

Accelerating Diffusion-based Text-to-Speech Model Training with Dual Modality Alignment

Jeongsoo Choi^{1*}, Zhikang Niu^{2,3*}, Ji-Hoon Kim¹, Chunhui Wang⁴, Joon Son Chung¹, Xie Chen^{2,3}

¹School of Electrical Engineering, KAIST, South Korea

²MoE Key Lab of Artificial Intelligence, X-LANCE Lab, School of Computer Science, Shanghai Jiao Tong University, China ³Shanghai Innovation Institute, China ⁴Geely, China

{jeongsoo.choi, joonson}@kaist.ac.kr, {zhikangniu, chenxie95}@sjtu.edu.cn

Abstract

The goal of this paper is to optimize the training process of diffusion-based text-to-speech models. While recent studies have achieved remarkable advancements, their training demands substantial time and computational costs, largely due to the implicit guidance of diffusion models in learning complex intermediate representations. To address this, we propose **A-DMA**, an effective strategy for **Accelerating training with Dual Modality Alignment**. Our method introduces a novel alignment pipeline leveraging both text and speech modalities: text-guided alignment, which incorporates contextual representations, and speech-guided alignment, which refines semantic representations. By aligning hidden states with discriminative features, our training scheme reduces the reliance on diffusion models for learning complex representations. Extensive experiments demonstrate that A-DMA doubles the convergence speed while achieving superior performance over baselines.

Index Terms: text-to-speech synthesis, diffusion model, modality alignment

1. Introduction

Text-to-Speech (TTS) aims to synthesize natural human speech x from text inputs y . Recent advancements, driven by scaling up both data and model size, have enabled the generation of natural speech that closely mimics the voice of any given reference speech x' , a capability known as zero-shot TTS. Existing methods for zero-shot TTS can be broadly categorized into two approaches: Autoregressive (AR) [1–5] and Non-Autoregressive (NAR) [6–10] models. While AR models have shown exceptional performance in reproducing the acoustic characteristics of reference speaker, they suffer from slow inference speed and error accumulation owing to their structural limitations [11–13].

NAR models, primarily built on diffusion networks, address the inherent limitations of AR models, enhancing both inference speed and output quality. For example, VoiceBox adopts a flow-matching algorithm [14], outperforming AR-based VALL-E [1] in quality with 20 times faster inference speed. DiTTo-TTS [10] employs a Diffusion Transformer (DiT) [15] within latent diffusion models (LDMs), further advancing both speech quality and inference speed. Most recently, E2 TTS [8] and F5-TTS [9] have achieved notable improvements with simpler architectures, demonstrating the drawbacks of external aligners or explicit duration modeling, which are commonly used in prior studies. Despite these advances, training such methods remains time-consuming due to the implicit guidance of diffusion networks in learning complex representations [16–18].

In this paper, we present A-DMA, a simple yet effective strategy for Accelerating training through Dual Modality Alignment. Without altering the model architecture, our approach presents two novel alignment pipelines: text-guided alignment and speech-guided alignment, each leveraging modality-specific knowledge from text and speech, respectively. Text-guided alignment utilizes a Connectionist Temporal Classification (CTC)-based compressor to align text information with hidden states, thereby enriching contextual representations within diffusion networks. In parallel, speech-guided alignment incorporates self-supervised speech features into intermediate states through convolutional layers, enhancing the semantic representations within diffusion models.

