et al. 2023c) and linearly projected from 512 dimensions to 1024 dimensions for GPT input. The training data consists of approximately 400 hours of audio content sourced from the AudioCaps dataset (Kim et al. 2019) and the AudioSet SL subset from the WavsCaps dataset (Mei et al. 2024). All audio samples are uniformly resampled to a 16kHz sampling rate. The first four tokens from the VQ layers are pre-processed and flattened to configure the GPT model’s block size to 2000, corresponding to a processing rate of 50Hz. The training process spans 80,000 steps on four NVIDIA 4090 GPUs, with a batch size of 8 and a learning rate of 3e-4. A warmup phase of 2000 steps is employed to optimize the training process.

Evaluation Metrics: following (Huang et al. 2023) and

Model/Datasets	ESC-50	US8K	FSD50K	VIVAE	FMA	MTT	IRMAS	MS-DB	RAVDESS	A-MNIST	SLURP	EMOVO
DAC (Kumar et al. 2024)	27.65	45.16	7.08	30.80	38.50	27.69	30.33	51.79	37.50	73.59	7.72	23.46
Encodec (Défossez et al. 2022)	30.60	64.47	8.63	31.59	39.5	26.57	28.16	64.47	31.60	78.71	8.44	25.51
Baseline Acoustic Codec	40.00	55.38	11.10	37.70	48.75	32.26	36.26	62.77	47.22	82.58	9.28	24.83
Hubert-general-audio	69.95	74.87	34.27	48.35	64.50	43.35	49.80	74.09	69.10	99.43	21.27	35.03
X-Codec	69.85	75.37	34.05	49.40	64.63	42.95	52.24	75.63	68.40	99.48	21.65	35.20

Table 5: Performance of semantic representation on the ARCH benchmark. The table shows the performance of various models across different domains. ESC-50 (Piczak 2015), US8K (Salamon, Jacoby, and Bello 2014), FSD50K (Fonseca et al. 2022), and VIVAE (Holz, Larrouy-Maestri, and Poeppel 2022) represent performance on Acoustic Events. FMA (Defferrard et al. 2017), MTT (Law et al. 2009), IRMAS (Bosch et al. 2012), and MS-DB (Bittner et al. 2014) indicate performance in the Music domain. RAVDESS (Livingstone and Russo 2018), AudioMNIST (Becker et al. 2024), SLURP (Bastianelli et al. 2020), and EMOVO (Costantini et al. 2014) reflect performance in the Speech domain. Higher values indicate better performance across these tasks.

(Liu et al. 2023), we calculate Frechet Distance (FD), Inception Score (IS), Frechet Audio Distance (FAD) for text-to-audio generation. In addition, CLAP score (Huang et al. 2023) is used to evaluate the correspondence between the generated audio and the text prompt.

Experiment Results: As shown in Table 4, the proposed X-Codec significantly outperforms the baseline acoustic codec across all metrics. These results demonstrate that semantic information integration significantly enhances the codec’s capability, underscoring the value of semantic enrichment in audio generation tasks.

Analysing the Effect of Codec We hypothesize that the enhanced audio generation capabilities of the audio LLMs are attributed to the improved semantic understanding facilitated by the X-Codec. To validate this hypothesis, we employ the ARCH benchmark (La Quatra et al. 2024) to evaluate the audio semantic understanding, and the benchmark is a comprehensive framework specifically designed to evaluate automatic recognition learning methods across a diverse range of audio classification domains, including acoustic events, music, and speech. The results from this benchmark are shown in Table 5.

Our findings indicate that HuBERT-general-audio significantly outperforms traditional acoustic codecs such as DAC, Encodec, and the baseline acoustic codec across all metrics. This improvement highlights the enhanced semantic understanding of X-Codec for general audio, which appears to be lacking in conventional acoustic audio codecs.

Moreover, X-Codec achieves performance that is comparable or even superior to HuBERT-general-audio, confirming the effectiveness of our approach to enhancing semantic processing within codecs. This equivalence or improvement indicates the capability of X-Codec to integrate semantic information robustly.

