

Code Drift: Towards Idempotent Neural Audio Codecs

Patrick O’Reilly
Northwestern University

Prem Seetharaman
Adobe Research

Jiaqu Su
Adobe Research

Zeyu Jin
Adobe Research

Bryan Pardo
Northwestern University

Abstract—Neural codecs have demonstrated strong performance in high-fidelity compression of audio signals at low bitrates. The token-based representations produced by these codecs have proven particularly useful for generative modeling. While much research has focused on improvements in compression ratio and perceptual transparency, recent works have largely overlooked another desirable codec property – *idempotence*, the stability of compressed outputs under multiple rounds of encoding. We find that state-of-the-art neural codecs exhibit varied degrees of idempotence, with some degrading audio outputs significantly after as few as three encodings. We investigate possible causes of low idempotence and devise a method for improving idempotence through fine-tuning a codec model. We then examine the effect of idempotence on a simple conditional generative modeling task, and find that increased idempotence can be achieved without negatively impacting downstream modeling performance – potentially extending the usefulness of neural codecs for practical file compression and iterative generative modeling workflows.

Index Terms—audio compression, neural audio compression, audio codecs, neural audio codecs, generative modeling

I. INTRODUCTION

Neural network-based codecs have shown promise in lossy compression of digital media across modalities such as video, images, and audio [1]–[3]. Whereas signal-processing codecs rely on hand-crafted algorithms for identifying and removing redundant information within a signal, neural codecs learn to encode and decode signals from data directly.

For audio, neural codecs typically encode waveform signals as one or more sequences of discrete tokens, each drawn from a codebook. High compression ratios can be achieved with a combination of small codebooks (e.g. 2^{10} – 2^{12} possible tokens) and low token sequence rates (e.g. 100 – 1000 times less than the source audio). Because these token representations are far more compact than waveform audio, while retaining most perceptually relevant information, it is common to train generative models of audio directly on token sequences [4], [5]. Predicted token sequences can then be mapped to waveform audio via the decoder of the chosen codec.

Idempotence – the stability of a codec’s decoded output, and thus its encoded representation, under multiple rounds of encoding and decoding – is a desirable property for codecs. This is because codecs with low idempotence significantly degrade media under multiple encodings, and may therefore be poorly suited for real-world use. Such multiple-encoding scenarios can arise when digital media is repeatedly compressed for transmission, easier storage, and retargeting for different platforms. The increasing adoption of neural codec-based generative models for audio postprocessing and editing [6]–[8] may also result in repeated encoding and decoding of an audio file before a final version reaches listeners. While some signal-processing codecs explicitly address idempotence or “tandem coding” [9]–[12], the idempotence of neural audio codecs has not been studied in depth. Rather, existing research has focused on *perceptual transparency* – the preservation of output quality under a single encoding.

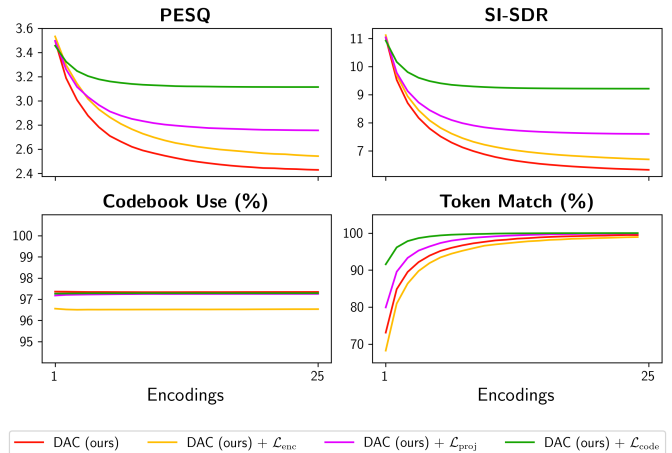


Fig. 1. **Effect of idempotent training:** audio quality (in PESQ and SI-SDR), average codebook use, and match rate between tokens encoded at the n^{th} and $(n + 1)^{\text{th}}$ iterations as a function of encoding iterations for our DAC reproduction and fine-tuned variants with idempotence losses described in Section IV. Fine-tuning with \mathcal{L}_{code} results in the highest overall idempotence through 25 encoding iterations, while preserving codebook use.