Given the transformation from text input into the corresponding speech output, we employ text-guided alignment in the earlier diffusion blocks, while speech-guided alignment is applied in the later diffusion blocks. With extensive ablation studies, we validate the design choices of our architecture. Furthermore, experimental results show that our approach achieves state-of-the-art performance on zero-shot TTS benchmarks, confirming the effectiveness of A-DMA from various perspectives. Code and demo samples are available at: <https://github.com/ZhikangNiu/A-DMA>

2. Method

As illustrated in Figure 1, our model introduces an improved training paradigm for a diffusion-based text-to-speech (TTS) system, leveraging dual-modality alignment to accelerate convergence while maintaining a straightforward inference process and high speech fidelity. In this section, we provide a detailed introduction to our proposed method, highlighting its key components and underlying motivations. Section 2.1 introduces our baseline TTS architecture and training scheme. Section 2.2 presents a text-based representation alignment method aimed at improving intelligibility, and Section 2.3 outlines our strategy for enhancing internal speech representations through pre-trained self-supervised learning (SSL) speech models.

2.1. Diffusion Model for Text-to-Speech

Recent advancements in non-autoregressive TTS have established critical foundations for our work. Voicebox [6] is a non-autoregressive flow matching model with a UNet-style Diffusion Transformer, outperforming the previous state-of-the-art TTS model VALL-E [1] in intelligibility, speaker similarity, and inference latency. However, its dependency on external duration prediction and forced text alignment introduces pre-processing complexity. Therefore, E2 TTS [8] eliminates these additional components (duration model, grapheme-to-phoneme conversion, and forced alignment) and converts the input text directly

*These authors contributed equally.

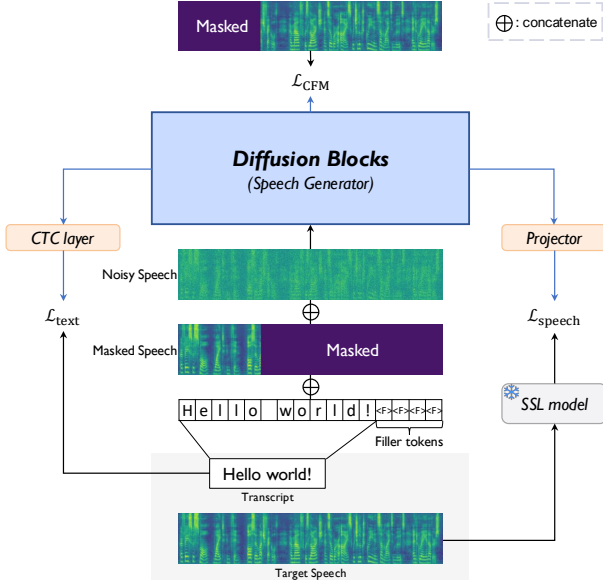


Figure 1: Overall training framework. The model uses masked speech, noisy speech, and transcript to predict the masked part of the speech via flow matching. The SSL model and CTC layer are utilized only during training, with the SSL parameters frozen throughout the process.

into a character sequence with filler tokens, but E2 TTS suffers from slow convergence and low robustness. To address these limitations, F5-TTS [9] introduces three key innovations: 1) A ConvNeXt-based text refinement module to facilitate alignment with speech; 2) An inference-time Sway Sampling strategy, significantly enhancing both the performance and efficiency compared with E2 TTS; 3) Replacing the backbone from U-Net [19] style Transformer [20] to Diffusion Transformer (DiT) [15].

The flow matching framework aims to construct a probability path p_t between simple distribution p_0 (typically Gaussian) and data distribution q through a time-dependent vector field v_t , generating flows ψ_t for sampled flow step $t \in [0, 1]$. The Conditional Flow Matching (CFM) loss is defined as:

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t, q(x_1), p(x_0)} \|v_t(\psi_t(x)) - \frac{d}{dt}\psi_t(x)\|^2,$$

where $x_1 \sim q$ represents the training audio data and x_0 denotes sampled Gaussian noise. We also employ the optimal transport conditional flow matching (OT-CFM) configuration $\psi_t(x) = (1-t)x_0 + tx_1$, simplifying the loss to:

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{t, q(x_1), p(x_0)} \|v_t((1-t)x_0 + tx_1) - (x_1 - x_0)\|^2.$$

In our work, we adopt F5-TTS as the baseline since it serves as a strong benchmark with state-of-the-art text-to-speech capabilities. The F5-TTS is trained on a text-guided speech-infilling task using optimal transport flow matching, where the model predicts masked speech segments based on noisy speech $\psi_t(x) = (1-t)x_0 + tx_1$, masked speech $x_m = (1-m) \odot x_1$, and the transcript y with filler tokens, and $m \in \{0, 1\}^{F \times N}$ represents a binary temporal mask, where F is mel dimension and N is the sequence length.