Limitation

While our method significantly enhances the performance of codecs for LLMs by integrating semantic information, it does come with certain trade-offs. According to the principle of “no free lunch,” improving one aspect of a system often involves compromises in others. In the case of our enhanced codecs, the primary limitation lies in their potential impact on the original functionality of codecs, which is compression

for information transmission. The introduction of a semantic extraction layer adds additional computational overhead, potentially increasing the time required for processing. This can affect the efficiency of the codec when used in applications where rapid data compression and transmission are critical. Consequently, while our approach offers substantial benefits for semantic understanding and audio processing, it may not be as effective in contexts where high-speed data compression is paramount.

Model	Mel DT.	STFT DT.	UTMOS
Baseline ($n_q=1$)	0.79	0.73	2.96
Baseline ($n_q=8$)	0.54	0.58	3.72
X-codec ($n_q=1$)	0.86	0.77	3.71
X-codec ($n_q=8$)	0.62	0.63	4.01

Table 6: Comparison of reconstruction based on Mel DT., STFT DT., and UTMOS metrics using 1000 LibriTTS speech samples. “DT.” is short for distance

Furthermore, the integration of semantic layers can slightly impair certain acoustic metrics such as Mel and STFT distance, which are crucial for maintaining the fidelity of compressed audio. However, it is essential to note that these trade-offs are counterbalanced by significant improvements in human auditory perception, as evidenced by the UTMOS scores.

Conclusion

In this paper, we introduced X-codec, an advanced audio codec that integrates semantic information through self-supervised learning models to enhance performance in large language models, specifically in text-to-speech synthesis, music continuation, and general audio classification tasks. Our evaluations demonstrate that X-codec significantly improves semantic understanding and audio generation quality across a variety of domains.

Acknowledgments

The research was supported by NSFC (No. 62206234), Early Career Scheme (ECS-Hong Kong University of Science and Technology22201322), Theme-based Research Scheme (T45-205/21-N) from Hong Kong RGC, and Generative AI Research and Development Centre from InnoHK.