Neural audio codecs are typically designed not only for data compression, but to facilitate generative modeling on encoded token representations. While there has been some study on the relationship between representation idempotence and generative models for image codecs [13], we are aware of no such study for audio codecs.

Motivated by the importance of idempotence and its potential impact on generative modeling of encoded representations, we undertake an empirical study of the idempotence of neural audio codecs. Our contributions are as follows:

- We measure the idempotence of state-of-the-art neural audio codecs across speech datasets and find that many codecs demonstrate low stability under re-encoding.
- We develop a method for improving codec idempotence with minimal audio quality degradation through codec fine-tuning and evaluate the impact of different objective functions in encouraging idempotent representations.
- We explore the links between idempotence and properties such as phase sensitivity, codebook use, and reconstruction quality.
- We evaluate the impact of codec idempotence on language modeling and find that idempotence can be increased without negatively affecting downstream modeling performance.

Through this work, we hope to encourage additional study of neural audio codec representations, and of their impact on both lossy compression and downstream generative modeling. We provide audio examples at <https://oreillyp.github.io/codedrift/>.

II. RELATED WORK

Neural Audio Codecs: Recent works have proposed to use deep neural networks for audio compression [3], [16], [21]. These neural

TABLE I
COMPARISON OF CODECS

Codec	Sample Rate (kHz)	Bitrate (kbps)	VCTK		Expresso	
			PESQ @ 1 ↑	PESQ @ 25 ↑	PESQ @ 1 ↑	PESQ @ 25 ↑
MP3 [14]	48	24 / 12 / 6	3.79 / 3.79 / 3.41	1.99 / 1.99 / 2.29	3.49 / 3.49 / 3.32	1.93 / 1.93 / 2.13
OPUS [15]	48	24 / 12 / 6	4.17 / 3.72 / 2.30	1.82 / 1.40 / 1.04	4.13 / 3.58 / 2.12	1.69 / 1.33 / 1.04
DAC [16]	44.1 / 24 / 16	7.7 / 24 / 6	3.97 / 4.30 / 4.04	3.17 / 1.28 / 1.32	3.78 / 4.29 / 3.91	2.89 / 1.24 / 1.27
Encodec [3], [7]	48 / 24 / 16	24 / 24 / 2.2	2.96 / 3.65 / 3.03	1.05 / 1.05 / 2.21	2.87 / 3.34 / 2.91	1.05 / 1.05 / 2.08
SpeechTokenizer [17]	16	4	2.65	1.10	2.41	1.07
FACodec [18]	16	4.8	3.09	1.03	3.03	1.04
Spectral-Codec [19]	44.1 / 22.05	6.9 / 6.9	3.24 / 3.49	1.18 / 1.36	3.16 / 3.41	1.14 / 1.27
ESC [20]	16	9	3.85	2.50	3.67	2.24

codecs encode waveform audio into a hierarchy of token sequences through residual vector quantization (RVQ), and are trained end-to-end with objectives for both audio reconstruction quality and regularization of the encoded representations. Subsequent works have sought to disentangle encoded representations to correspond to “semantic” attributes of audio such as prosody and linguistic content [17], [18], and to augment encoded token sequences with time-invariant global information to better capture qualities such as speaker identity [18], [22]. While the aforementioned codecs operate directly on waveform audio representations, recent works have found advantages in modeling magnitude and phase-magnitude spectrogram representations [19], [20]. In this work, we evaluate neural audio codecs from each of the above categories – RVQ waveform codecs, disentangled “semantic” codecs, and spectrogram codecs.

Generative Modeling with Codec Representations: While initial work on neural audio codecs focused on traditional lossy compression applications such as music and video streaming [3], [21], researchers have found that the encoded representations produced by these systems facilitate generative modeling of audio for tasks such as music generation [5], [8], text-to-speech synthesis [17], [23], sound effect generation [24], and speech enhancement [6]. These generative models are trained to predict token series using a language-modeling (categorical cross-entropy) objective in token space. In this work, we evaluate the performance of neural audio codecs on a simple conditional generation task (vocoding) in which tokens are predicted from a heavily compressed mel-frequency cepstrum coefficient (MFCC) representation.

