



Accelerating Autoregressive Speech Synthesis Inference With Speech Speculative Decoding

Zijian Lin¹, Yang Zhang², Yougen Yuan², Yuming Yan², Jinjiang Liu¹, Zhiyong Wu^{1,*}, Pengfei Hu², Qun Yu²

¹Shenzhen International Graduate School, Tsinghua University, China

²Platform and Content Group, Tencent, China

linzj24@mails.tsinghua.edu.cn, autozhang@tencent.com, zywu@sz.tsinghua.edu.cn

Abstract

Modern autoregressive speech synthesis models leveraging language models have demonstrated remarkable performance. However, the sequential nature of next token prediction in these models leads to significant latency, hindering their deployment in scenarios where inference speed is critical. In this work, we propose Speech Speculative Decoding (SSD), a novel framework for autoregressive speech synthesis acceleration. Specifically, our method employs a lightweight draft model to generate candidate token sequences, which are subsequently verified in parallel by the target model using the proposed SSD framework. Experimental results demonstrate that SSD achieves a significant speedup of 1.4x compared with conventional autoregressive decoding, while maintaining high fidelity and naturalness. Subjective evaluations further validate the effectiveness of SSD in preserving the perceptual quality of the target model while accelerating inference.¹

Index Terms: speech synthesis, speculative decoding, autoregressive model

1. Introduction

Modern autoregressive speech synthesis systems [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11] often employ four core components to achieve human-like speech generation, as shown in Figure 1a. The text encoder tokenizes input text into fine-grained sequences, and in some cases, further transforms them into high-dimensional semantic representations. The speech encoder discretizes continuous audio into semantic speech token sequences through vector quantization. The autoregressive language model (LM) then processes these tokens, capturing long-range temporal dependencies and contextual patterns to predict sequential subsequent acoustic units. Finally, the speech decoder reconstructs these discrete speech tokens into high-fidelity waveform signals. This architecture enables precise prosodic control and exceptional coherence in synthesized speech, setting new benchmarks for naturalness through the powerful contextual modeling capabilities of the language model.

Despite the great success, Autoregressive (AR) models generate speech strictly sequentially, leading to computationally intensive and time-consuming inference. This inefficiency is caused by two key factors: first, speech token sequences are inherently long, as each token typically represents a short segment of speech (e.g., 25 ms or 50 ms) [12, 13, 14], resulting in significantly longer sequences compared to text tokens; and second, the increasing parameter size of neural networks, driven by the pursuit of higher model performance, further amplifies computational demands. The inference time scales quadratically with

*Corresponding author.

¹Speech samples: <https://thuhcsi.github.io/interspeech2025-SSD/>

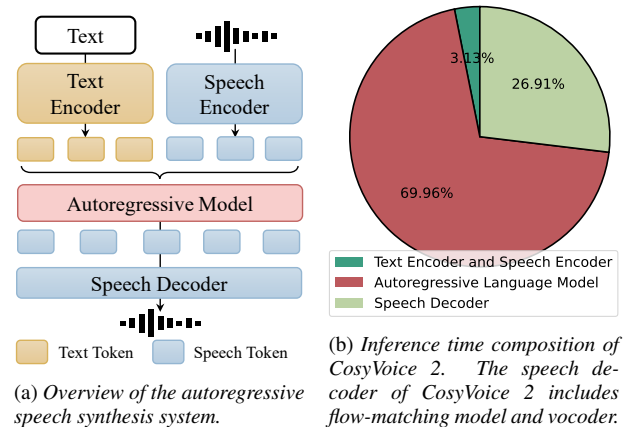


Figure 1: Visualization of autoregressive speech synthesis system and inference time composition of CosyVoice 2.

sequence length [15], while the computational complexity of the Transformer backbone grows quadratically, further exacerbating inefficiency. For instance, in CosyVoice 2 [5], a state-of-the-art LM-based speech synthesis model, the autoregressive inference process of the LM takes up approximately 70% of the total synthesis time, as illustrated in Figure 1b. This bottleneck has made inference acceleration a critical research focus in AR speech synthesis, driving exploration of non-autoregressive (NAR) alternatives [16, 17] such as diffusion models [18, 19] or flow matching [20, 21], which enable parallel generation and significantly improve inference speed.