References

- Agostinelli, A.; Denk, T. I.; Borsos, Z.; Engel, J.; Verzetti, M.; Caillon, A.; Huang, Q.; Jansen, A.; Roberts, A.; Tagliasacchi, M.; et al. 2023. MusiclM: Generating music from text. *arXiv preprint arXiv:2301.11325*.
- Baevski, A.; Zhou, Y.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33: 12449–12460.
- Bastianelli, E.; Vanzo, A.; Swietojanski, P.; and Rieser, V. 2020. SLURP: A Spoken Language Understanding Resource Package. In *EMNLP*. ACM.
- 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*.
- Bittner, R. M.; Salamon, J.; Tierney, M.; Mauch, M.; Cannam, C.; and Bello, J. P. 2014. Medleydb: A multitrack dataset for annotation-intensive mir research. In *Proc. Intl. Soc. Music Information Retrieval Conf. (ISMIR)*, volume 14, 155–160.
- Borsos, Z.; Marinier, R.; Vincent, D.; Kharitonov, E.; Pietquin, O.; Sharifi, M.; Roblek, D.; Teboul, O.; Grangier, D.; Tagliasacchi, M.; et al. 2023. AudiolM: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Bosch, J. J.; Janer, J.; Fuhrmann, F.; and Herrera, P. 2012. A Comparison of Sound Segregation Techniques for Predominant Instrument Recognition in Musical Audio Signals. In *Proc. Intl. Soc. Music Information Retrieval Conf. (ISMIR)*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chen, Q.; Chu, Y.; Gao, Z.; Li, Z.; Hu, K.; Zhou, X.; Xu, J.; Ma, Z.; Wang, W.; Zheng, S.; et al. 2023. Lauragpt: Listen, attend, understand, and regenerate audio with gpt. *arXiv preprint arXiv:2310.04673*.
- Chen, S.; Wu, Y.; Wang, C.; Liu, S.; Tompkins, D.; Chen, Z.; and Wei, F. 2022. Beats: Audio pre-training with acoustic tokenizers. *arXiv preprint arXiv:2212.09058*.
- Copet, J.; Kreuk, F.; Gat, I.; Remez, T.; Kant, D.; Synnaeve, G.; Adi, Y.; and Défossez, A. 2024. Simple and controllable music generation. *Advances in Neural Information Processing Systems*, 36.
- Costantini, G.; Iaderola, I.; Paoloni, A.; and Todisco, M. 2014. EMOVO Corpus: an Italian Emotional Speech Database. In *LREC*. European Language Resources Association (ELRA).
- Defferrard, M.; Benzi, K.; Vandergheynst, P.; and Bresson, X. 2017. FMA: A Dataset For Music Analysis. In *Proc. Intl. Soc. Music Information Retrieval Conf. (ISMIR)*.
- Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2022. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12873–12883.
- Fonseca, E.; Favory, X.; Pons, J.; Font, F.; and Serra, X. 2022. FSD50K: An Open Dataset of Human-Labeled Sound Events. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Holz, N.; Larrouy-Maestri, P.; and Poeppel, D. 2022. The variably intense vocalizations of affect and emotion (VI-VAE) corpus prompts new perspective on nonspeech perception. *Emotion*.
- 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, P.-Y.; Xu, H.; Li, J.; Baevski, A.; Auli, M.; Galuba, W.; Metze, F.; and Feichtenhofer, C. 2022. Masked autoencoders that listen. *Advances in Neural Information Processing Systems*, 35: 28708–28720.
- Huang, R.; Huang, J.; Yang, D.; Ren, Y.; Liu, L.; Li, M.; Ye, Z.; Liu, J.; Yin, X.; and Zhao, Z. 2023. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. *arXiv preprint arXiv:2301.12661*.
- Kharitonov, E.; Vincent, D.; Borsos, Z.; Marinier, R.; Girgin, S.; Pietquin, O.; Sharifi, M.; Tagliasacchi, M.; and Zeghidour, N. 2023. Speak, read and prompt: High-fidelity text-to-speech with minimal supervision. *Transactions of the Association for Computational Linguistics*, 11: 1703–1718.
- Kim, C. D.; Kim, B.; Lee, H.; and Kim, G. 2019. AudioCaps: Generating captions for audios in the wild. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 119–132.
- Kim, J.; Lee, K.; Chung, S.; and Cho, J. 2024. CLaM-TTS: Improving Neural Codec Language Model for Zero-Shot Text-to-Speech. *arXiv preprint arXiv:2404.02781*.
- Kong, Q.; Cao, Y.; Iqbal, T.; Wang, Y.; Wang, W.; and Plumbley, M. D. 2020. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28: 2880–2894.
- Kumar, R.; Seetharaman, P.; Luebs, A.; Kumar, I.; and Kumar, K. 2024. High-fidelity audio compression with improved rvqgan. *Advances in Neural Information Processing Systems*, 36.
- 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. *arXiv preprint arXiv:2405.00934*.
- Law, E.; West, K.; Mandel, M. I.; Bay, M.; and Downie, J. S. 2009. Evaluation of algorithms using games: The case of music tagging. In *Proc. Intl. Soc. Music Information Retrieval Conf. (ISMIR)*.