2.2. Text Representation Alignment

A key requirement in TTS is the generation of intelligible speech. One challenge in TTS is the large discrepancy in

sequence lengths between text and speech. To accelerate the learning of correlations between these modalities, we introduce a text representation alignment loss by leveraging an auxiliary CTC [21] loss at an intermediate layer of our model. Let $\Phi_i(\psi_t(x), x_m, y)$ denote the hidden representation at the i -th Transformer layer, where the inputs include the noisy speech $\psi_t(x)$, the masked speech x_m , and the corresponding text embedding y . Unlike conventional auxiliary CTC approaches such as InterCTC [22, 23] primarily used in speech recognition tasks, our approach uniquely applies CTC supervision to the generative model’s intermediate states. This early supervision ensures that the model learns the alignment between semantic (text) and acoustic (speech) features at the onset of processing. Formally, the text alignment loss is defined as:

$$\mathcal{L}_{\text{text}} = -\log p_{\text{CTC}}(y | \Phi_i(\psi_t(x), x_m, y)),$$

where p_{CTC} is the probability of generating the target text y from the intermediate representation. This loss encourages the model to maintain high fidelity to the input text throughout the transformation process, ultimately leading to more intelligible synthesized speech.

2.3. Speech Representation Alignment

Recent studies [24, 25] have demonstrated that incorporating representations from pre-trained self-supervised learning (SSL) speech models, such as HuBERT [26], WavLM [27] and W2V-Bert [28] can significantly enhance the quality of internal representations. Leveraging this insight, we incorporate a speech representation alignment loss into our TTS framework. Specifically, we employ a pre-trained SSL speech model f to extract high-quality representations from the reference speech x . To align the temporal and feature dimensions of the intermediate representations from our TTS model with the output of the SSL model, we apply an interpolation layer followed by a 1D convolution layer. This can be formulated as:

$$h_t = \text{Conv1D}(\text{Interp}(\Phi_i(\psi_t(x), x_m, y))).$$

The speech alignment loss is then defined as:

$$\mathcal{L}_{\text{speech}} = -\frac{1}{N} \sum_{n=1}^N \cos(h_t^{[n]}, f(x)^{[n]}),$$

where N denotes the total sequence length, and $\cos(\cdot, \cdot)$ computes the cosine similarity between the aligned TTS feature and the SSL feature at each time step. This loss term regularizes the TTS model by encouraging its internal representations to mimic the robust, high-quality features learned by the SSL model.

2.4. Overall Training Objective

To jointly optimize our model for both high-quality speech generation and multimodal alignment, we integrate the conditional flow matching loss with the two auxiliary alignment losses. The overall loss function is defined as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{CFM}} + \lambda_{\text{text}} \mathcal{L}_{\text{text}} + \lambda_{\text{speech}} \mathcal{L}_{\text{speech}},$$

where λ_{text} and λ_{speech} are hyperparameters that control the relative importance of the text and speech alignment losses. In our experiments, they are empirically set to 0.1 and 1, respectively, to balance their contributions. By optimizing this joint objective, the model learns to synthesize speech that is not only natural and expressive but also well-aligned with the input text and the reference speaker’s characteristics.

