

# FASTENHANCER: SPEED-OPTIMIZED STREAMING NEURAL SPEECH ENHANCEMENT

Sunghwan Ahn    Jinmo Han    Beom Jun Woo    Nam Soo Kim

Department of Electrical and Computer Engineering and INMC  
Seoul National University, Seoul, South Korea  
{shahn, jmhan, bjwoo}@hi.snu.ac.kr, nkim@snu.ac.kr

## ABSTRACT

Streaming speech enhancement is a crucial task for real-time applications such as online meetings, smart home appliances, and hearing aids. Deep neural network-based approaches achieve exceptional performance while demanding substantial computational resources. Although recent neural speech enhancement models have succeeded in reducing the number of parameters and multiply-accumulate operations, their sophisticated architectures often introduce significant processing latency on common hardware. In this work, we propose *FastEnhancer*, a streaming neural speech enhancement model designed explicitly to minimize real-world latency. It features a simple encoder-decoder structure with efficient RNNFormer blocks. Evaluations on various objective metrics show that *FastEnhancer* achieves state-of-the-art speech quality and intelligibility while simultaneously demonstrating the fastest processing speed on a single CPU thread. Code and pre-trained weights are publicly available<sup>1</sup>.

**Index Terms**— Speech enhancement, noise suppression, low latency, neural network.

## 1. INTRODUCTION

Monaural speech enhancement aims to estimate a clean speech signal from a single-channel mixture corrupted by noise. Deploying enhancement algorithms in real-time applications like live communication or smart home assistants presents a significant challenge, as it demands not only high-quality output but also fast processing speed on resource-limited devices.

Early traditional approaches like spectral subtraction or Wiener filtering [1] focused on estimating spectral amplitude of clean speech. However, they rely on noise stationarity assumption and often fail in real-world environments. Other techniques exploited the periodicity of voiced speech [2, 3], but they require accurate pitch information. Furthermore, while many of these traditional techniques improve speech quality, they often fail to enhance speech intelligibility [4]. In contrast, deep neural network (DNN)-based models have demonstrated a remarkable ability to handle diverse noise types, significantly surpassing traditional methods [5]. Nonetheless, the high computational demands of these models result in significant processing latency on resource-constrained devices, hindering their deployment on real-world products.

In response, many recent studies have attempted to decrease the number of parameters or multiply-accumulate (MAC) operations of neural speech enhancement models. These attempts include dual-path recurrent neural network (DPRNN) structures [6–8], elaborate




Fig. 1. RTF vs. SISDR and RTF vs. STOI of various models.

encoder-decoder designs [9], and sub-band decomposition strategies [10–13]. Although these approaches successfully reduce the number of parameters and MAC operations, we observe that they frequently increase memory operations and architectural complexity, which slows down the actual inference speed (Section 2.2). Furthermore, despite the widespread adoption of DPRNN, it remains an open question whether it is the optimal structure, particularly concerning its assumption of a sequential relationship between frequency bands (Section 2.1).

To address these limitations, we focus on developing a neural speech enhancement model that satisfies the following properties:

- **High quality and intelligibility:** The enhanced speech sounds clear and natural (high quality), with the words being easy to understand (high intelligibility).
- **Streaming:** The model can process input spectrograms in a frame-by-frame manner.
- **Fast processing speed:** The model achieves a low real-time factor (RTF) on resource-constrained environments, where RTF is the ratio of processing time to the input audio duration. We consider a single CPU thread as a key benchmark for such environments, as many real-world applications lack dedicated hardware accelerators like GPUs or NPUs.

To meet these objectives, we introduce *FastEnhancers*, a family of streaming neural speech enhancement models optimized for speed, quality, and intelligibility. Our key contributions are:

- We demonstrate that a streamlined encoder-decoder structure, in contrast to the complex, sub-band-focused designs of prior works, improves processing speed without sacrificing performance.
- We propose an RNNFormer block, an efficient dual-path block that combines a time-axis recurrent neural network (RNN) for

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2025-00554289)

<sup>1</sup><https://github.com/aask1357/fastenhancer>

low-complexity temporal modeling with a frequency-axis transformer for effectively capturing global, non-sequential relationships across frequency bands.

