

Accelerating Codec-based Speech Synthesis with Multi-Token Prediction and Speculative Decoding

Tan Dat Nguyen¹, Ji-Hoon Kim¹, Jeongsoo Choi¹, Shukjae Choi², Jinseok Park², Younglo Lee², Joon Son Chung¹

¹Korea Advanced Institute of Science and Technology, South Korea, ²42dot Inc, South Korea

{tandat.kaist, jh.kim, jeongsoo.choi, joonson}@kaist.ac.kr

{shukjae.choi, jinseok.park, younglo.lee}@42dot.ai

Abstract—The goal of this paper is to accelerate codec-based speech synthesis systems with minimum sacrifice to speech quality. We propose an enhanced inference method that allows for flexible trade-offs between speed and quality during inference without requiring additional training. Our core idea is to predict multiple tokens per inference step of the AR module using multiple prediction heads, resulting in a linear reduction in synthesis time as the number of heads increases. Furthermore, we introduce a novel speculative decoding technique that utilises a Viterbi-based algorithm to select the optimal sequence of generated tokens at each decoding step. In our experiments, we demonstrate that the time required to predict each token is reduced by a factor of 4 to 5 compared to baseline models, with minimal quality trade-off or even improvement in terms of speech intelligibility. Audio samples are available at: mm.kaist.ac.kr/projects/mtp.

Index Terms—speech synthesis, speculative decoding

I. INTRODUCTION

Text-to-Speech (TTS) [1] has garnered significant attention within the research community. Amongst the various approaches, deep neural networks have shown considerable advancements in modeling natural speech [2]–[7]. Recent models based on this framework have shown the potential to generate fluent speech and replicate a speaker’s voice only using a few seconds of audio input [7]–[16]. Furthermore, LLMs also possess outstanding ability on several speech task such as recognizing emotions, identifying speakers, enhancing speech, or even understanding the complex context from background sounds [16]–[19]. This ability is important to create a unified model that can naturally communicate with humans, reflecting emotion, accent, and style [18].

The LLM-based TTS model exhibits outstanding generative ability due to the capacity of the Transformer, as described by the scaling law [20], the advanced development of audio tokenizers [15], [21]–[23], and the availability of large speech dataset [24]–[27]. This class of TTS models can be viewed as an extension of large language models, where the input speech is considered as a new language. The vocabulary of this new language is the set of discrete representations, also referred to as tokens, which are generated by encoder of audio codec. The TTS model then leverages its sequence generation ability to generate sequences of tokens that are transformed back to waveforms by the decodec module [7], [11] or the vocoder module [28]–[31]. Since speech itself also

has linguistic content in the form of a sequence, a promising approach is to harness the power of pretrained LLMs on text-only domains [17], [32], [33].

Although AR inference brings significant improvements to codec-based speech synthesis, however it substantially reduces the generation speed of TTS systems. This slowdown is a notable drawback compared to other types of speech synthesis models. For example, non-AR models [4], [5] generate an entire sentence in a single step, and models using progressive inference require only a few function evaluations (NFE) to produce high-quality speech output [6], [34], [35]. In contrast, models using AR inference with a large language model architecture must generate sequences step by step, with the number of steps being equal to the sequence length. Consequently, as the sequence length increases, the time required for generation also increase. This problem is compounded by the fact that the computational complexity of Transformer-based models increases quadratically with sequence length, leading to an increase in the total floating-point operations per sequence.

In this paper, we introduce a multi-token prediction method that significantly reduces the NFE while maintaining comparable quality. Building on the observation that consecutive speech tokens are often similar, our architecture is trained to predict multiple future tokens simultaneously, rather than just one at a time. Additionally, we propose a Viterbi-like approach to capture the statistical relationships between the predicted tokens. A key advantage of our method is that it allows users to explicitly control the quality-speed trade-off by adjusting the number of future tokens predicted per function evaluation, without the need for re-training or fine-tuning the model. To our knowledge, this is the first instance of Viterbi-based speculative decoding in this context.