Speculative decoding [22, 23] accelerates text-based large language models (LLMs) by employing a small draft model to autoregressively generate candidate token sequences, which are then verified in parallel by a larger target model. Inspired by these approaches, there have been notable attempts [24, 25] to apply speculative decoding to accelerate autoregressive speech synthesis. However, they integrate multiple draft heads directly into the target model and tend to require fine-tuning the target model itself to enhance draft head quality. In contrast, we introduce Speech Speculative Decoding (SSD) with an independent draft architecture that maintains high-quality predictions without requiring target model fine-tuning. Moreover, unlike text tokens, which often have clear one-to-one mappings with words or phonemes, speech tokens exhibit diverse and variable mappings, as audio segments that sound nearly identical can have different token-level representations. This variability stems from the continuous and dynamic nature of speech signals. Traditional speculative decoding [22, 23], designed for text, may result in suboptimal efficiency for speech synthesis, as their acceptance criteria fail to account for these nuanced relationships. To address this, we propose an improved acceptance

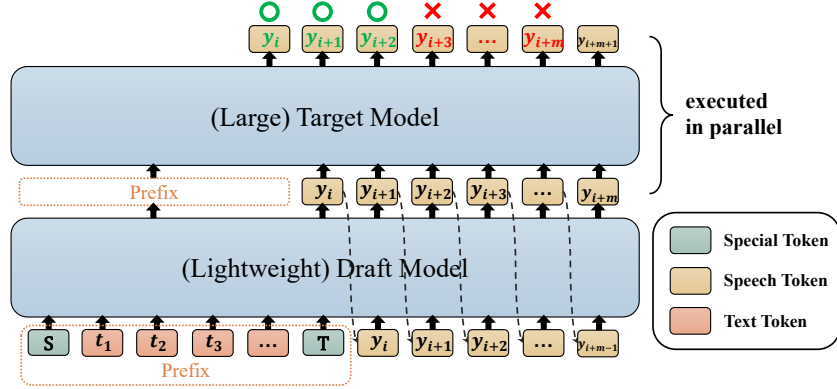


Figure 2: *The Inference Architecture of proposed Speech Speculative Decoding (SSD) framework. The dashed line depicts autoregressive decoding in the draft model, followed by parallel verification by the target model. The prefix comprises all prompt tokens and previously generated tokens during inference. Green indicates accepted tokens, while red represents rejected tokens.*

criterion tailored to speech tokens, enabling more efficient verification of candidate sequences generated by the draft model. The adaptation in SSD finally achieves a significant speedup of $1.4\times$ compared to conventional autoregressive decoding, while maintaining high fidelity and naturalness.

Our contributions can be summarized as follows. First, we present a novel SSD method targeting speech synthesis acceleration through two key mechanisms: a draft model to generate candidate token sequence, and a target model for parallel verification. Second, we design an efficient way to construct the lightweight draft model for alignment. Third, a tolerance factor is used to influence the acceptance rate of generated tokens, thus balancing the speech synthesis effect and inference speed.

2. Methodology

In this section, we demonstrate the proposed SSD acceleration framework, which comprises two core components: a target model (the large model requiring acceleration) and a draft model (a lightweight auxiliary model). As illustrated in Figure 2, the process unfolds in three phases: 1) Autoregressively generating m candidate tokens using the draft model, 2) Parallel verification of these tokens in the target model via SSD, 3) Applying a speculative decoding-based acceptance strategy to validate the token sequences.

2.1. Autoregressive Language Model for Speech Synthesis

In an autoregressive speech synthesis system, a language model is applied to predict discrete tokens of speech. Our approach builds upon the open-source CosyVoice framework [4, 5] for autoregressive speech synthesis, which uses language model to predict single stream of discrete speech tokens and then uses flow-matching model as the speech decoder.

Given a text-speech pair $\{T, A\}$, where T denotes text sequence and A represents the speech signal, text T is tokenized into a more fine-grained sequence $T = \{t_1, t_2, \dots, t_N\}$ of length N via some methods such as grapheme-to-phoneme (g2p) transformation or BPE-based text tokenizer, and the speech encoder discretizes the continuous speech signal A into a sequence of discrete tokens $Y = \{y_1, y_2, \dots, y_L\}$ of length L . Then the language model predicts Y autoregressively taken text T as input. To input all the information into the language model, text tokens and speech tokens are typically concatenated into a single sequence via some special tokens and then uniformly fed into the model.

