

MAGE: A COARSE-TO-FINE SPEECH ENHANCER WITH MASKED GENERATIVE MODEL

The Hieu Pham^{1*}, Tan Dat Nguyen^{2*}, Phuong Thanh Tran¹, Joon Son Chung², Duc Dung Nguyen^{1†}

¹AITech Lab, Ho Chi Minh City University of Technology, VNU-HCM, Vietnam

²Korea Advanced Institute of Science and Technology, South Korea

ABSTRACT

Speech enhancement remains challenging due to the trade-off between efficiency and perceptual quality. In this paper, we introduce **MAGE**, a Masked Audio Generative Enhancer that advances generative speech enhancement through a compact and robust design. Unlike prior masked generative models with random masking, MAGE employs a scarcity-aware coarse-to-fine masking strategy that prioritizes frequent tokens in early steps and rare tokens in later refinements, improving efficiency and generalization. We also propose a lightweight corrector module that further stabilizes inference by detecting low-confidence predictions and re-masking them for refinement. Built on BigCodec and finetuned from Qwen2.5-0.5B, MAGE is reduced to 200M parameters through selective layer retention. Experiments on DNS Challenge and noisy LibriSpeech show that MAGE achieves state-of-the-art perceptual quality and significantly reduces word error rate for downstream recognition, outperforming larger baselines. Audio examples are available at <https://hieugiaosu.github.io/MAGE>.

Index Terms— Speech Enhancement, Masked Generative Model, Generative Model, Masking Strategy

1. INTRODUCTION

Speech enhancement (SE) [1, 2] is a fundamental task in speech processing that seeks to separate clean and intelligible signals from audio corrupted by diverse distortions. These degradations arise from background noise, reverberation, microphone clipping, transmission artifacts, and resampling effects, and they can severely impair both human perception and machine processing. Achieving reliable SE is therefore crucial for a wide range of downstream applications, including personalized enhancement [3], robust speech recognition [4], and scalable data collection [5].

Modern approaches fall broadly into two categories. *Discriminative models* directly map noisy inputs to clean signals by optimizing losses aligned with intrusive metrics such as SI-SDR, PESQ, and STOI. These models are computationally efficient, with DPRCN [6] achieving *PESQ* 2.86 using only

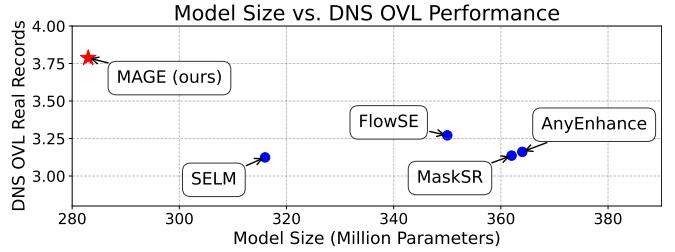


Fig. 1. A comparison of DNS OVL scores on a real-world test set against model size for several different methods

0.8M parameters and TF-GridNet [7] reaching *PESQ* 3.78 and *STOI* 0.994 with 9.8M parameters. Nonetheless, despite progress in augmentation and adversarial training [8], their performance often degrades when acoustic conditions differ from training data. *Generative models* instead seek to model the underlying speech distribution. Diffusion- and flow-based approaches such as SGMSE [1] and StoRM [2] improve perceptual quality by reversing the corruption process, yet they remain computationally heavy and data-demanding, even with efficiency-oriented variants like SpeechFlow [9]. Discrete representations facilitate language-model-driven SE, as demonstrated by SELM [10], which leverages WavLM [11], and MaskSR [12] in combination with DAC [13]. However, these systems are still large and impractical for deployment. Recently, masked generative models (MGMs) such as MaskSR and AnyEnhance [14] have shown state-of-the-art results, but they typically require hundreds of millions to billions of parameters and rely on random masking strategies that introduce inefficiency and redundancy. This gap motivates the need for generative SE models that deliver both perceptual quality and efficiency.

In this work, we propose **MAGE**, a Masked Audio Generative Enhancer that combines the perceptual advantages of generative modeling with the efficiency needed for practical deployment. Unlike prior masked generative models that rely on random masking and large parameter counts, MAGE introduces a scarcity-aware coarse-to-fine masking strategy and a lightweight corrector module, enabling more effective training and stable inference. Finetuned from Qwen2.5-0.5B and speech tokens from BigCodec, our proposed model is reduced to only 200M parameters while

* These authors contributed equally to this work.