Studies of Neural Audio Codec Representations: Recent works have evaluated the utility of neural audio codec representations for both discriminative and generative downstream tasks [25]–[28]. Others have focused on identifying specific properties within these representations and understanding their impact on downstream performance. Lemerrier et al. [29] find that the codebooks learned by Encodec [3] exhibit conditional dependence and show that downstream generative modeling performance can be improved by reducing this dependence via an additional training objective. Similarly, we evaluate one specific codec property (idempotence) and examine its relationship to downstream generative modeling performance.

Idempotence in Neural Image Codecs: Neural codec idempotence has been studied in the image modality. Zhang et al. [30] introduce the notion of “strong idempotence” – stability to re-encoding across different bandwidths – and evaluate both traditional and neural image codec models. Xu et al. [13] explore the relationship between perceptual image compression and codec idempotence, and show that high-quality compression can be achieved by pairing a neural codec with an unconditional generative model through an

“inversion” optimization. Li et al. [31] relax invertibility constraints on normalizing flow-based image compression to obtain a fully idempotent codec and improve the trade-off between compression ratio and perceptual transparency in near-idempotent variants. Kim et al. [32] evaluate the idempotence of neural image codecs and propose an auxiliary training objective to improve idempotence; we evaluate similar objectives in our experiments. In this work, we extend the study of neural codec idempotence to the audio modality.

III. MEASURING IDEMPOTENCE

We begin by evaluating the idempotence of state-of-the-art neural audio codecs in terms of the stability of both audio outputs and encoded representations under multiple rounds of encoding.

Codecs: We consider the waveform RVQ codecs Encodec (in the original configurations, as well as the 16kHz speech-domain variant trained by Peng et al.) [3], [7] and Descript Audio Codec (DAC) [16]; the “semantic” codecs SpeechTokenizer [17] and FACodec [18]; and the spectrogram-based spectral codecs of Langman et al. [19] and ESC [20]. We use the available pre-trained weights for all codecs, and evaluate each codec at the maximum supported bitrate. We note that these codecs differ substantially in native sample rate, compression ratio, training data, and architecture, making exact performance comparisons difficult; nevertheless, we find strong shared trends in the idempotence of audio outputs and encoded representations. Finally, as baselines we evaluate the signal-processing MP3 [14] and OPUS [15] codecs at low bitrates (see Table I).

Datasets While many codecs are trained for both speech and general-purpose audio compression [3], [16], we focus on speech data to facilitate comparisons. We evaluate codecs on the full-bandwidth speech datasets Expresso [33], containing spontaneous and read speech across four speakers; and VCTK [34], containing read speech across 109 speakers. For each dataset, we select 500 one-second excerpts for evaluation and resample audio to the native sample rate of each codec. For the signal-processing baselines, we select 100 excerpts, due to longer encoding times.

Metrics: To measure the stability of output audio under multiple rounds of encoding, we use the full-reference metrics *Scale-Invariant Signal-to-Distortion ratio (SI-SDR)* [35] and *Perceptual Evaluation of Speech Quality (PESQ)* [36]. PESQ estimates “perceptually-weighted” similarity between the original audio and encoded/decoded audio (up to 8kHz), while SI-SDR operates on the available bandwidth. Given the original audio input as a reference, codecs with high idempotence should produce reconstructions with high PESQ and SI-SDR scores even after multiple encoding iterations. To analyze drift in the internal representations of neural codecs after multiple encoding iterations, we compute the *match rate* between tokens at

each separate codebook level of the encoded representation after n and $(n + 1)$ encoding iterations. This value ranges from 100% when all tokens match to 0% when none do. While this is a coarse measure that does not account for codebook size or latent similarity, it allows us to observe the convergence of certain codecs towards “stable tokenizations” over multiple iterations. We also measure *codebook use* (the entropy of encoded token distributions as a percentage of the maximum achievable codebook entropy) to study whether codecs converge to small subsets of tokens after multiple encoding iterations. Higher entropy indicates broader codebook use.