- We evaluate *FastEnhancers* against existing models using various metrics for speech quality and intelligibility on the VCTK-Demand dataset [14]. Experimental results show that *FastEnhancers* achieve state-of-the-art performances and the lowest RTFs among all evaluated models. By configuring models with various speeds, we establish a new Pareto frontier for the trade-off between the model performance and inference speed (Fig. 1).

## 2. RELATED WORKS

### 2.1. Dual-Path Networks

The original DPRNN [6] architecture was proposed for waveform-domain speech separation. It divides an input sequence of  $C$  channels and  $L$  samples into short chunks with a chunk size of  $K$  and a hop size of  $P$ , resulting in a three-dimensional tensor of shape  $[C, L/P, K]$ . The model then iteratively performs an intra-chunk RNN along the third dimension and an inter-chunk RNN along the second dimension. This effectively models long sequences and outperforms previous waveform-domain approaches. DPRNN has also been widely adopted for time-frequency (TF)-domain speech enhancement [7–12]. It processes a spectrogram by alternating between a time RNN that models temporal dependencies and a frequency RNN that models dependencies across frequency bands. However, applying an RNN to the frequency axis assumes that frequency bands have a sequential relationship. This is suboptimal because spectral relationships (e.g. harmonics) exhibit global rather than sequential dependencies. We argue that a model should be capable of handling all frequency interactions concurrently.

Some prior works have attempted to replace the RNNs in dual-path networks with other powerful modules like Mamba and xLSTM [15, 16]. However, these models are designed for non-streaming, large-scale scenarios, and still impose a sequential relationship on the frequency axis. Meanwhile, others have explored dual-path transformers (DPT) [17, 18]. Transformers are intuitively well-suited for modeling global information across frequency bands. On the other hand, a time-axis transformer is problematic for low-latency streaming. It requires caching previous key and value frames, demanding additional memory operations. Moreover, to improve speech quality, the number of cache frames should be increased [19], requiring more memory and arithmetic operations.

In this paper, we propose RNNFormer, which is composed of a time RNN and a frequency transformer (Section 3). The time RNN maintains the low-complexity streaming nature of our model, while the frequency transformer captures global information across all frequency bands. We experimentally prove that RNNFormers achieve a superior RTF-performance curve compared to DPRNNs and DPTs (Section 5.2).

### 2.2. Processing Speed of Speech Enhancement Models

Many prior works have focused on reducing the number of parameters and MACs to improve model efficiency. DPCRNN [7] and TRU-Net [8] introduced a TF-domain U-Net structure with a convolutional encoder, decoder, and DPRNNs. Building upon this, more elaborate methods such as sub-band processing and enhanced DPRNN structures have emerged.

BSRNN [13] introduced a band-split encoder and decoder. FSPEN [10] proposed a sub-band encoder and decoder and a time

RNN with a path extension. GTCRN [9] and LiSenNet [11] utilized sub-band feature extraction, grouped DPRNNs, and channel split-concatenate-shuffle operations. While these methods successfully reduced the number of parameters and MACs, their increased architectural complexity and memory operations paradoxically led to higher RTFs (Section 5.1).

In contrast to these approaches, we directly aim to minimize the RTF for practical scenarios. We experimentally prove that a simple network design – without complex sub-band operations or RNN grouping – is sufficient to obtain a powerful and fast model.

## 3. MODEL ARCHITECTURE

This section details the architecture of *FastEnhancer* as shown in Fig. 3. We start by describing the input and output. We then present convolutional and normalization layers characterized by low latency for streaming inference. Finally, using those building blocks, we explain a simple encoder-decoder structure and our novel RNNFormer blocks. Note that we apply weight normalization [20] and use the SiLU activation function [21] across the entire model.