† Corresponding author

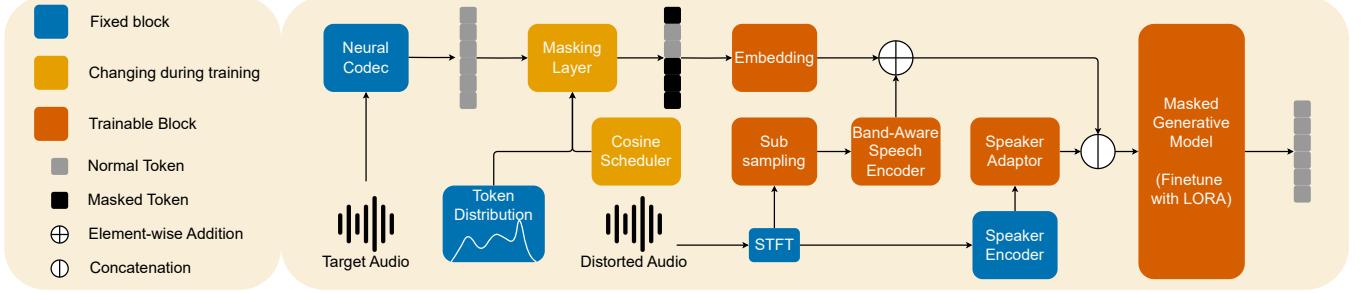


Fig. 2. Training pipeline and model design of MAGE. Target audio is first converted into sequence of tokens using a Neural Codec. These tokens are then masked according to their distribution to form a coarse-to-fine masking strategy, as described in Sec. 2.3. Besides, speaker identity is extracted by a Band-Aware Encoder and a pretrained Speaker Encoder, enabling it to capture the acoustic characteristics. The model is optimized using cross-entropy loss applied only on the masked tokens.

maintaining strong performance. Extensive experiments on DNS Challenge and noisy LibriSpeech show that MAGE achieves state-of-the-art perceptual quality and significantly improves word error rate (WER) for downstream automatic speech recognition, outperforming larger flow-based and mask-based baselines. These results highlight MAGE as a compact, generalizable, and robust framework for speech enhancement, advancing the development of resource-efficient generative models for real-world applications.

2. MAGE

2.1. Neural Codec and Encoder

To extract discrete representations from target audio, we adopt BigCodec [15], which provides stable tokenization and high-quality reconstruction using a single codebook with 80 tokens per second. To match this token rate, the distorted audio $w_{\text{distorted}}$ is converted into a complex spectrogram $S \in \mathbb{C}^{F \times T}$ and downsampled with 2D convolutions. For conditioning, we leverage a TF-GridNet block [7], which efficiently models cross-band frequency interactions while remaining lightweight. The encoder output is projected into the embedding space $x_{\text{cond}} \in \mathbb{R}^D$ to condition the generative model. In addition, this complex spectrogram is used to extract speaker identity x_e with a pretrained speaker encoder¹.

2.2. Masked Generative Model (MGM)

Given a token sequence $x = [x^1, \dots, x^T]$ derived from $w_{\text{distorted}}$, the MGM learns to reconstruct masked tokens over N denoising steps. At each step i , tokens are masked according to a binary vector $m_{(i)}$ sampled from a Bernoulli distribution:

$$p(i) = \cos\left(\frac{\pi}{2} \frac{i}{N}\right), \quad (1)$$

where $p(i)$ follows a cosine schedule [14, 12]. We obtain the masked sequence $\tilde{x}_{(i)}$ by filling the mask token M into

the positions selected for masking, while keeping the remaining tokens unchanged. The masked sequence is then added element-wise with x_{cond} , and concatenated with the speaker embedding, which is projected from x_e using a lightweight adaptor. The resulting sequence is used as the input to the masked generative model. The model parameters θ are optimized by predicting the masked tokens:

$$\mathcal{L}_{\text{mask}} = - \sum_{t=1}^T m_{(i)}^t \log P(x^t | \tilde{x}_{(i)}, x_{\text{cond}}, x_e; \theta). \quad (2)$$

Intuitively, this objective forces the model to learn conditional distributions of tokens given both observed context and acoustic information.