In Table I, we compare the performance of our selected codecs in terms of PESQ scores after a single encoding (indicating perceptual transparency) and after 25 encodings (indicating idempotence). Values are grey for datasets seen during codec training. Where we evaluate multiple variants of a codec (e.g. at different bitrates) we separate entries with slashes. In all experiments we normalize codec outputs to match the input volume. In Figure 2 we select the variant of each neural codec with the highest demonstrated idempotence and plot audio quality, codebook use, and token match rate as a function of the number of encoding iterations. In Figure 3 we plot the token match rate between encodings of an input under small time shifts, indicating the sensitivity of the token representation to phase, and compute the correlation between average observed sensitivity and idempotence as measured by PESQ score after 25 encoding iterations.

The results of our experiments show that the idempotence of neural audio codecs varies greatly; DAC, ESC, and the Peng et al. variant of Encodec fare best, with DAC in particular demonstrating high audio quality even after 25 encoding iterations. Codec type alone (RVQ, semantic, or spectrogram) does not dictate performance. As can be seen in Figure 2, the behavior of token representations under multiple encodings also varies significantly across codecs. In codecs exhibiting the highest idempotence (DAC, ESC), codebook use remains virtually unchanged, while token match rate increases – indicating that even as these codecs converge to nearly fixed tokenizations, they do not collapse to a small subset of possible tokens. By contrast, Spectral-Codec and FACodec converge more slowly or not at all, while codebook use declines noticeably.

The observed differences in idempotence are not completely explained by the single-encoding perceptual transparency of codecs, as indicated by both Figure 2 and Table I. For instance, DAC 24kHz and DAC 16kHz have high perceptual transparency (as indicated by PESQ and SI-SDR scores after one encoding iteration) but exhibit very low idempotence after 25 iterations. Phase sensitivity seems to play a larger role: as shown in Figure 3, the sensitivity of token representations to phase shift positively correlates with idempotence. This is counterintuitive, as we might expect more idempotent codecs to be invariant to imperceptible perturbations such as phase shift. However, from our experiments it appears that fine-grained encoding of phase information is linked to the preservation of audio quality over multiple encoding iterations.

IV. ENCOURAGING IDEMPOTENCE

Given the wide range of idempotence observed across neural audio codecs, we explore whether it is possible to improve idempotence without sacrificing perceptual transparency or fundamentally altering the compression scheme. We focus our efforts on DAC, as it shows the highest overall idempotence of audio and encoded representations in our experiments.

For a neural codec f and input signal x , idempotence can be expressed as the relationship

$$f(x) = f(f(x))$$

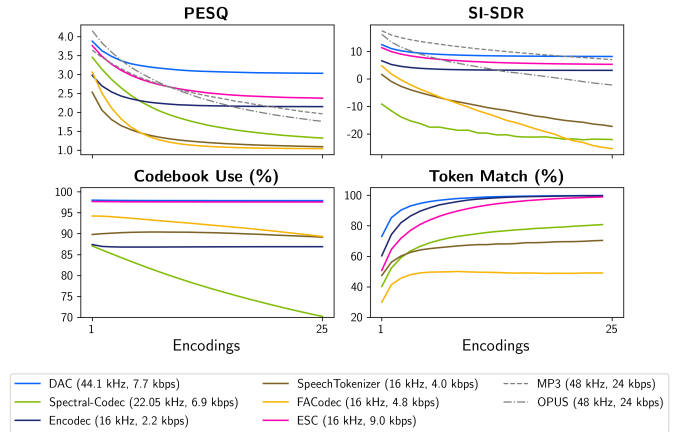


Fig. 2. **Audio and code stability over encodings:** audio quality (in PESQ and SI-SDR) of selected neural codecs by encoding iteration, average codebook use by encoding iteration, and match rate between tokens encoded at the n^{th} and $(n + 1)^{\text{th}}$ iterations.