II. RELATED WORKS

Neural Audio Codec: A codec model consists of an encoder, a decoder, and multiple codebooks that are quantized from the latent space of the encoder. These models are often built on the RVQ-GAN framework [21], [22] which provide flexible bitrate for the task of audio compression. Recently, several approaches have been proposed for designing an optimal discrete representation space. Early works focused on improving compression and reconstruction quality [7], [21], [22]. As codec-based TTS systems have demonstrated strong

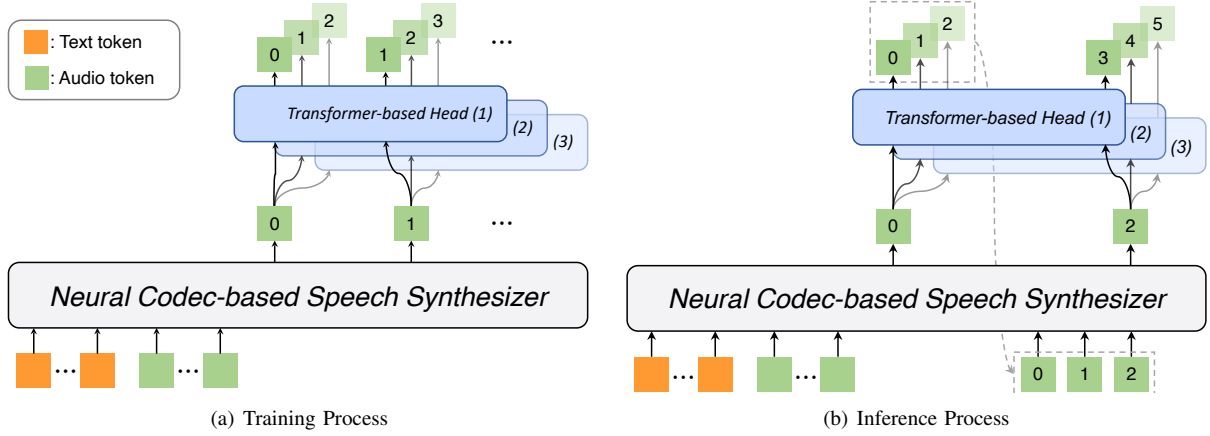


Fig. 1. AR module with $n = 3$ addition heads. Multiple heads are attached into current model that are simultaneously optimized to predict few future tokens given current state. Note that all the heads are trained once but can be flexibly ignored during inference for faster inference speed. During inference, few future tokens are predicted given current state per inference step. The generated tokens are concatenated into input sequence for the next inference step, respectively. Speculative decoding is also applied to enhance to quality of output tokens sequences.

performance, more recent efforts have focused on developing factorized and distinguishable features for each codebook [13], [15], [23], [36]–[39]. Moreover, current studies [36], [40]–[42] introduce codecs for mel-spectrograms, significantly increasing token compression rates and leading to faster AR module inference.

Multi-tokens prediction for AR module: Multi-token prediction is not a new concept and can be related to speculative decoding [43]–[45]. These methods use a smaller draft model to generate an initial token sequence, which is then refined by the larger or original model for a coherent continuation. Recent works explore this approach in large language models (LLMs) like incorporating techniques to reduce hardware memory requirements [46] or using tree search algorithms to enhance the quality of the output sequence [45]. To our knowledge, the most related work to ours is VALL-E 2 [8], but the multi-token prediction methods used in these two approaches are fundamentally different. Additionally, our work not only offers explicit, training-free control over the quality-speed trade-off but also improves the relationship between the predicted tokens through a Viterbi-like algorithm.