Table 1: *Experiments on LibriSpeech-PC test-clean dataset. The best results among low-resource setting are marked in **bold**.*

Method	# Param.	Data	WER(%)↓	SIM↑
Ground Truth			2.23	0.69
<i>High-resource</i>				
CosyVoice [29]	300M	170kh	3.59	0.66
FireRedTTS [30]	580M	248kh	2.69	0.47
E2 TTS [8]	333M	100kh	2.95	0.69
F5-TTS [9]	336M	100kh	2.42	0.66
<i>Low-resource</i>				
USLM [24]	361M	0.6kh	6.11	0.43
E2 TTS [8]	157M	0.6kh	3.10	0.63
+ A-DMA (Ours)	157M	0.6kh	2.25	0.65
F5-TTS [9]	159M	0.6kh	2.68	0.60
+ A-DMA (Ours)	159M	0.6kh	1.97	0.62

3. Experiments

3.1. Dataset

We train our model on the LibriTTS [31] dataset, a multi-speaker English corpus containing approximately 585 hours of read speech sampled at 24 kHz. The dataset is derived from the LibriSpeech [32] corpus and curated specifically for TTS research. Utterances with significant background noise are excluded to ensure high-quality data for training.

3.2. Preprocessing

For text, we use character sequences for both input and text-guided alignment. For speech, we use audio with 24 kHz sampling rate and extract mel-spectrograms with 100 mel filter banks, a window size of 1024, and a hop size of 256, resulting in approximately 94 Hz. Since the mel-spectrogram sequence is generally much longer than the character sequence, filler tokens are concatenated to the character sequence to match the length of the mel-spectrogram. For speech-guided alignment, we re-sample the audio to 16 kHz before feeding it into the speech SSL model.

3.3. Model Setup

Our baseline model is F5-TTS Small, which consists of 18 layers, 12 attention heads, and a 768-dimensional embedding for the DiT block; and 4 layers with a 512-dimensional embedding for the ConvNeXt V2 [33] module, totaling 159M parameters. We train the model with a total batch size of approximately 1k seconds, using the AdamW [34] optimizer with a learning rate of 7.5×10^{-5} . The learning rate is scheduled with a warmup phase of 20,000 updates, followed by linear decay. Training is conducted on 4 RTX A6000 GPUs.

For inference, we apply Exponential Moving Average (EMA) weights and the sway sampling strategy with a coefficient of -1, utilizing the Euler ODE solver for F5-TTS, and the midpoint method for E2-TTS, as described in [9]. To synthesize audio from the generated mel-spectrograms, we use the pre-trained vocoder Vocos [35], which was trained on LibriTTS for 1.2M steps.

3.4. Evaluation

We follow the evaluation setup used in F5-TTS to assess our model on the LibriSpeech-PC [36] test set. This test set con-




Figure 2: *The WER and SIM results showing the effectiveness of the proposed dual modality alignment.*

sists of 2.2 hours of data from the LibriSpeech [32] test-clean subset, comprising 1,127 samples, each ranging from 4 to 10 seconds in length. It is specifically designed to facilitate comparison across different models in the research community and is publicly available for benchmarking. Specifically, we assess its ability to generate speech that aligns with the content of the provided text and the speaker characteristics of the given speech prompt. We perform the cross-sentence generation task using two utterances of the same speaker, where one serves as the speech prompt and the other is the target speech to be synthesized. The evaluation is conducted using the following three metrics:

Word Error Rate (WER) measures the intelligibility of synthesized speech by comparing its transcription to the ground-truth text. We employ Whisper-large-V3 [37] for automatic transcription and compute WER accordingly.

Speaker Similarity (SIM) quantifies the resemblance between the synthesized and ground-truth speech in terms of speaker identity. We use the WavLM-large-based [27] speaker verification model to extract speaker embeddings and compute their cosine similarity between synthesized and ground truth speeches.

UTMOS [38] is an objective metric for assessing the naturalness of synthesized speech. It employs an open-source MOS prediction model to estimate audio quality without requiring reference recordings or labels. While not an absolute measure, UTMOS provides a practical and efficient way to gauge the perceived naturalness of synthetic speech.