Therefore, in the setting of the zero-shot TTS, the language

model θ_{LM} learns to model the distribution of the target speech token sequence Y_{tgt} :

$$P(Y_{tgt}|T_{tgt}, Y_{pt}, T_{pt}; \theta_{LM}) = \prod_{t=1}^L P(Y_{tgt}^t | Y_{tgt}^{<t}, T_{tgt}, Y_{pt}, T_{pt}; \theta_{LM}), \quad (1)$$

where Y_{pt} denotes the audio prompt, T_{pt} represents the transcription of the audio prompt, and T_{tgt} means the target text. Traditional autoregressive models suffer from speed limitations due to their token-by-token serial generation process. This sequential dependency significantly reduces computational efficiency and brings to high latency.

Algorithm 1: SSD for Speech Synthesis

Input: Prompt T_{in} , target model M_q , draft model M_p , draft length L_d , target length L_t , tolerance factor β

Output: Generated speech signal

Initialize $L_{acc} \leftarrow 0$

while $L_{acc} < L_t$ **do**

$p_1^d, \dots, p_{L_d}^d \leftarrow M_p(T_{in})$

$x_1, \dots, x_{L_d} \sim p_1^d, \dots, p_{L_d}^d$

$q_1^s, \dots, q_{L_d+1}^s \leftarrow M_q(T_{in}, x_1, \dots, x_{L_d})$

for $i = 1$ **to** L_d **do**

Sample $r_i \sim U[0, 1]$

if $r_i < \min(1, q_i^s / p_i^d) + \beta$ **then**

Append x_i to T_{in}

else

Resample $x'_i \sim \text{normlize}(\max(0, q_i^s - p_i^d))$

Append x'_i to T_{in}

break

end

end

if All tokens accepted **then**

Append $x_{L_d+1} \sim q_{L_d+1}^s$ to T_{in}

end

Update L_{acc} based on accepted tokens

end

2.2. Speech Speculative Decoding

As shown in Algorithm 1, speculative decoding accelerates autoregressive generation by leveraging a draft model M_p to propose candidate tokens and a target model M_q to verify them. In speech synthesis, we adapt this framework to address the inherent ambiguity of speech tokens, which lack explicit alignment with linguistic units. To enhance the flexibility of speculative decoding for speech synthesis without compromising output quality, we introduce a tolerance factor into the criteria.

Given an input prompt T_{in} in concatenated by text and speech tokens, the draft model first generates L_d candidate tokens autoregressively:

$$p_1^d, \dots, p_{L_d}^d = M_p(T_{in}), \quad (2)$$

$$x_1, \dots, x_{L_d} \sim p_1^d, \dots, p_{L_d}^d. \quad (3)$$

where x_i is a draft token, p_i^d is the probability distribution of the draft model, and $i \in [1, L_d]$.

After we obtain the draft tokens, the target model M_q then processes all draft tokens in parallel to produce verification distributions q^s :

$$q_1^s, \dots, q_{L_d+1}^s = M_q(T_{in}, x_1, \dots, x_{L_d}). \quad (4)$$

Unlike text generation where tokens have discrete semantics and one-to-one mappings with words or phonemes, speech tokens exhibit continuous and overlapping characteristics. Strict acceptance criteria in standard speculative decoding often lead to unnecessary rejections, however minor distribution discrepancies may not perceptually affect speech quality. Therefore, we introduce a tolerance factor $\beta \geq 0$ that relaxes the acceptance threshold:

$$\text{Accept } x_i \text{ if } r_i < \min(1, q_i^s/p_i^d) + \beta, \quad (5)$$

where $r_i \sim U[0, 1]$. This modification allows acoustically plausible variations to be accepted, increasing decoding throughput without compromising perceptual quality.

When rejection occurs at position i , we resample token x_i' from the adjusted distribution:

$$\text{Resample } x_i' \sim p_{new} = \text{normlize}(\max(0, q_i^s - p_i^d)), \quad (6)$$

ensuring the corrected token follows the target model’s distribution. If all L_d tokens are accepted, we append an additional token sampled from $q_{L_d+1}^s$.