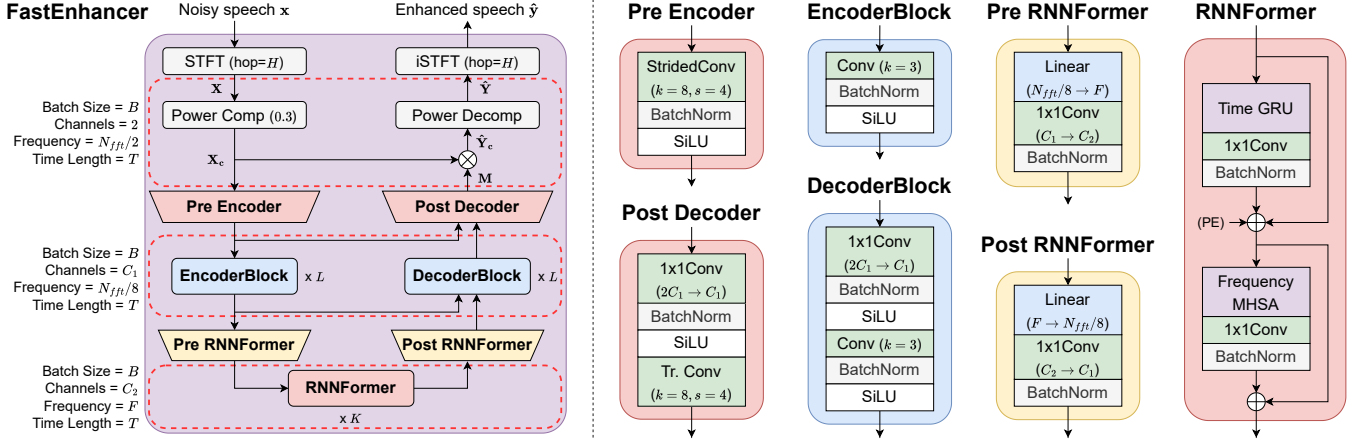
**Input and Output** – An input noisy speech  $\mathbf{x}$  is first transformed into a complex spectrogram  $\mathbf{X} = \text{STFT}(\mathbf{x})$  with a Fourier transform size of  $N_{fft} = 512$  and a hop size of  $H$ . To mimic the dynamic compression nature of the human ear, we apply power compression as  $\mathbf{X}_c = |\mathbf{X}|^c \cdot e^{j\angle\mathbf{X}}$ , where  $c = 0.3$ . We then interpret  $\mathbf{X}_c$  as a tensor with two channels (real and imaginary part) and  $N_{fft}/2$  frequencies (for simplicity, we discard the highest frequency bin). Given  $\mathbf{X}_c$ , *FastEnhancer* predicts a two-channel mask  $\mathbf{M}$  and generates a power-compressed enhanced spectrogram  $\hat{\mathbf{Y}}_c = \mathbf{M} \odot \mathbf{X}_c$ . After power decompression  $\hat{\mathbf{Y}} = |\hat{\mathbf{Y}}_c|^{1/c} \cdot e^{j\angle\hat{\mathbf{Y}}_c}$ , the highest frequency bin is zero-padded and the final enhanced speech  $\hat{\mathbf{y}} = \text{iSTFT}(\hat{\mathbf{Y}})$  is synthesized.  $\hat{\mathbf{y}}$  is then compared to a target clean speech  $\mathbf{y}$ .

**Convolutional Layers** – Causal convolutional layers with a time-axis kernel size greater than one introduce latency, because they require caching past frames to compute the current output frame. For a time-axis kernel size of  $k$ , a buffer of  $k - 1$  past frames is stored and concatenated with the current frame before the convolution operation. The latest  $k - 1$  frames are then stored in memory for the next computation. These memory operations significantly increase processing time (Section 5.2). To ensure minimum latency, *FastEnhancer* exclusively uses convolutions with a time-axis kernel size of one. Throughout this paper, all specified kernel sizes and strides for convolutions refer to the frequency axis.

**Normalization Layers** – Speech enhancement models typically use either batch normalization (BN) [22] or layer normalization (LN) [23]. While both methods introduce negligible number of parameters and MACs, BN can be converted into a simple multiplication-addition operation and fused into adjacent layers at inference time. In contrast, LN cannot be fused in the same manner, increasing the RTF (Section 5.2). For this reason, we use BN layers throughout our model and fuse them into neighboring layers at inference time.