2.3. Coarse-to-Fine (CTF) Strategy and Corrector

Uniform masking in Eq. 1 treats all tokens equally, ignoring the fact that token frequencies are highly non-uniform. This creates a bias: frequent tokens dominate training, while rare tokens are predicted less reliably. In other words, the model learns ‘popular’ predictions more often than ‘rare’ predictions, which leads to suboptimal generalization.

To mitigate this imbalance, we introduce a **coarse-to-fine (CTF) masking strategy**. We compute a frequency vector $f = [f^1, \dots, f^T]$ for each input x , where f^i is the document frequency [18] of token x^i in the training corpus. We then calculate an IDF-like score:

$$z = \log\left(\frac{N + 1}{f + 1}\right), \quad (3)$$

where N is the total number of samples of dataset. Higher z indicates rarer tokens. The base masking probability for each token is

$$p_{\text{base}} = \sigma\left(\frac{z - \bar{z}}{\text{std}(z)}\right), \quad (4)$$

where $\sigma(\cdot)$ is the sigmoid function, and \bar{z} and $\text{std}(z)$ denote the mean and standard deviation of z . To preserve the advantages of the cosine schedule while making it token-dependent,

¹<https://github.com/resemble-ai/Resemblyzer>

System	With Reverb				Without Reverb				Real Recordings		
	SIG↑	BAK↑	OVL↑	SSIM↑	SIG↑	BAK↑	OVL↑	SSIM↑	SIG↑	BAK↑	OVL↑
BigCodec Resyn. GT	4.473	4.471	4.190	0.857	4.473	4.471	4.190	0.857	–	–	–
Noisy	1.760	1.497	1.392	–	3.392	2.618	2.483	–	3.053	2.510	2.255
Conv-TasNet [16]	2.415	2.710	2.010	0.939	3.092	3.341	3.001	0.945	3.102	2.975	2.410
SGMSE [1]	2.730	2.741	2.430	0.899	3.501	3.710	3.137	0.934	3.297	2.894	2.793
StoRM [2]	2.947	3.141	2.516	<u>0.934</u>	3.514	3.941	3.205	<u>0.943</u>	3.410	3.379	2.940
ANYENHANCE [14]	3.500	4.040	3.204	–	3.640	4.179	3.418	–	3.488	3.977	3.161
MaskSR-M [12]	3.531	4.065	3.253	0.827	3.586	4.116	3.339	0.929	3.430	4.025	3.136
FlowSE [17]	3.614	4.110	3.340	0.809	3.690	4.200	3.451	0.940	3.643	4.100	3.271
MAGE	3.530	4.149	3.107	0.724	4.407	4.515	4.151	0.817	3.830	4.302	3.500
+ Corrector	3.525	<u>4.146</u>	3.081	0.724	4.441	4.557	4.201	0.800	4.098	<u>4.309</u>	<u>3.744</u>
+ CTF	3.876	3.901	3.653	0.799	<u>4.559</u>	4.408	4.235	0.819	4.206	4.145	3.787
+ CTF & Corrector	<u>3.864</u>	3.961	<u>3.372</u>	0.789	4.580	4.338	<u>4.223</u>	0.821	<u>4.191</u>	3.924	3.666

Table 1. Performance comparison on DNS Challenge test set. DNSMOS scores (SIG/BAK/OVL) and speaker cosine similarity (SSIM) are reported. Higher values (↑) indicate better performance. BigCodec Resyn. GT denotes the ground-truth signal re-encoded and decoded to illustrate the upper bound of the method.

we define the final CTF probability:

$$p_{\text{CTF}} = \min\left(\frac{E_{\cos}}{E_{\text{base}}} \cdot p_{\text{base}}, 1\right), \quad (5)$$

where $E_{\cos} = T \sin\left(\frac{\pi i}{2N}\right)$ is the expected number of masked tokens under the cosine schedule, and $E_{\text{base}} = \sum_{t=1}^T p_{\text{base}}^t$ under the base distribution. This adjustment ensures that frequent tokens are predicted earlier, while rare tokens are emphasized in later steps, creating a natural curriculum.