2.3. Lightweight Draft Model Construction

Modern autoregressive speech synthesis systems rely on a single large target model, posing a challenge for speculative decoding: constructing a lightweight draft model that aligns with the vocabulary of target model. To address this, we design the draft model as a shallower variant of the target architecture, sharing its core structure but with fewer Transformer layers. Crucially, the draft model is initialized using the pre-trained upper layers of the target model (e.g., layers 19–24 in a 24-layer setup), enabling immediate compatibility with its linguistic representations while avoiding cold-start training.

To enhance coordination between the models without costly full fine-tuning, we employ a layer-specific adaptation strategy inspired by adapter tuning [26, 27]. During training, only lower layers (e.g., layers 1–2) of the draft model and classification head are updated, while the initialized upper layers and embeddings remain frozen. This approach preserves high-level semantic knowledge of the target model while allowing the draft model to adapt to speech-specific patterns through tunable lower layers. Unlike traditional adapter methods that insert new modules, our strategy achieves parameter efficiency through selective layer inheritance and hierarchical fine-tuning. To keep the training stage simple and efficient, we only use the cross entropy loss to fine-tune the draft model:

$$L_{draft} = -\frac{1}{L_{tgt}} \sum_{l=1}^{L_{tgt}} \log(M_p(u_{<l})) \quad (7)$$

where $U_{tgt} = \{u_1, \dots, u_{L_{tgt}}\}$ is the target sequence, $u_{<l} = \{u_1, \dots, u_{l-1}\}$, and L_{tgt} is the length of the target sequence.

3. Experiments

3.1. Implementation Details

We use CosyVoice 2 [5] as the target model, which is a highly efficient TTS model consisting of 24 layers of Transformer adapted from Qwen2.5 (0.5B) [28]. Based on this model, we design a draft model consisting of 8 layers of Transformer. We follow the setting of CosyVoice 2: each Transformer layer typically employs a multi-head attention mechanism with the number of heads set to 14, and the feed-forward layer in each Transformer layer has a hidden dimension of 896. We train the draft model on LibriTTS [29] while keeping the parameters of CosyVoice 2 frozen. To accelerate the training process, we load the pre-trained parameters of CosyVoice 2 to initialize the model. For the Transformer layers, we utilize the parameters of the lower two layers and the upper six layers of CosyVoice 2. During the training of the draft model, we only train the first two layers of the Transformer and the classification layer. The length of draft tokens as L_d is set to 3 and the tolerance factor β is set to 0.4. We conduct all the experiments on NVIDIA A100-SXM4-40GB GPU.

3.2. Dataset

LibriTTS, introduced by [29], is a well-known multi-speaker English speech dataset with transcriptions. It encompasses a substantial amount of speech, sourced from 2,456 English speakers, 585 hours in total.

For data partitioning, the subsets “train-clean-100”, “train-clean-360”, and “train-other-500” are grouped together to form the training data. This combined training set is approximately 580 hours long and is contributed by 2,306 speakers. And we merge the subsets “dev-clean” and “dev-other” for model selection, based on token prediction accuracy of the speech token sequences on this set.

Following [30], we evaluate SSD on the subset of LibriTTS. This testset, which covers 37 speakers, contains 500 utterances selected from the “test-clean” subset of LibriTTS. Its purpose is to evaluate the zero-shot adaptation ability for entirely new and unseen speakers. Moreover, for each speaker in the testset, there is a short voice prompt lasting approximately 3 seconds.

3.3. Evaluation Metrics

Referring to [5], we take word error rate (WER), speaker similarity (SS), and NMOS score as the objective metrics for a comprehensive evaluation of SSD. The Word Error Rate can measure content consistency. We use Whisper-large V3 as the Automatic Speech Recognition (ASR) model to calculate the WER. Regarding the Speaker Similarity, we employ the ERes2Net model to extract the speaker embedding vectors of the prompt speech and the generated one, and adopt their raw cosine similarity as the speaker similarity. The NMOS score is a widely-adopted metric for evaluating the objective quality of speech. We compute LM-RTF to assess the efficiency of the LM, which is defined as the ratio of the time taken by the LM to generate the speech to the duration of the synthesized speech.