III. METHOD

A. Problem modeling

The popular codec-based speech synthesis model architecture comprises two main modules: an auto-regressive (AR) module and a non-auto-regressive (NAR) module, following the VALL-E framework [7]. Specifically, given a training dataset with pairs s, p , where s represents the speech signal and $p = \{p_1, p_2, \dots, p_T\}$ is the corresponding phoneme sequence of length T , the codec model compresses the speech signal into discrete tokens \mathbf{A} using eight quantizers: $\text{Codec}(s) = \mathbf{A}^{8 \times L} = \{a^1, a^2, \dots, a^8\}$, where each $a^i = \{a^i_1, a^i_2, \dots, a^i_L\}$. In models using the VALL-E architecture, the AR module is responsible for predicting a_1 , while the NAR module regresses the entire sequence $a^i_{i=2}^8$.

Specifically, for the discrete token sequence \mathbf{a}^1 , an auto-regressive Transformer Decoder-only θ_{AR} is trained to predict the next tokens on a text prompt p , and an acoustic prompt condition $\tilde{\mathbf{a}}^1$ extracted from the reference audio, and all previously predicted tokens.

$$p(\mathbf{a}^1 | p, \tilde{\mathbf{a}}^1; \theta_{AR}) = \prod_{t=0}^L p(a^1_{t+1} | \mathbf{a}^1_{<t}, \tilde{\mathbf{a}}^1, p; \theta_{AR}) \quad (1)$$

Since all data are discrete, we concatenate them all into a single sequence without specific tokens to distinguish them. For tokens of the remain layers $a^i_{i=2}^8$, we train a NAR LM θ_{NAR} to iteratively predict each token sequence. Each prediction process is conditioned on phoneme sequence p and generated token sequence of previous layers $\mathbf{a}^{<j}$.

$$p(\mathbf{a}^{2:8} | p, \tilde{\mathbf{a}}; \theta_{NAR}) = \prod_{j=2}^8 p(a^j | \mathbf{a}^{<j}, \tilde{\mathbf{a}}, p; \theta_{NAR}) \quad (2)$$

B. Multiple Tokens Prediction (MTP)

We take inspiration from [45], [46], which utilizes parallel decoding to accelerate the decoding process of LLM. We divide the AR module into two parts: first, encode all previous tokens in to latent space $p(\mathbf{z}_{1:t} | \mathbf{a}^1_{1:t}, \tilde{\mathbf{a}}, p; \theta_{AR})$, then, using this latent $\mathbf{z}_{1:t}$ to predict a few n consecutive future tokens $p(\mathbf{a}^1_{t+1:t+n} | \mathbf{z}_{1:t})$ by n distinguished heads. This idea leads to a new negative log-likelihood objective function of the whole AR module with multiple heads:

$$\mathcal{L} = - \sum_t \sum_{i=1}^n \log P(a^1_{t+i} | \mathbf{z}_{1:t}) P(\mathbf{z}_{1:t} | \mathbf{a}^1_{1:t}, \tilde{\mathbf{a}}, p) \quad (3)$$

$$\geq - \sum_t \log P(\mathbf{a}^1_{t+1:t+n} | \mathbf{z}_{1:t}) P(\mathbf{z}_{1:t} | \mathbf{a}^1_{1:t}, \tilde{\mathbf{a}}, p) \quad (4)$$

$$= - \sum_t \log P(\mathbf{a}^1_{t+1:t+n} | \mathbf{a}^1_{1:t}, \tilde{\mathbf{a}}, p) \quad (5)$$




Fig. 2. Viterbi-based Speculative Decoding is illustrated as follows: (1) Multiple prediction heads generate several distributions per timestep simultaneously. (2) To optimize memory and computational efficiency, the dimensions of the transition matrix and state probabilities are reduced by selecting only the necessary rows and columns. (3) The best sequence is determined using Speculative Decoding, as described in Algorithm 1. The transition matrix computation for LibriTTS is completed in just 3 minutes. Additionally, top_k sampling is employed to preserve diversity.