A neural codec f typically consists of an encoder network f_e , decoder network f_d , and quantization module f_q such that its output can be expressed as $f(x) = f_d \circ f_q \circ f_e(x)$, where $f_q(x)$ is the quantized encoded representation. For DAC codec f and input x , we can decompose the compression and decompression process into individual steps. Taking $x' := f(x)$, we have:

$$x' = f_d \circ f_{dp} \circ f_{dq} \circ f_q \circ f_p \circ f_e(x) \quad (1)$$

where f_e is a neural network encoding waveform audio x into frame-level continuous latents; f_p is a linear projection into a low-dimension space for codebook lookup; f_q is a L_2 -normalized nearest-neighbors lookup over codebook vectors, returning codebook indices; f_{dq} indexes the codebook and to obtain codebook vectors; f_{dp} is a linear projection of the indexed codebook vectors back to the latent space; and f_d is a neural network that upsamples the resulting latent sequence to obtain waveform audio. Because the vector quantization operation $f_{dq} \circ f_q$ is non-differentiable with respect to the input, the straight-through gradient estimator is used during backpropagation. DAC is trained with a combination of adversarial and spectrogram losses to encourage perceptual transparency as well as the standard codebook and commitment losses [37]. For additional details we refer the reader to Kumar et al. [16].

Similar to Kim et al. [32], we encourage idempotent representations through the use of an additional regularizing objective during codec training. We explore three regularizing losses for enforcing idempotence at different stages of the compression process.

- 1) **Encoder idempotence**, the distance between continuous latents produced by the first and second encoding of a signal:

$$\mathcal{L}_{\text{enc}} = \|f_e(x') - f_e(x)\|_2$$

- 2) **Projected idempotence**, the distance between projected latents produced by the first and second encoding of a signal:

$$\mathcal{L}_{\text{proj}} = \|f_p \circ f_e(x') - f_p \circ f_e(x)\|_2$$

- 3) **Token idempotence**, the distance between projected latents from the second encoding and the codebook vectors selected in the first encoding:

$$\mathcal{L}_{\text{code}} = \|f_p \circ f_e(x') - f_{dq} \circ f_q \circ f_p \circ f_e(x)\|_2$$

All losses are averaged over the time dimension of the representation at which they are computed. We note that the encoder idempotence loss is most similar to the idempotence loss proposed by Kim et al.




Fig. 3. **Phase sensitivity correlates with idempotence:** For each selected codec, we compute the token match rate between encodings of audio inputs under time shifts between ± 2 ms. We then compute the correlation between average observed match rate and PESQ score after 25 encoding iterations. In general, codecs with higher phase sensitivity (as indicated by lower match rates) exhibit higher idempotence (as indicated by higher PESQ scores).

[32], which is computed on the quantized latents of the first encoding; we find this degrades audio quality significantly and thus use the unquantized latents instead.

To evaluate the effect of our proposed idempotence losses, we train a reproduction of DAC from scratch for 500k steps on a mixture of openly-licensed speech data (CommonVoice [38], LibriVox [39], VoxPopuli [40]) and proprietary licensed sound effects and music data. Our DAC operates at 48kHz with a hop length of 640 samples, resulting in a token rate of 75Hz and bitrate of 6.7kbps; all other aspects of the model and training are identical to that of Kumar et al. After pretraining we freeze the quantizer module, disable codebook dropout, introduce one of the idempotence losses described above weighted by a constant λ , and fine-tune for a further 50k steps. To select λ we sweep candidate values of $\{1, 10, 100, 1000\}$ and select the largest value for which the pretrained model’s audio quality is preserved (as measured by validation loss). We use $\lambda = 1$ for \mathcal{L}_{enc} , $\lambda = 10$ for \mathcal{L}_{proj} , and $\lambda = 100$ for \mathcal{L}_{code} . In Figure 1 we perform the multiple-encoding experiment from Section III with our fine-tuned DAC variants and find that the choice of loss has a significant effect on the stability of output audio and encoded representations. All idempotent fine-tuned variants exhibit improved stability in audio quality under re-encoding when compared to our DAC reproduction. The codebook-space objective \mathcal{L}_{code} performs best, resulting in superior audio quality to all evaluated codecs (including the original DAC model of Kumar et al.) after 25 encoding iterations. The projected objective \mathcal{L}_{proj} is less effective, followed by the latent objective \mathcal{L}_{enc} .