For subjective evaluation, we conduct two Mean Opinion Score (MOS) tests on a 5-point scale with 1-point intervals (1: lowest, 5: highest), involving 30 raters each evaluating 10 randomized samples. First, we focus on the similarity (SIM-MOS), which provides insights into how closely the synthetic voice mimics the unique vocal characteristics of the target speaker. Then, we evaluate naturalness of the speech (NAT-MOS), meaning how close the synthesized voice is to a real-

human voice and detecting any signs of artificial synthesis.

4. Results and Analysis

We apply SSD to CosyVoice 2 and compare the result with draft model and CosyVoice 2 using autoregressive decoding to synthesize speech.

4.1. Objective Evaluation

As shown in Table 1, our Speech Speculative Decoding (SSD) framework achieves 1.4 \times faster synthesis than the conventional autoregressive decoding used by CosyVoice 2. Notably, this acceleration is attained despite SSD being trained on a small-scale open-source dataset (approximately 1/300 the size of CosyVoice 2’s proprietary industrial dataset [5]), highlighting its robust compatibility with large-scale systems. While the WER of our system shows a slight increase compared to CosyVoice 2, the negligible differences in NMOS and SS, coupled with the performance in SIM-MOS and NAT-MOS as discussed in Section 4.2, suggest that the WER discrepancy stems primarily from extrapolation generalization limitations due to limited training data.

Critically, despite that the draft model achieves the worst WER (16.13%), SSD maintains robust synthesis quality through its verification mechanism. This decoupling between draft model performance and final output quality underscores robustness of SSD: even with a suboptimal draft model, the framework maintains strong synthesis quality while delivering significant acceleration. These results collectively validate the effectiveness of SSD as a promising solution for autoregressive speech synthesis acceleration.

Table 1: Objective evaluation result of SSD with $\beta = 0.4$.

Model	WER (%) \downarrow	NMOS \uparrow	SS \uparrow	LM-RTF \downarrow
Draft Model	16.13	3.95	0.61	0.222
CosyVoice 2	3.67	3.96	0.62	0.504
w. SSD	5.70	3.94	0.63	0.360

4.2. Subjective Evaluation

Results in Table 2 demonstrate that SSD achieves comparable performance to the original CosyVoice 2 using autoregressive decoding in terms of both similarity (SIM-MOS) and naturalness (NAT-MOS). While the draft model alone achieves the lowest scores, highlighting its limitations in independent high-quality speech generation, SSD bridges this gap by leveraging the target model for verification. These results confirm that SSD effectively preserves the synthesis quality of the target model while accelerating inference, validating it as a practical solution for efficient autoregressive speech synthesis without compromising perceptual quality.

Table 2: Subjective evaluation result of SSD with $\beta = 0.4$.

Model	SIM-MOS \uparrow	NAT-MOS \uparrow
Draft Model	3.737 \pm 0.141	3.820 \pm 0.143
CosyVoice 2	3.789 \pm 0.148	3.930 \pm 0.147
w. SSD	3.784 \pm 0.152	3.925 \pm 0.139

4.3. Analysis of Tolerance Factor

We further analyze how tolerance factors influence model acceleration and performance in this section. Table 3 shows the objective evaluation result with different tolerance factors β . We conduct this experiment with the length of draft tokens L_d set to 3. For the efficiency of the model, as the tolerance factor increases, more words are accepted during the speculative decoding process, and the efficiency of the model in synthesizing

speech gradually improves.

Table 3: Objective evaluation result of SSD with different tolerance factor β .

β	WER (%) \downarrow	NMOS \uparrow	SS \uparrow	LM-RTF \downarrow
$\beta = 0.0$	6.34	3.86	0.62	0.509
$\beta = 0.1$	5.71	3.90	0.50	0.444
$\beta = 0.2$	6.08	3.92	0.62	0.408
$\beta = 0.3$	6.29	3.93	0.62	0.386
$\beta = 0.4$	5.70	3.94	0.63	0.360

Notably, experimental results indicate that varying tolerance factors do not lead to significant differences in speech quality, naturalness, or speaker similarity. Specifically, as the tolerance factor increases, more tokens generated by the draft model are accepted, leading to divergent token distributions across different β . However, the speech quality across different β remains consistent, suggesting that a single speech token does not exhibit a one-to-one correspondence with a word or a specific type of paralinguistic information. Instead, multiple combinations of speech tokens can reconstruct similar speech signals through the speech decoder. Thus, this characteristic of the autoregressive speech synthesis system makes SSD a highly rational and promising approach for accelerating the inference process of autoregressive speech synthesis models, as SSD allows for more diverse token generation.