The equality in 4 holds when n consecutive tokens are independent given z_1 . This highlights the need for a tokenizer that supports such inference. To the best of our knowledge, SpeechTokenizer [15] one is the most suitable candidate, as its first-layer tokens encode only linguistic information, effectively serving as pseudo labels for phonemes. This makes future token prediction easier compared to Encodec [22], where linguistic and acoustic information are entangled at every quantizer layer. Furthermore, by applying n heads in parallel during inference, we reduce the time complexity from $\mathcal{O}(L)$ to $\mathcal{O}(\frac{L}{n} + \alpha)$, where α represents the additional computational overhead introduced by the extra heads.

C. Vitebi-based Speculative Decoding

Predicting multiple tokens $a_{t+1:t+n}^1$ at once, based on the context from $a_{\leq t}^1$, introduces a context mismatch among the predicted tokens. Specifically, except for a_{t+1}^1 , all tokens $a_{t+2:t+n}^1$ are predicted without full awareness of the preceding tokens, which can lead to incorrect predictions that accumulate over time and compromise the entire sequence. In this work, we propose a Viterbi-based speculative decoding technique to bridge the contextual gap between successive tokens predicted simultaneously. The key objective of this algorithm is to select the best possible sequence given top_k tokens of the n heads.

Let V represent the set of tokens from the first layer of the RVQ model. We assume each token depends only on its previous token, allowing us to model the token transitions using a Markov chain. Our goal is to model the likelihood of token a_i occurring immediately before token a_j , represented by the probability $Q(i, j)$, where $Q \in \mathbb{R}^{|V| \times |V|}$ is the transition matrix derived from the dataset. Additionally, we obtain $S_t(a_i)$ from the t heads of the multi-token prediction (MTP), representing the probability of a_i at time t . Therefore, the probability of transitioning from a_i at time $t - 1$ to a_j at time t is $Q(a_i, a_j) \cdot S_t(a_j)$.

Algorithm 1 Speculative Decoding

```

1: Initialization:
2: for each state  $a_i$  do
3:    $\delta_1(i) \leftarrow S_1(a_i)$ 
4:    $\psi_1(i) \leftarrow 0$ 
5: end for
6: Recursion:
7: for each time step  $t = 2, 3, \dots, T$  do
8:   for each state  $j$  do
9:      $\delta_t(j) \leftarrow \max_i \delta_{t-1}(i) \cdot Q(a_i, a_j) \cdot S_t(a_j)$ 
10:     $\psi_t(j) \leftarrow \arg \max_i [\delta_{t-1}(i) \cdot Q(a_i, a_j)]$ 
11:   end for
12: end for
13: Termination:
14:  $P^* \leftarrow \max_j \delta_T(j)$ 
15:  $a_T^* \leftarrow \arg \max_j \delta_T(j)$ 
16: Path Backtracking:
17: for each time step  $t = T - 1, \dots, 1$  do
18:    $a_t^* \leftarrow \psi_{t+1}(a_{t+1}^*)$ 
19: end for
20: Return:  $A^* = (a_1^*, a_2^*, \dots, a_T^*)$ 

```

To determine the optimal sequence of tokens A^* , we apply Viterbi-based speculative decoding, as shown in Fig. 2 and Algorithm 1. However, implementing this algorithm has a computational complexity of $\mathcal{O}(n * V^2)$, leading to high computational costs. To mitigate this, we reduce the dimensions of both the transition matrix and the state matrix by selecting only the relevant elements during each inference step. This reduces the size of the transition matrix from $\mathbb{R}^{|V| \times |V|}$ to $\mathbb{R}^{m \times m}$, where $m \leq k \cdot n$ is the total number of unique tokens from the top k highest-probability predictions across the n heads of the MTP. In practice, m is much smaller than V , resulting in a new computational complexity of $\mathcal{O}(n * m^2)$, which significantly reduces the overhead computation time for each MTP inference.