V. DOWNSTREAM GENERATIVE MODELING

While idempotence is in itself valuable for applications such as file compression, many neural audio codecs are explicitly designed to facilitate generative modeling on encoded representations. Accordingly, we evaluate the impact of our idempotent training configurations on a simple generative modeling task. Using token representations obtained from DAC, our DAC reproduction, and our fine-tuned variants, we train SoundStorm-like [23] masked generative transformer models to generate speech conditioned on 30-bin MFCC representations computed at the codec frame rate. This can be thought of as a vocoding task, albeit conditioned on a more compressed speech representation than the typical 80-120 band mel-spectrogram features [41], [42], making token prediction difficult. We apply a proprietary bandwidth extension algorithm to the LJSpeech dataset [43] to obtain roughly 24 hours of high-quality single-speaker read speech at 48kHz, allowing us to fully utilize the bandwidth of our DAC-based codecs. We train all models for 100k iterations on 5-

TABLE II
GENERATIVE MODELING ON CODEC TOKENS

Codec	SQUIM-MOS \uparrow	PESQ \uparrow	ASR-CER \downarrow
DAC (ours)	4.39	1.68	0.07
DAC (ours) + \mathcal{L}_{enc}	4.34	1.66	0.06
DAC (ours) + \mathcal{L}_{proj}	4.36	1.67	0.07
DAC (ours) + \mathcal{L}_{code}	4.26	1.65	0.06

second random excerpts at a batch size of 32 using the AdamW [44] optimizer with learning rate $1e-4$, and perform conditioning dropout with probability 0.2. We use small 43m-parameter transformer models consisting of 12 standard transformer encoder blocks with 8 attention heads each and apply rotary positional encodings in each block. During inference, we proceed from the lowest to highest codebook level following the “MaskGIT” strategy used by Yang et al. [6]; we use the confidence-based masking scheme from Garcia et al. [8] with mask temperature 10.5, sample with temperature 1.0 and nucleus 0.85 [45], and apply classifier-free guidance with weight 2.0 [46]. For each model we generate 100 samples using MFCC feature sequences from unseen recordings and compute the following metrics: *SQUIM Mean Opinion Score (SQUIM-MOS)* [47], a deep neural network-based estimate of speech quality that does not require precise alignment to the reference audio; *Character Error Rate (ASR-CER)* of a Wav2Vec2.0 [48] speech recognition system’s predicted transcripts on generated audio, which measures the intelligibility of generated speech; and PESQ. The results of our evaluations are presented in Table V. Due to the information bottleneck of the MFCC representation, generations do not match the precise spectral structure of the source audio, as indicated by low PESQ scores. However, we find the generations retain high audio quality and intelligibility, as indicated by SQUIM-MOS and ASR-CER scores. Generative models trained on codecs fine-tuned with idempotence objectives produce outputs of similar quality to a model trained on the base codec, indicating that the proposed idempotence losses do not have a strong negative effect on generative modeling.

VI. CONCLUSION

In this work we evaluate the idempotence of neural audio codecs and highlight possible connections to perceptual transparency, codebook use, and phase sensitivity. We then demonstrate that neural codec idempotence can be significantly increased via fine-tuning without sacrificing perceptual transparency or suitability for downstream generative modeling. This suggests that the idempotence of neural audio codecs can be improved without compromising desirable distributional properties of their encoded representations, allowing practical use in low-bitrate lossy file-compression scenarios without significant architectural or training modifications. Moreover, as codec-based generative models see wider use for audio editing tasks [6] and in iterative creative interfaces [49], idempotent fine-tuning may help to prevent the deterioration of audio quality under repeated applications of these models. We leave the study of the effects of codec idempotence on specific generative applications to future work. Future work may also explore incorporating idempotence objectives into codec training more fully; examining the impact of training data and architecture on idempotence; and applying idempotent codec architectures from the image modality such as invertible networks [31] to naturally invertible audio representations (e.g. STFT).