Corrector. Finally, to further improve robustness, we add a correction stage. Unlike standard MGM which generates tokens once and commits to them, MAGE allows re-masking and regeneration of previously decoded tokens. This prevents error accumulation and enables iterative refinement. Instead of random re-masking as in MaskSR [12], we train a lightweight 4-layer BLSTM corrector that identifies low-confidence tokens and re-masks them. During training, up to 30% of ground-truth tokens are randomly corrupted, teaching the corrector to detect inconsistencies. At inference, it selectively re-masks problematic tokens and passes them back to the generative model for correction. This design improves perceptual quality by enabling the model to self-revise.

3. EXPERIMENTAL SETUP

We finetune the masked language model from Qwen2.5-0.5B using LoRA [19]. To reduce computational cost, only half of the original layers are retained, resulting in a compact model with 200M parameters. The MAGE speech encoder applies Short-time Fourier Transform with $n_{\text{fft}} = 256$, $window = 256$, $hop_size = 100$, followed by two TF-GridNet blocks with embedding size of 48, BLSTM hidden

System	LibriSpeech			
	SIG↑	BAK↑	OVL↑	WER↓
SGMSE [1]	4.254	4.109	3.813	28.52
StoRM [2]	4.030	4.241	3.986	27.34
FlowSE [17]	3.539	2.923	2.634	35.53
MAGE+CTF	4.449	4.301	4.076	25.27
MAGE+CTF+Corrector	4.517	4.301	4.141	23.45

Table 2. Performance on the noisy LibriSpeech test set. Reported metrics: DNSMOS (SIG/BAK/OVL, ↑) and WER (↓).

size of 192, and 4 heads attention. The language model is a reduced Qwen2.5-0.5B, where only the odd-numbered layers are kept (layer 1 is the first), and attention is configured in a non-autoregressive mode. LoRA is applied to `q_proj`, `v_proj`, `o_proj`, `up_proj`, and `down_proj` with $r = 16$, `lora_alpha = 32`, and `dropout = 0.1`. Training is performed with AdamW (learning rate and weight decay 1×10^{-4}), batch size 8, on a single RTX 4090 GPU.

The training corpus is constructed by augmenting clean speech from LibriSpeech [20] and the DNS Challenge [21] with noise from WHAM! [22] and DNS Challenge, and reverberation from OpenSLR28. The final dataset contains 512k four-second utterances at 16 kHz, with a composition of 50% noise-only, 30% noise+reverb, and 20% noise+reverb combined with resampling and spectrogram augmentation.

We evaluate MAGE against both discriminative and generative baselines using standard multiple metrics: SIG (signal distortion), BAK (background intrusiveness), and OVL (overall quality) from ITU-T P.835 [23], as well as Speaker Similarity (SSIM), computed as the cosine similarity between

System	Without Reverb		
	SIG↑	BAK↑	OVL↑
SSL Model (HuBERT) [25]	4.362	4.518	4.220
Transformer (8 layers)	3.414	4.002	3.276
Transformer (6 layers)	3.322	3.842	3.155
Band Aware	4.407	4.515	4.151

Table 3. Ablation study on different Speech Encoder

speaker embeddings from Wespeaker [24]. Results are reported under three conditions: with reverberation, without reverberation, and on real recordings. Higher values indicate better performance for all metrics. We also use the released ASR model² to compute WER, reflecting intelligibility gains.

4. RESULTS AND ANALYSIS

Table 1 compares MAGE with both discriminative and generative baselines on the DNS Challenge test set. Conventional discriminative models such as Conv-TasNet achieve strong SSIM scores but lag behind in perceptual quality (SIG/BAK/OVL). Generative methods including SGMSE, StoRM, and FlowSE offer improvements, yet remain limited in overall robustness. Mask-based approaches such as MaskSR-M and ANYENHANCE further advance BAK and OVL, but require considerably larger model sizes. In contrast, **MAGE** achieves competitive or superior performance with only 200M parameters. Notably, incorporating the coarse-to-fine (CTF) masking strategy yields substantial gains, especially in OVL (up to 4.235 without reverb and 3.787 on real recordings). The corrector module further stabilizes inference, and the combined CTF+Corrector configuration achieves the best SIG score of 4.580 without reverb. For reference, the *BigCodec Resyn. GT* row reports scores obtained by reconstructing ground-truth clean audio directly through the codec, representing the upper bound of codec fidelity, especially speaker similarity.