4. Results

4.1. Quantitative Comparison

Table 1 shows the evaluation results of previous methods and our proposed approach **A-DMA** on LibriSpeech-PC [36] test-clean dataset. As shown in high-resource data setting, E2 TTS and F5-TTS which are diffusion-based NAR models demonstrate comparable performance to state-of-the-art AR models such as CosyVoice [29] and FireRedTTS [30]. To assess the effectiveness of A-DMA, we apply it to E2 TTS and F5-TTS while keeping all training settings identical, with only our method added. The results demonstrate that our method achieves superior performance in both WER and SIM metrics than existing methods including USLM [24]. Notably, our results highlight the effectiveness of the proposed dual alignment method in enhancing the alignment between input text and generated speech, even with limited training data. Importantly, our approach does not modify the inference pipeline, ensuring that the inference speed remains unchanged while achieving improved performance.

Table 2: Ablation study of modality aligners on LibriSpeech-PC test-clean dataset. All of the models are trained on LibriTTS for about 100 epochs. Here, “speech align” uses a HuBERT-large alignment, and “text align” uses a CTC alignment.

Method	$\mathcal{L}_{\text{text}}$	$\mathcal{L}_{\text{speech}}$	WER(%)↓	UTMOS↑	SIM↑
Baseline	\times	\times	7.474	4.040	0.578
+ text align	2nd layer	\times	7.188	4.053	0.591
	4th layer	\times	4.558	4.052	0.592
	8th layer	\times	2.602	4.076	0.586
	12th layer	\times	2.483	4.079	0.589
	16th layer	\times	2.965	4.033	0.587
+ speech align	\times	2nd layer	8.288	4.056	0.591
	\times	4th layer	7.421	4.091	0.592
	\times	8th layer	4.900	4.121	0.606
	\times	12th layer	3.521	4.089	0.609
	\times	16th layer	4.470	4.012	0.600
+ dual align	8th layer	8th layer	3.063	4.087	0.606
	12th layer	12th layer	2.688	4.076	0.604
	8th layer	12th layer	2.353	4.054	0.609
	8th layer	16th layer	2.226	4.024	0.600

To further analyze the effectiveness of our approach, Figure 2 illustrates the performance trends across training steps. The results indicate that our method significantly accelerates the alignment process between the input text and generated speech while maintaining higher performance upon convergence, demonstrating its robustness and efficiency in low-resource scenarios.

4.2. Layer-wise Analysis

We conduct extensive layer-wise experiments to investigate the impact of alignment at different layers as shown in Table 2. Our findings indicate that each modality achieves optimal performance when aligned at the 12th layer. As expected, text alignment significantly enhances intelligibility (WER), while speech alignment improves speaker similarity. Applying alignment at the early layers for both modalities does not lead to substantial improvements, suggesting that these layers are insufficient to effectively align with discriminative representations.

Furthermore, we compare alignment at the 8th and 16th layers and observe that applying text guidance at a lower layer and speech guidance at a higher layer yields superior performance. These results suggest that optimal alignment occurs at different depths for each modality. To further investigate, we apply our proposed dual alignment approach. We find that aligning the same 12th layer for both modalities leads to sub-optimal performance, whereas single modality alignment achieves better results. This is reasonable given the inherent representation gap between text and speech features, causing a trade-off when aligning both simultaneously. To address this, we explore aligning text at a lower layer than speech alignment, finding that this strategy significantly improves performance, outperforming both single-modality alignment and the baseline.

Notably, our method maintains a high UTMOS score, demonstrating that the improvements in alignment do not come at the cost of naturalness. The UTMOS scores remain comparable with the baseline, ensuring that our dual alignment approach preserves speech quality and naturalness while achieving enhanced intelligibility and speaker similarity.