REFERENCES

- [1] J. Li, B. Li, and Y. Lu, "Deep contextual video compression," in *Neural Information Processing Systems*, 2021.
- [2] F. Mentzer, E. Agustsson, M. Tschannen, R. Timofte, and L. Van Gool, "Conditional probability models for deep image compression," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [3] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *arXiv preprint arXiv:2210.13438*, 2022.
- [4] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, D. Roblek, O. Teboul, D. Grangier, M. Tagliasacchi, and N. Zeghidour, "Audiolm: a language modeling approach to audio generation," *arXiv preprint arXiv:2209.03143*, 2023.
- [5] J. Copet, F. Kreuk, I. Gat, T. Remez, D. Kant, G. Synnaeve, Y. Adi, and A. Défossez, "Simple and controllable music generation," in *Neural Information Processing Systems*, 2024.
- [6] H. Yang, J. Su, M. Kim, and Z. Jin, "Genhancer: High-fidelity speech enhancement via generative modeling on discrete codec tokens," in *Interspeech*, 2024.
- [7] P. Peng, P.-Y. Huang, D. Li, A. Mohamed, and D. Harwath, "Voicecraft: Zero-shot speech editing and text-to-speech in the wild," *arXiv preprint arXiv:2403.16973v1*, 2024.
- [8] H. F. Garcia, P. Seetharaman, R. Kumar, and B. Pardo, "Vampnet: Music generation via masked acoustic token modeling," in *ISMIR*, 2023.
- [9] G. K. Wallace, "The jpeg still picture compression standard," *Communications of the ACM*, 1991.
- [10] AudioLabs Erlangen, "Tandem Coding," AES Coding Tutorial. [Online]. Available: <https://www.audiolabs-erlangen.de/content/resources/aesCodingTutorial/tandem.html>
- [11] F. Kurth, "An audio codec for multiple compression without loss of perceptual quality," in *Proceedings of the 17th International AES Conference on High Quality Audio Coding*, Florence, Italy, 1999.
- [12] S. Ritscher and U. Felderhoff, "Cascading of different audio codecs," in *Proceedings of the 100th AES Convention*, Copenhagen, Denmark, 1996, preprint 4174.
- [13] T. Xu, Z. Zhu, D. He, Y. Li, L. Guo, Y. Wang, Z. Wang, H. Qin, Y. Wang, J. Liu, and Y.-Q. Zhang, "Idempotence and perceptual image compression," in *ICLR*, 2024.
- [14] "Mp3 (mpeg layer iii audio encoding)," 1993. [Online]. Available: <https://www.iso.org/standard/22412.html>
- [15] J.-M. Valin, K. Vos, and T. B. Terriberry, "Definition of the opus audio codec," *IETF RFC 6716*, 2012. [Online]. Available: <https://tools.ietf.org/html/rfc6716>
- [16] R. Kumar, P. Seetharaman, I. K. Alejandro Luebs, and K. Kumar, "High-fidelity audio compression with improved rvqgan," in *NeurIPS*, 2023.
- [17] X. Zhang, D. Zhang, S. Li, Y. Zhou, and X. Qiu, "Speectokenizer: Unified speech tokenizer for speech large language models," in *ICLR*, 2024.
- [18] Z. Ju, Y. Wang, K. Shen, X. Tan, D. Xin, D. Yang, Y. Liu, Y. Leng, K. Song, S. Tang, Z. Wu, T. Qin, X.-Y. Li, W. Ye, S. Zhang, J. Bian, L. He, J. Li, and S. Zhao, "Naturalspeech 3: Zero-shot speech synthesis with factorized codec and diffusion models," *arXiv preprint arXiv:2403.03100*, 2024.
- [19] R. Langman, A. Jukić, K. Dhawan, N. R. Koluguri, and B. Ginsburg, "Spectral codecs: Spectrogram-based audio codecs for high quality speech synthesis," 2024.
- [20] Y. Gu and E. Diao, "Esc: Efficient speech coding with cross-scale residual vector quantized transformers," 2024.
- [21] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, "Soundstream: An end-to-end neural audio codec," *arXiv preprint arXiv:2107.03312*, 2021.
- [22] Y. Ren, T. Wang, J. Yi, L. Xu, J. Tao, C. Zhang, and J. Zhou, "Fewer-token neural speech codec with time-invariant codes," in *ICASSP*, 2024.
- [23] Z. Borsos, M. Sharifi, D. Vincent, E. Kharitonov, N. Zeghidour, and M. Tagliasacchi, "Soundstorm: Efficient parallel audio generation," *arXiv preprint arXiv:2305.09636*, 2023.