Table 2 reports results on the noisy LibriSpeech benchmark, highlighting downstream ASR performance. While prior generative models such as SGMSE and StoRM reduce word error rate (WER) relative to the noisy baseline, **MAGE with CTF and Corrector** achieves a WER of 23.45%, a relative improvement of over 5% absolute compared to SGMSE. These findings confirm that the enhanced audio produced by MAGE not only improves perceptual quality but also provides tangible benefits for recognition tasks.

Table 3 presents an ablation on the choice of speech encoder. SSL models, such as HuBERT, deliver strong performance but are computationally intensive. Lightweight

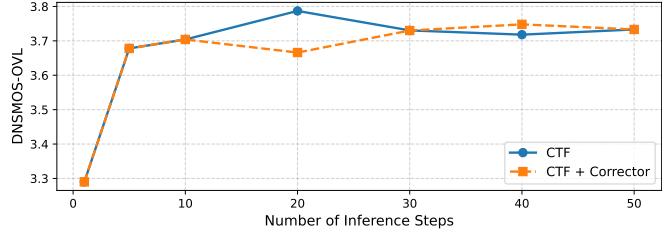


Fig. 3. Ablation study on the number of inference steps for CTF and CTF + Corrector. The overall performance is measured using DNSMOS-OVL on Real Recording DNS dataset

transformer encoders reduce complexity but exhibit degraded quality. However, the band-aware TF-GridNet encoder achieves a balance, reaching SIG 4.407, BAK 4.515, and OVL 4.151, comparable to HuBERT while being more efficient. This demonstrates that explicitly modeling cross-band spectral dependencies is more parameter-efficient than adopting SSL encoders. Fig. 3 presents an ablation on the number of inference steps for CTF and CTF+Corrector, evaluated on real recordings with DNSMOS-OVL. Performance improves rapidly up to 10 steps and stabilizes beyond 20 steps, with both variants maintaining strong quality across the range. CTF alone reaches a peak at 20 steps, while CTF+Corrector provides more stable performance at higher step counts, indicating that corrective refinement helps mitigate error accumulation without requiring excessive iterations.

Overall, these results show that MAGE effectively narrows the gap between perceptual quality and efficiency. The combination of CTF masking and corrective refinement delivers consistent improvements across datasets and metrics, establishing MAGE as a compact yet robust generative framework for speech enhancement. While MAGE achieves strong results, some limitations remain. Its dependence on simulated distortions may reduce generalization to real-world conditions, and the evaluation focuses mainly on DNSMOS instead of metrics such as PESQ or STOI.

5. CONCLUSION

In this work, we introduced MAGE, a lightweight difficulty-aware masked generative model for speech enhancement, finetuned from Qwen2.5-0.5B and coupled with BigCodec. Using a coarse-to-fine masking strategy with an auxiliary corrector, MAGE achieves strong perceptual quality and robustness with only 200M parameters. Experiments on DNS Challenge and LibriSpeech show that it outperforms discriminative and generative baselines in noisy and reverberant conditions. Future work includes extending to multilingual and streaming scenarios, joint training with ASR/TTS, and scaling to larger codecs and multimodal inputs, aiming to make MAGE a practical, generalizable solution for real-world deployment.

²https://huggingface.co/nvidia/stt_en_conformer_transducer_xlarge

6. ACKNOWLEDGEMENT

We acknowledge Ho Chi Minh City University of Technology (HCMUT), VNU-HCM for supporting this study. This work was also supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT, RS-2023-00212845).