Table 3: Ablation study of speech SSL model and speech align equation after 100 epochs trained on LibriTTS dataset - $\mathcal{L}_{\text{speech}}$ is applied to 12th layer. “Avg.” refers to the average of the intermediate outputs of the SSL model, while “Last” refers to using only the output of the last layer of the SSL model.

SSL model	Output	$\mathcal{L}_{\text{speech}}$	WER(%)↓	UTMOS↑	SIM↑
Baseline		\times	7.474	4.040	0.578
HuBERT	Last	$ \Phi_n - f(x) $	6.386	4.068	0.605
HuBERT	Last	$-\log(\sigma(\cos(\Phi_n, f(x))))$	5.104	4.056	0.606
HuBERT	Last	$-\cos(\Phi_n, f(x))$	3.521	4.089	0.609
HuBERT	Avg.	$-\cos(\Phi_n, f(x))$	4.716	4.070	0.618
WavLM	Last	$-\cos(\Phi_n, f(x))$	3.992	4.100	0.612

4.3. Ablation Study

Table 3 presents an ablation study on the speech alignment method. We experiment with various losses and determine that the negative cosine similarity loss (line #3, #4, and #5) yields the best results compared to L1 (line #1) and the negative cosine similarity loss with sigmoid activation (line #2). When we use the average layer output of the SSL speech model, we find it improves speaker similarity (SIM) slightly, but degrades the WER compared to using the last layer alone. This trade-off suggests that averaging across all layers may introduce information redundancy, which can bias the model by incorporating less relevant or noisy features from earlier layers. For all other experiments, we adopt the last layer of HuBERT with negative cosine similarity loss.

5. Conclusion

In this paper, we propose A-DMA, a novel modality alignment strategy that effectively guides diffusion-based text-to-speech models by leveraging discriminative representations. This approach significantly accelerates model convergence while enhancing the quality of generated speech. Experimental results validate that our approach outperforms the baseline system in both training efficiency and speech quality while maintaining a straightforward inference pipeline. We further validate the effectiveness of each proposed component through comprehensive layer-wise analysis and ablation studies. We believe that these promising findings pave the way for future research in advanced modality alignment techniques and further optimization of text-to-speech systems.

6. Limitation

Although our extensive layer-wise experiments demonstrate that the proposed dual modality alignment leads to superior generation performance with faster convergence, this work does not target model efficiency. In future work, we plan to further optimize the framework by reducing model complexity and accelerating inference.

7. Acknowledgements

J. Choi, J.-H. Kim, and J. S. Chung were supported by IITP grant funded by the Korea government (MSIT) (RS-2024-00457882, National AI Research Lab Project). Z. Niu and X. Chen were supported by the National Natural Science Foundation of China (No. U23B2018 and No. 62206171), Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102.