- [24] F. Kreuk, G. Synnaeve, A. Polyak, U. Singer, A. Défossez, J. Copet, D. Parikh, Y. Taigman, and Y. Adi, "Audiogen: Textually guided audio generation," in *ICLR*, 2023.
- [25] P. Mousavi, L. D. Libera, J. Duret, A. Ploujnikov, C. Subakan, and M. Ravanelli, "Dasb - discrete audio and speech benchmark," *arXiv preprint arXiv:2406.14294*, 2024.
- [26] H. Wu, H.-L. Chung, Y.-C. Lin, Y.-K. Wu, X. Chen, Y.-C. Pai, H.-H. Wang, K.-W. Chang, A. H. Liu, and H. yi Lee, "Codec-superb: An in-depth analysis of sound codec models," *arXiv preprint arXiv:2402.13071*, 2024.
- [27] H. Wu, X. Chen, Y.-C. Lin, K. wei Chang, H.-L. Chung, A. H. Liu, and H. yi Lee, "Towards audio language modeling – an overview," 2024.
- [28] H. Wu, X. Chen, Y.-C. Lin, K. Chang, J. Du, K.-H. Lu, A. H. Liu, H.-L. Chung, Y.-K. Wu, D. Yang, S. Liu, Y.-C. Wu, X. Tan, J. Glass, S. Watanabe, and H. yi Lee, "Codec-superb @ slt 2024: A lightweight benchmark for neural audio codec models," 2024.
- [29] J.-M. Lemerrier, S. Rouard, J. Copet, Y. Adi, and A. Défossez, "An independence-promoting loss for music generation with language models," in *ICML*, 2024.
- [30] Q. Zhang, T. Xu, Y. Li, and Y. Wang, "Evaluating strong idempotence of image codec," *arXiv preprint arXiv:2304.08269*, 2023.
- [31] Y. Li, T. Xu, Y. Wang, J. Liu, and Y.-Q. Zhang, "Idempotent learned image compression with right-inverse," in *NeurIPS*, 2024.
- [32] J.-H. Kim, S. Jang, J.-H. Choi, and J.-S. Lee, "Instability of successive deep image compression," in *ACM MultiMedia*, 2020.
- [33] T. A. Nguyen, W.-N. Hsu, A. D'Avirro, B. Shi, I. Gat, M. Fazel-Zarani, T. Remez, J. Copet, G. Synnaeve, M. Hassid, F. Kreuk, Y. Adi, and E. Dupoux, "Expresso: A benchmark and analysis of discrete expressive speech resynthesis," *arXiv preprint arXiv:2308.05725*, 2023.
- [34] J. Yamagishi, "English multi-speaker corpus for cstr voice cloning toolkit," 2012. [Online]. Available: <https://datashare.ed.ac.uk/handle/10283/2950>
- [35] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr – half-baked or well done?" in *ICASSP*, 2019.
- [36] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *ICASSP*, 2001.
- [37] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," in *NeurIPS*, 2017.
- [38] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *LREC*, 2020.
- [39] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An asr corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [40] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, 2021, pp. 993–1003. [Online]. Available: <https://aclanthology.org/2021.acl-long.80>
- [41] M. Morrison, R. Kumar, K. Kumar, P. Seetharaman, A. Courville, and Y. Bengio, "Chunked autoregressive gan for conditional waveform synthesis," in *ICLR*, 2022.
- [42] S. Lee, W. Ping, B. Ginsburg, B. Catanzaro, and S. Yoon, "Bigvgan: A universal neural vocoder with large-scale training," in *ICLR*, 2023.
- [43] K. Ito and L. Johnson, "The lj speech dataset." [Online]. Available: <https://keithito.com/LJ-Speech-Dataset/>
- [44] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [45] A. Holtzman, J. Buys, L. Du, M. Forbes, and Y. Choi, "The curious case of neural text degeneration," in *ICLR*, 2020.
- [46] J. Ho and T. Salimans, "Classifier-free diffusion guidance," in *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [47] A. Kumar, K. Tan, Z. Ni, P. Manocha, X. Zhang, E. Henderson, and B. Xu, "Torchaudio-squim: Reference-less speech quality and intelligibility measures in torchaudio," in *ICASSP*, 2023.
- [48] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Neural Information Processing Systems (NeurIPS)*, 2020.
- [49] H. F. Garcia and S. Moore, "Token telephone," in *New Interfaces for Musical Expression (NIME)*, 2024, art installation.