- Li, Y.; Yuan, R.; Zhang, G.; Ma, Y.; Chen, X.; Yin, H.; Lin, C.; Ragni, A.; Benetos, E.; Gyenge, N.; et al. 2023. Mert: Acoustic music understanding model with large-scale self-supervised training. *arXiv preprint arXiv:2306.00107*.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv preprint arXiv:2301.12503*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024a. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liu, H.; Xu, X.; Yuan, Y.; Wu, M.; Wang, W.; and Plumbley, M. D. 2024b. SemantiCodec: An Ultra Low Bitrate Semantic Audio Codec for General Sound. *arXiv preprint arXiv:2405.00233*.
- 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*.
- Mei, X.; Meng, C.; Liu, H.; Kong, Q.; Ko, T.; Zhao, C.; Plumbley, M. D.; Zou, Y.; and Wang, W. 2024. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Panayotov, V.; Chen, G.; Povey, D.; and Khudanpur, S. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 5206–5210. IEEE.
- Peng, P.; Huang, P.-Y.; Li, D.; Mohamed, A.; and Harwath, D. 2024. VoiceCraft: Zero-Shot Speech Editing and Text-to-Speech in the Wild. *arXiv preprint arXiv:2403.16973*.
- Piczak, K. J. 2015. ESC: Dataset for Environmental Sound Classification. In *ACM Multimedia*, MM '15. New York, NY, USA: ACM. ISBN 9781450334594.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Rafii, Z.; Liutkus, A.; Stöter, F.-R.; Mimitakis, S. I.; and Bitner, R. 2017. The MUSDB18 corpus for music separation.
- Rubenstein, P. K.; Asawaroengchai, C.; Nguyen, D. D.; Bapna, A.; Borsos, Z.; Quitry, F. d. C.; Chen, P.; Badawy, D. E.; Han, W.; Kharitonov, E.; et al. 2023. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*.
- Saeki, T.; Xin, D.; Nakata, W.; Koriyama, T.; Takamichi, S.; and Saruwatari, H. 2022. Utmos: Utokyo-sarulab system for voicemos challenge 2022. *arXiv preprint arXiv:2204.02152*.
- Salamon, J.; Jacoby, C.; and Bello, J. P. 2014. A Dataset and Taxonomy for Urban Sound Research. In *ACM Multimedia*, MM '14. New York, NY, USA: ACM. ISBN 9781450330633.
- Schatz, T.; Peddinti, V.; Bach, F.; Jansen, A.; Hermansky, H.; and Dupoux, E. 2013. Evaluating speech features with the minimal-pair ABX task: Analysis of the classical MFC/PLP pipeline. In *INTERSPEECH 2013: 14th Annual Conference of the International Speech Communication Association*, 1–5.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Wang, C.; Chen, S.; Wu, Y.; Zhang, Z.; Zhou, L.; Liu, S.; Chen, Z.; Liu, Y.; Wang, H.; Li, J.; et al. 2023. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*.
- Wu, H.; Chang, K.-W.; Wu, Y.-K.; and Lee, H.-y. 2023a. Speechgen: Unlocking the generative power of speech language models with prompts. *arXiv preprint arXiv:2306.02207*.
- Wu, J.; Gaur, Y.; Chen, Z.; Zhou, L.; Zhu, Y.; Wang, T.; Li, J.; Liu, S.; Ren, B.; Liu, L.; et al. 2023b. On decoder-only architecture for speech-to-text and large language model integration. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 1–8. IEEE.
- Wu, Y.; Chen, K.; Zhang, T.; Hui, Y.; Berg-Kirkpatrick, T.; and Dubnov, S. 2023c. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Yang, D.; Liu, S.; Huang, R.; Tian, J.; Weng, C.; and Zou, Y. 2023a. Hifi-codec: Group-residual vector quantization for high fidelity audio codec. *arXiv preprint arXiv:2305.02765*.
- Yang, D.; Tian, J.; Tan, X.; Huang, R.; Liu, S.; Chang, X.; Shi, J.; Zhao, S.; Bian, J.; Wu, X.; et al. 2023b. Uniaudio: An audio foundation model toward universal audio generation. *arXiv preprint arXiv:2310.00704*.
- Zeghidour, N.; Luebs, A.; Omran, A.; Skoglund, J.; and Tagliasacchi, M. 2021. Soundstream: An end-to-end neural audio codec. *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 30: 495–507.
- Zhang, X.; Zhang, D.; Li, S.; Zhou, Y.; and Qiu, X. 2023. Spechtokenizer: Unified speech tokenizer for speech large language models. *arXiv preprint arXiv:2308.16692*.
- Zhang, Z.; Chen, S.; Zhou, L.; Wu, Y.; Ren, S.; Liu, S.; Yao, Z.; Gong, X.; Dai, L.; Li, J.; et al. 2024. Speechlm: Enhanced speech pre-training with unpaired textual data. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.