4.4. Analysis of the Length of Draft Tokens

To explore the trade-off between draft token length and model efficiency, we analyze acceleration results with varying L_d . As shown in Figure 3, the efficiency of the model initially ascends and then declines. This is attributed to the fact that as L_d grows larger, the quality of the draft tokens decrease, leading to the generation of many useless tokens due to error accumulation. On the other hand, shorter draft lengths result in higher token acceptance rates but fail to fully leverage the parallel verification capability for accelerating inference. Thus, draft length represents a trade-off between draft model quality and verification model acceleration.




Figure 3: LM-RTF results of SSD with different L_d .

5. Conclusions

In this work, we introduce Speculative Speech Decoding (SSD), a novel approach to accelerate autoregressive speech synthesis. We propose a lightweight draft model constructed by fine-tuning a few parameters from the target model. Experiments on CosyVoice 2 show that our proposed SSD achieves 1.4 \times acceleration while maintaining high synthesis quality. Subjective evaluations demonstrate that the speech generated by SSD achieves comparable performance to the original CosyVoice 2 in terms of similarity and naturalness. SSD demonstrates a reliable and promising solution for more efficient speech synthesis.

6. Acknowledgements

This work is supported by National Natural Science Foundation of China (62076144), National Social Science Foundation of China (13&ZD189) and Shenzhen Science and Technology Program (JCYJ20220818101014030).