IV. EXPERIMENTS

A. Dataset

We conduct experiments on the LibriTTS [27] corpus, using the all train subsets for training, and 'test-clean' for evaluation, following [13]. We also use SpeechTokenizer to extract audio tokens all experiments. For fair comparison, all models are trained from scratch. Both AR and NAR modules are optimized separately for 20 epochs on a single A6000 GPU, with a learning rate of 0.05 and 200 warmup steps.

B. Comparison with baselines

We integrate our method into state-of-the-art models, VALL-E [7] and USLM [15], and evaluate their zero-shot performance using various metrics. For content consistency, we compute Word Error Rate (WER) and Character Error Rate (CER) with a Conformer-Transducer speech-to-text system¹. Speaker similarity is assessed via cosine similarity of x -vector embeddings between prompts and generated utterances. We also conduct SMOS and MOS surveys for speaker similarity

¹nvidia/teams/nemo/models/stt_en_conformer_transducer_large

TABLE I

EXPERIMENTAL RESULTS ON LIBRITTS DATASET. TPT DENOTES TIME PER TOKEN (MS/TOKEN) AND SPEEDUP REFERS TO THE EXTENT OF SPEED IMPROVEMENT COMPARED TO THE BASELINE. MOS IS PRESENTED WITH 95% CONFIDENCE INTERVAL. \uparrow : HIGHER IS BETTER, \downarrow : LOWER IS BETTER.

Model	TPT \downarrow	Speedup \uparrow	MOS \uparrow	SMOS \uparrow	UTMOS \uparrow	SIM \uparrow	WER(%) \downarrow	CER(%) \downarrow
Ground truth	–	–	4.77 \pm 0.09	3.53 \pm 0.29	4.14	–	2.9	0.8
VALL-E [7] + Ours	20.5 4.5	– \times 4.56	2.89\pm0.17 2.79 \pm 0.18	3.20 \pm0.20 2.79 \pm 0.20	3.56 3.53	79.6 79.2	12.8 12.0	8.1 8.3
USLM [15] + Ours	23.5 4.4	– \times 5.34	3.72 \pm 0.15 4.02\pm0.12	3.00 \pm 0.20 3.25\pm0.19	3.82 3.82	79.3 78.8	11.5 8.7	7.7 5.2





Fig. 3. TPT \downarrow over num. of headsFig. 4. WER \downarrow over num. of headsFig. 5. UTMOS \uparrow over num. of heads

Fig. 6. Speedup ratios of MTP

TABLE II

ABLATION STUDIES ON THE EFFECT OF top_k IN SPEED AND QUALITY OF MULTIPLE TOKEN PREDICTION

	$n = 8$			$n = 4$		
top_k	TPT \downarrow	UTMOS \uparrow	WER \downarrow	TPT \downarrow	UTMOS \uparrow	WER \downarrow
3	4.4	3.82	8.6	8.0	3.82	8.1
5	4.4	3.79	9.3	8.0	3.78	8.7
7	4.4	3.76	9.7	8.0	3.72	9.9
9	4.4	3.74	10.2	8.0	3.65	10.9
15	4.5	3.67	11.1	8.0	3.49	12.2
25	4.5	3.58	12.2	8.1	3.31	13.9

and audio quality (5-point scale), and use UTMOS [47] for MOS estimation. Generation efficiency is measured using time-per-token (TPT). All metrics, except the MOS and SMOS surveys (50 random samples per model), are computed on the test-clean set. Results are shown in Table I.

Compared to the baselines, multiple token prediction with 8 prediction heads and speculative decoding accelerates the inference process by 4 to 5 times while maintaining comparable speech quality and speaker similarity. Notably, content consistency improves, with WER decreasing in both USLM and VALL-E, and CER improving in USLM. This demonstrates the potential of multiple token predictions to enhance content accuracy while speeding up the AR process. Specifically, USLM with speculative decoding achieved the best WER of 8.7% and was the fastest among all models. The difference in inference time per token between VALL-E and USLM arises from the variance in the length of sequences generated by each model, which lead to variance of computing complexity. However, the speedup ratios are both theoretically and experimentally validated, ensuring that the performance improvements are consistent across different conditions.