7. REFERENCES

- [1] Julius Richter, Simon Welker, Jean-Marie Lemercier, et al., “Speech enhancement and dereverberation with diffusion-based generative models,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2023.
- [2] Jean-Marie Lemercier, Julius Richter, et al., “StoRM: A diffusion-based stochastic regeneration model for speech enhancement and dereverberation,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2023.
- [3] Tanel Pärnamaa and Ando Saabas, “Personalized speech enhancement without a separate speaker embedding model,” in *Proc. Interspeech*, 2024.
- [4] Kazuma Iwamoto, Tsubasa Ochiai, Marc Delcroix, et al., “How bad are artifacts?: Analyzing the impact of speech enhancement errors on asr,” in *Proc. Interspeech*, 2022.
- [5] Adam Sabra, Cyprian Wronka, Michelle Mao, and Samer Hijazi, “Seep: A speech enhancement-based curation pipeline for scalable acquisition of clean speech,” in *Proc. ICASSP*, 2024.
- [6] Xiaohuai Le, Hongsheng Chen, Kai Chen, and Jing Lu, “DPCRN: Dual-path convolution recurrent network for single channel speech enhancement,” in *Proc. Interspeech*, 2021.
- [7] Zhong-Qiu Wang, Samuele Cornell, Shukjae Choi, et al., “TF-GridNet: integrating full- and sub-band modeling for speech separation,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2023.
- [8] Peidong Wang, Ke Tan, and DeLiang Wang, “Bridging the gap between monaural speech enhancement and recognition with distortion-independent acoustic modeling,” in *Proc. Interspeech*, 2019.
- [9] Alexander H. Liu, Matthew Le, Apoorv Vyas, et al., “Generative pre-training for speech with flow matching,” in *Proc. ICLR*, 2024.
- [10] Ziqian Wang, Xinfu Zhu, Zihan Zhang, YuanJun Lv, Ning Jiang, Guoqing Zhao, and Lei Xie, “SELM: Speech enhancement using discrete tokens and language models,” in *Proc. ICASSP*, 2024.
- [11] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, et al., “WavLM: Large-scale self-supervised pre-training for full stack speech processing,” *IEEE Journal of Selected Topics in Signal Processing*, 2022.
- [12] Xu Li, Qirui Wang, and Xiaoyu Liu, “MaskSR: Masked language model for full-band speech restoration,” in *Proc. Interspeech*, 2024.
- [13] Rithesh Kumar, Prem Seetharaman, Alejandro Luebs, et al., “High-fidelity audio compression with improved rvqgan,” in *NeurIPS*, 2023.
- [14] Junan Zhang, Jing Yang, Zihao Fang, et al., “AnyEnhance: A unified generative model with prompt-guidance and self-critic for voice enhancement,” *IEEE Trans. on Audio, Speech, and Language Processing*, 2025.
- [15] Shengpeng Ji, Ziyue Jiang, Wen Wang, et al., “WavTok-encoder: an efficient acoustic discrete codec tokenizer for audio language modeling,” in *Proc. ICLR*, 2025.
- [16] Yi Luo and Nima Mesgarani, “Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2019.
- [17] Ziqian Wang and Zikai Liu and Xinfu Zhu and others, “FlowSE: Efficient and high-quality speech enhancement via flow matching,” in *Proc. Interspeech*, 2025.
- [18] Karen Sparck Jones, *A statistical interpretation of term specificity and its application in retrieval*, Taylor Graham Publishing, 1988.
- [19] Edward J Hu, Yelong Shen, Phillip Wallis, et al., “LoRA: Low-rank adaptation of large language models,” in *Proc. ICLR*, 2022.
- [20] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “LibriSpeech: An asr corpus based on public domain audio books,” in *Proc. ICASSP*, 2015.
- [21] Maximilian Strake, Bruno Defraene, Kristoff Fluyt, et al., “Interspeech 2020 deep noise suppression challenge: A fully convolutional recurrent network (fcrn) for joint dereverberation and denoising,” in *Proc. Interspeech*, 2020.
- [22] Gordon Wichern, Joe Antognini, Michael Flynn, et al., “WHAM!: Extending speech separation to noisy environments,” in *Proc. Interspeech*, 2019.
- [23] Chandan K A Reddy, Vishak Gopal, and Ross Cutler, “Dnsmos p.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors,” in *Proc. ICASSP*, 2022.
- [24] Hongji Wang, Chengdong Liang, Shuai Wang, et al., “Wespeaker: A research and production oriented speaker embedding learning toolkit,” in *Proc. ICASSP*, 2023.
- [25] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, et al., “HuBERT: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2021.