7. References

- [1] 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 preprint arXiv:2301.02111*, 2023.
- [2] Z. Zhang, L. Zhou, C. Wang, S. Chen, Y. Wu, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li *et al.*, “Speak foreign languages with your own voice: Cross-lingual neural codec language modeling,” *arXiv preprint arXiv:2303.03926*, 2023.
- [3] P. Anastassiou, J. Chen, J. Chen, Y. Chen, Z. Chen, Z. Chen, J. Cong, L. Deng, C. Ding, L. Gao *et al.*, “Seed-tts: A family of high-quality versatile speech generation models,” *arXiv preprint arXiv:2406.02430*, 2024.
- [4] 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 preprint arXiv:2407.05407*, 2024.
- [5] Z. Du, Y. Wang, Q. Chen, X. Shi, X. Lv, T. Zhao, Z. Gao, Y. Yang, C. Gao, H. Wang *et al.*, “Cosyvoice 2: Scalable streaming speech synthesis with large language models,” *arXiv preprint arXiv:2412.10117*, 2024.
- [6] S. Liao, Y. Wang, T. Li, Y. Cheng, R. Zhang, R. Zhou, and Y. Xing, “Fish-speech: Leveraging large language models for advanced multilingual text-to-speech synthesis,” *arXiv preprint arXiv:2411.01156*, 2024.
- [7] H.-H. Guo, K. Liu, F.-Y. Shen, Y.-C. Wu, F.-L. Xie, K. Xie, and K.-T. Xu, “Fireredtts: A foundation text-to-speech framework for industry-level generative speech applications,” *arXiv preprint arXiv:2409.03283*, 2024.
- [8] E. Kharitonov, D. Vincent, Z. Borsos, R. Marinier, S. Girgin, O. Pietquin, M. Sharifi, M. Tagliasacchi, and N. Zeghidour, “Speak, read and prompt: High-fidelity text-to-speech with minimal supervision,” *Transactions of the Association for Computational Linguistics*, vol. 11, pp. 1703–1718, 2023.
- [9] M. Łajszczak, G. Cambara, Y. Li, F. Beyhan, A. van Korlaar, F. Yang, A. Joly, . Martın-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 preprint arXiv:2402.08093*, 2024.
- [10] P. Peng, P.-Y. Huang, S.-W. Li, A. Mohamed, and D. Harwath, “VoiceCraft: Zero-shot speech editing and text-to-speech in the wild,” in *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, L.-W. Ku, A. Martins, and V. Srikumar, Eds. Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 12442–12462. [Online]. Available: <https://aclanthology.org/2024.acl-long.673/>
- [11] 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,” *arXiv preprint arXiv:2310.00704*, 2023.
- [12] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 3451–3460, 2021.
- [13] N. Zeghidour, A. Luebs, A. Omran, J. Skoglund, and M. Tagliasacchi, “Soundstream: An end-to-end neural audio codec,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 495–507, 2022.
- [14] A. Defossez, J. Copet, G. Synnaeve, and Y. Adi, “High fidelity neural audio compression,” *Transactions on Machine Learning Research*.
- [15] A. Vaswani, “Attention is all you need,” *Advances in Neural Information Processing Systems*, 2017.
- [16] J. Kim, J. Kong, and J. Son, “Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 5530–5540.
- [17] J. Kong, J. Park, B. Kim, J. Kim, D. Kong, and S. Kim, “Vits2: Improving quality and efficiency of single-stage text-to-speech with adversarial learning and architecture design,” in *Interspeech 2023*, 2023, pp. 4374–4378.
- [18] Y. Chen, Z. Niu, Z. Ma, K. Deng, C. Wang, J. Zhao, K. Yu, and X. Chen, “F5-tts: A fairytaler that fakes fluent and faithful speech with flow matching,” *arXiv preprint arXiv:2410.06885*, 2024.
- [19] Z. Ju, Y. Wang, K. Shen, X. Tan, D. Xin, D. Yang, Y. Liu, Y. Leng, K. Song, S. Tang *et al.*, “Naturalspeech 3: Zero-shot speech synthesis with factorized codec and diffusion models,” *arXiv preprint arXiv:2403.03100*, 2024.
- [20] S. E. Eskimez, X. Wang, M. Thakker, C. Li, C.-H. Tsai, Z. Xiao, H. Yang, Z. Zhu, M. Tang, X. Tan *et al.*, “E2 tts: Embarrassingly easy fully non-autoregressive zero-shot tts,” in *2024 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2024, pp. 682–689.
- [21] M. Le, A. Vyas, B. Shi, B. Karrer, L. Sari, R. Moritz, M. Williamson, V. Manohar, Y. Adi, J. Mahadeokar *et al.*, “Voicebox: Text-guided multilingual universal speech generation at scale,” *Advances in neural information processing systems*, vol. 36, 2024.
- [22] C. Chen, S. Borgeaud, G. Irving, J.-B. Lespiau, L. Sifre, and J. Jumper, “Accelerating large language model decoding with speculative sampling,” *arXiv preprint arXiv:2302.01318*, 2023.
- [23] Y. Leviathan, M. Kalman, and Y. Matias, “Fast inference from transformers via speculative decoding,” in *International Conference on Machine Learning*. PMLR, 2023, pp. 19274–19286.
- [24] T. D. Nguyen, J.-H. Kim, J. Choi, S. Choi, J. Park, Y. Lee, and J. S. Chung, “Accelerating codec-based speech synthesis with multi-token prediction and speculative decoding,” in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [25] B. Li, H. Wang, S. Zhang, Y. Guo, and K. Yu, “Fast and high-quality auto-regressive speech synthesis via speculative decoding,” in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [26] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly, “Parameter-efficient transfer learning for nlp,” in *International conference on machine learning*. PMLR, 2019, pp. 2790–2799.
- [27] B. Thomas, S. Kessler, and S. Karout, “Efficient adapter transfer of self-supervised speech models for automatic speech recognition,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7102–7106.
- [28] A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei *et al.*, “Qwen2.5 technical report,” *arXiv preprint arXiv:2412.15115*, 2024.
- [29] 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,” *arXiv preprint arXiv:1904.02882*, 2019.
- [30] C. Du, Y. Guo, F. Shen, Z. Liu, Z. Liang, X. Chen, S. Wang, H. Zhang, and K. Yu, “Unicats: A unified context-aware text-to-speech framework with contextual vq-diffusion and vocoding,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 16, 2024, pp. 17924–17932.