8. 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:2301.02111*, 2023.
- [2] 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:2406.02430*, 2024.
- [3] P. Peng, P.-Y. Huang, D. Li, A. Mohamed, and D. Harwath, “Voicecraft: Zero-shot speech editing and text-to-speech in the wild,” in *Proc. ACL*, 2024.
- [4] M. Łajszczak, G. C  mbara, Y. Li, F. Beyhan, A. van Korlaar, F. Yang, A. Joly,   . Mart  n-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.
- [5] L. Meng, L. Zhou, S. Liu, S. Chen, B. Han, S. Hu, Y. Liu, J. Li, S. Zhao, X. Wu *et al.*, “Autoregressive speech synthesis without vector quantization,” *arXiv:2407.08551*, 2024.
- [6] 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,” in *NeurIPS*, 2024.
- [7] K. Shen, Z. Ju, X. Tan, Y. Liu, Y. Leng, L. He, T. Qin, S. Zhao, and J. Bian, “Naturspeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers,” in *Proc. ICLR*, 2024.
- [8] 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 *IEEE Spoken Language Technology workshop*, 2024.
- [9] 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:2410.06885*, 2024.
- [10] K. Lee, D. W. Kim, J. Kim, and J. Cho, “Ditto-tts: Efficient and scalable zero-shot text-to-speech with diffusion transformer,” in *Proc. ICLR*, 2025.
- [11] 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 *Proc. ICASSP*, 2025.
- [12] 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.
- [13] 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,” in *Proc. ICASSP*, 2025.
- [14] Y. Lipman, R. T. Chen, H. Ben-Hamu, M. Nickel, and M. Le, “Flow matching for generative modeling,” in *Proc. ICLR*, 2023.
- [15] W. Peebles and S. Xie, “Scalable diffusion models with transformers,” in *Proc. ICCV*, 2023.
- [16] T. Hang, S. Gu, C. Li, J. Bao, D. Chen, H. Hu, X. Geng, and B. Guo, “Efficient diffusion training via min-snr weighting strategy,” in *Proc. ICCV*, 2023.
- [17] J. Yao, C. Wang, W. Liu, and X. Wang, “Fasterdit: Towards faster diffusion transformers training without architecture modification,” in *NeurIPS*, 2024.
- [18] S. Yu, S. Kwak, H. Jang, J. Jeong, J. Huang, J. Shin, and S. Xie, “Representation alignment for generation: Training diffusion transformers is easier than you think,” in *Proc. ICLR*, 2025.
- [19] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Proc. MICCAI*, 2015.
- [20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez,   . Kaiser, and   . Polosukhin, “Attention is all you need,” in *NeurIPS*, 2017.
- [21] A. Graves, S. Fern  ndez, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in *Proc. ICML*, 2006.
- [22] J. Lee and S. Watanabe, “Intermediate loss regularization for ctc-based speech recognition,” in *Proc. ICASSP*, 2021.
- [23] M. Burchi and R. Timofte, “Audio-visual efficient conformer for robust speech recognition,” in *Proc. WACV*, 2023.
- [24] X. Zhang, D. Zhang, S. Li, Y. Zhou, and X. Qiu, “Speechtokenizer: Unified speech tokenizer for speech large language models,” in *Proc. ICLR*, 2024.
- [25] Z. Ju, Y. Wang, K. Shen, X. Tan, D. Xin, D. Yang, Y. Liu, Y. Leng, K. Song, S. Tang *et al.*, “Naturspeech 3: Zero-shot speech synthesis with factorized codec and diffusion models,” in *Proc. ICML*, 2024.
- [26] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2021.
- [27] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, “Wavlm: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [28] Y.-A. Chung, Y. Zhang, W. Han, C.-C. Chiu, J. Qin, R. Pang, and Y. Wu, “W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training,” in *IEEE Automatic Speech Recognition and Understanding workshop*, 2021.
- [29] 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.
- [30] H.-H. Guo, K. Liu, F.-Y. Shen, Y.-C. Wu, F.-L. Xie, K. Xie, and K.-T. Xu, “Firedtts: A foundation text-to-speech framework for industry-level generative speech applications,” *arXiv:2409.03283*, 2024.
- [31] 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.
- [32] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *Proc. ICASSP*, 2015.
- [33] S. Woo, S. Debnath, R. Hu, X. Chen, Z. Liu, I. S. Kweon, and S. Xie, “Convnext v2: Co-designing and scaling convnets with masked autoencoders,” in *Proc. CVPR*, 2023.
- [34] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *Proc. ICLR*, 2019.
- [35] H. Siuzdak, “Vocos: Closing the gap between time-domain and fourier-based neural vocoders for high-quality audio synthesis,” in *Proc. ICLR*, 2024.
- [36] A. Meister, M. Novikov, N. Karpov, E. Bakhturina, V. Lavrukhin, and B. Ginsburg, “Librispeech-pc: Benchmark for evaluation of punctuation and capitalization capabilities of end-to-end asr models,” in *IEEE Automatic Speech Recognition and Understanding workshop*, 2023.
- [37] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proc. ICML*, 2023.
- [38] 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.