We further assess the effectiveness of Viterbi-based speculative decoding by comparing quality and content consistency across different numbers of prediction heads. Fig. 3 shows that increasing the number of heads accelerates the AR process, with minimal impact on speed due to the reduced transition matrix dimensions, which shrink the search space for the Viterbi algorithm. The slight overhead from speculative decoding is highlighted in Fig. 6, comparing scenarios with Viterbi enabled (w/ Viterbi) or disabled (w/o Viterbi). To examine the trade-off between speed and quality, we evaluate WER, CER, and UTMOS for prediction heads ranging from 2 to 8. Without Viterbi-based decoding, content consistency remains acceptable (WER) with a few heads, but UTMOS drops rapidly. With more than 4 heads, both speech quality and content accuracy degrade significantly because of accuracy reduction in far future heads. However, speculative decoding maintains consistent synthesis performance across different numbers of heads, as shown in Figs. 4 and 5. Notably, higher top_k expands the Viterbi algorithm’s search space with trivial additional time. According to Table II, although higher top_k increases diversity in generated speech, it tends to lower overall quality in terms of WER and UTMOS.

V. CONCLUSION

In this work, we propose an inference method that combines multi-token prediction with Viterbi-based speculative decoding, significantly accelerating the AR module during inference. Our approach not only achieves faster inference but also maintains, or even surpasses, baseline performance in terms of generation quality. Additionally, the model is trained only once, yet offers explicit control over the trade-off between quality and speed. However, exploring the compatibility of more tokenizers with Viterbi-based speculative decoding is a promising direction for future work.

REFERENCES

- [1] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, “A survey on neural speech synthesis,” *arXiv:2106.15561*, 2021.
- [2] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *SSW*, 2016.
- [3] R. J. Skerry-Ryan, E. Battenberg, Y. Xiao, Y. Wang, D. Stanton, J. Shor, R. J. Weiss, R. Clark, and R. A. Saurous, “Towards end-to-end prosody transfer for expressive speech synthesis with tacotron,” in *Proc. ICML*, 2018.
- [4] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, “FastSpeech: Fast, robust and controllable text to speech,” in *NeurIPS*, 2019.
- [5] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *Proc. ICML*, 2021.
- [6] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov, “GradTTS: A diffusion probabilistic model for text-to-speech,” in *Proc. ICML*, 2021.
- [7] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li *et al.*, “Neural codec language models are zero-shot text to speech synthesizers,” *arXiv:2301.02111*, 2023.
- [8] S. Chen, S. Liu, L. Zhou, Y. Liu, X. Tan, J. Li, S. Zhao, Y. Qian, and F. Wei, “Vall-e 2: Neural codec language models are human parity zero-shot text to speech synthesizers,” *arXiv:2406.05370*, 2024.
- [9] C. Du, Y. Guo, H. Wang, Y. Yang, Z. Niu, S. Wang, H. Zhang, X. Chen, and K. Yu, “Vall-t: Decoder-only generative transducer for robust and decoding-controllable text-to-speech,” *arXiv:2401.14321*, 2024.
- [10] B. Han, L. Zhou, S. Liu, S. Chen, L. Meng, Y. Qian, Y. Liu, S. Zhao, J. Li, and F. Wei, “Vall-e r: Robust and efficient zero-shot text-to-speech synthesis via monotonic alignment,” *arXiv:2406.07855*, 2024.
- [11] 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:2403.16973*, 2024.
- [12] Y. Song, Z. Chen, X. Wang, Z. Ma, and X. Chen, “Ella-v: Stable neural codec language modeling with alignment-guided sequence reordering,” *arXiv:2401.07333*, 2024.
- [13] Z. Du, Q. Chen, S. Zhang, K. Hu, H. Lu, Y. Yang, H. Hu, S. Zheng, Y. Gu, Z. Ma *et al.*, “Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens,” *arXiv:2407.05407*, 2024.
- [14] Z. Borsos, M. Sharifi, D. Vincent, E. Kharitonov, N. Zeghidour, and M. Tagliasacchi, “Soundstorm: Efficient parallel audio generation,” *arXiv:2305.09636*, 2023.
- [15] X. Zhang, D. Zhang, S. Li, Y. Zhou, and X. Qiu, “SpeechTokenizer: Unified speech tokenizer for speech large language models,” *arXiv:2308.16692*, 2023.
- [16] D. Yang, J. Tian, X. Tan, R. Huang, S. Liu, X. Chang, J. Shi, S. Zhao, J. Bian, X. Wu *et al.*, “Uniaudio: An audio foundation model toward universal audio generation,” in *Proc. ICML*, 2024.
- [17] E. Nachmani, A. Levkovich, R. Hirsch, J. Salazar, C. Asawaroengchai, S. Mariooryad, E. Rivlin, R. J. Skerry-Ryan, and M. T. Ramanovich, “Spoken question answering and speech continuation using spectrogram-powered LLM,” in *Proc. ICLR*, 2024.
- [18] D. Zhang, S. Li, X. Zhang, J. Zhan, P. Wang, Y. Zhou, and X. Qiu, “Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities,” in *Proc. EMNLP*, 2023.
- [19] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. R. Glass, “Listen, think, and understand,” in *Proc. ICLR*, 2024.
- [20] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei, “Scaling laws for neural language models,” *arXiv:2001.08361*, 2020.
- [21] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, “Soundstream: An end-to-end neural audio codec,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2022.
- [22] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, “High fidelity neural audio compression,” *Transactions on Machine Learning Research*, 2023.
- [23] Z. Ju, Y. Wang, K. Shen, X. Tan, D. Xin, D. Yang, E. Liu, Y. Leng, K. Song, S. Tang *et al.*, “Naturspeech 3: Zero-shot speech synthesis with factorized codec and diffusion models,” in *ICML*, 2024.
- [24] A. Nagrani, J. S. Chung, and A. Zisserman, “Voxceleb: a large-scale speaker identification dataset,” in *Proc. Interspeech*, 2017.
- [25] W. Kang, X. Yang, Z. Yao, F. Kuang, Y. Yang, L. Guo, L. Lin, and D. Povey, “Libriheavy: a 50,000 hours asr corpus with punctuation casing and context,” 2023.
- [26] 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.
- [27] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, “LibriTTS: A corpus derived from librispeech for text-to-speech,” in *Proc. Interspeech*, 2019.
- [28] H. Siuzdak, “Vocos: Closing the gap between time-domain and fourier-based neural vocoders for high-quality audio synthesis,” in *Proc. ICLR*, 2024.
- [29] M. Łajszczak, G. Cămbăra, Y. Li, F. Beyhan, A. van Korfelaar, F. Yang, A. Joly, Á. Martín-Cortinas, A. Abbas, A. Michalski *et al.*, “Base tts: Lessons from building a billion-parameter text-to-speech model on 100k hours of data,” *arXiv:2402.08093*, 2024.
- [30] J.-H. Kim, S.-H. Lee, J.-H. Lee, and S.-W. Lee, “Fre-gan: Adversarial frequency-consistent audio synthesis,” in *Interspeech 2021*, 2021.
- [31] T. D. Nguyen, J. Kim, Y. Jang, J. Kim, and J. S. Chung, “Fregrad: Lightweight and fast frequency-aware diffusion vocoder,” in *ICASSP*, 2024.
- [32] P. K. Rubenstein, C. Asawaroengchai, D. D. Nguyen, A. Bapna, Z. Borsos, F. d. C. Quirry, P. Chen, D. E. Badawy, W. Han, E. Kharitonov *et al.*, “Audiotape: A large language model that can speak and listen,” *arXiv:2306.12925*, 2023.
- [33] D. Yang, H. Guo, Y. Wang, R. Huang, X. Li, X. Tan, X. Wu, and H. Meng, “Uniaudio 1.5: Large language model-driven audio codec is a few-shot audio task learner,” *arXiv:2406.10056*, 2024.
- [34] R. Huang, M. W. Y. Lam, J. Wang, D. Su, D. Yu, Y. Ren, and Z. Zhao, “Fastdiff: A fast conditional diffusion model for high-quality speech synthesis,” in *Proc. IJCAI*, 2022.
- [35] Z. Ye, W. Xue, X. Tan, J. Chen, Q. Liu, and Y. Guo, “Comospeech: One-step speech and singing voice synthesis via consistency model,” in *Proc. ACM MM*, 2023.
- [36] R. Langman, A. Jukić, K. Dhawan, N. R. Koluguri, and B. Ginsburg, “Spectral codecs: Spectrogram-based audio codecs for high quality speech synthesis,” *arXiv:2406.05298*, 2024.
- [37] Z. Ye, P. Sun, J. Lei, H. Lin, X. Tan, Z. Dai, Q. Kong, J. Chen, J. Pan, Q. Liu, Y. Guo, and W. Xue, “Codec does matter: Exploring the semantic shortcoming of codec for audio language model,” *arXiv preprint arXiv:2408.17175*, 2024.
- [38] H. Liu, X. Xu, Y. Yuan, M. Wu, W. Wang, and M. D. Plumbley, “Semanticcodec: An ultra low bitrate semantic audio codec for general sound,” *arXiv:2405.00233*, 2024.
- [39] S. Ji, Z. Jiang, X. Cheng, Y. Chen, M. Fang, J. Zuo, Q. Yang, R. Li, Z. Zhang, X. Yang *et al.*, “Wavtokenizer: an efficient acoustic discrete codec tokenizer for audio language modeling,” *arXiv:2408.16532*, 2024.
- [40] J. Kim, K. Lee, S. Chung, and J. Cho, “Clam-tts: Improving neural codec language model for zero-shot text-to-speech,” in *ICLR*, 2024.
- [41] H. Bai, T. Likhomanenko, R. Zhang, Z. Gu, Z. Aldeneh, and N. Jaitly, “dmel: Speech tokenization made simple,” *arXiv:2407.15835*, 2024.
- [42] Y. Takida, Y. Ikemiya, T. Shibuya, K. Shimada, W. Choi, C.-H. Lai, N. Murata, T. Uesaka, K. Uchida, W.-H. Liao, and Y. Mitsufuji, “Hqvae: Hierarchical discrete representation learning with variational bayes,” *Transactions on Machine Learning Research*, 2024.
- [43] Y. Leviathan, M. Kalman, and Y. Matias, “Fast inference from transformers via speculative decoding,” in *Proc. ICML*, 2023.
- [44] C. Chen, S. Borgeaud, G. Irving, J.-B. Lespiau, L. Sifre, and J. Jumper, “Accelerating large language model decoding with speculative sampling,” *arXiv:2302.01318*, 2023.
- [45] T. Cai, Y. Li, Z. Geng, H. Peng, J. D. Lee, D. Chen, and T. Dao, “Medusa: Simple llm inference acceleration framework with multiple decoding heads,” in *Proc. ICML*, 2024.
- [46] F. Gloeckle, B. Y. Idrissi, B. Rozière, D. Lopez-Paz, and G. Synnaeve, “Better & faster large language models via multi-token prediction,” in *Proc. ICML*, 2024.
- [47] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, “UTMOS: utokyo-sarulab system for voicemos challenge 2022,” in *Proc. Interspeech*, 2022.