**Encoder and Decoder** – The encoder starts with a pre-encoder network containing a strided convolution that reduces the frequency resolution from  $N_{fft}/2$  to  $N_{fft}/8$  while enlarging the channels from 2 to  $C_1$ . It is followed by a stack of  $L$  encoder blocks. The decoder contains  $L$  decoder blocks, followed by a post-decoder network that restores the frequency resolution using a transposed convolution. Skip connections are used to link the encoder and decoder.

After the encoder blocks and before the decoder blocks, we append pre-RNNFormer and post-RNNFormer networks to bridge to the RNNFormer blocks. The pre-RNNFormer network is composed



**Fig. 2.** *FastEnhancer* model architecture. Conv( $k$ ), StridedConv( $k, s$ ), and Tr. Conv( $k, s$ ) denote a convolution, a strided convolution, and a transposed convolution respectively, with a kernel size of  $k$  and a stride of  $s$ .  $1\times 1\text{Conv}(A \rightarrow B)$  represents a  $1\times 1$  convolution with  $A$  input channels and  $B$  output channels. Linear( $A \rightarrow B$ ) represents a linear layer with  $A$  input features and  $B$  output features.

of a linear layer that reduces frequency resolution from  $N_{fft}/8$  to  $F$ , followed by a convolution that reduces channels from  $C_1$  to  $C_2$ . The post-RNNFormer network performs the reverse operations. The weight of the linear layer in pre-RNNFormer is initialized as a linear filterbank, and the weight of the linear layer in post-RNNFormer is initialized as a linear interpolation. We empirically found that training these weights offered no performance benefit over keeping them fixed. Therefore, we fixed those weights in all experiments.

**RNNFormer** – *FastEnhancer* contains  $K$  RNNFormer blocks. Each block consists of two sub-blocks: a time-axis Gated Recurrent Unit (GRU) [24] block, and a frequency-axis Multi-Head Self-Attention (MHSA) [25] block. This hybrid design is critical for low-latency and high-performance. For the time axis, a GRU is well-suited for streaming inference with minimal latency, whereas an MHSA requires a key-value caching that increases delay (as described in Section 2.1). For the frequency axis, an MHSA is ideal as it can model global relationships and doesn’t require caching.

The GRU block contains a sequence of a unidirectional GRU, a convolution, a BN layer, and a residual skip connection. The MHSA block follows an identical structure but replaces the GRU layer with an MHSA with four heads. We add a trainable positional encoding to the input of the first MHSA block.

#### 4. EXPERIMENTAL SETUP

We train our model using a composite loss function:

$$\mathcal{L} = \lambda_1 \cdot \mathcal{L}_{mag} + \lambda_2 \cdot \mathcal{L}_{comp} + \lambda_3 \cdot \mathcal{L}_{con} + \lambda_4 \cdot \mathcal{L}_{wav} + \lambda_5 \cdot \mathcal{L}_{pesq}, \quad (1)$$

where  $\lambda_1 = 0.3$ ,  $\lambda_2 = 0.2$ ,  $\lambda_3 = 0.3$ ,  $\lambda_4 = 0.2$ , and  $\lambda_5 = 0.001$ .

- **Magnitude loss** calculates the mean squared error (MSE) between the magnitudes of  $\hat{\mathbf{Y}}_c$  and  $\mathbf{Y}_c$ :

$$\mathcal{L}_{mag} = \left\| |\mathbf{Y}_c| - |\hat{\mathbf{Y}}_c| \right\|_2^2, \quad (2)$$

where  $\mathbf{Y} = \text{STFT}(\mathbf{y})$  is a spectrogram of a clean speech  $\mathbf{y}$  and  $\mathbf{Y}_c = |\mathbf{Y}|^c \cdot e^{j\angle \mathbf{Y}}$  is a compressed spectrogram of  $\mathbf{y}$ .

- **Complex spectrogram loss** interprets two complex spectrograms  $\hat{\mathbf{Y}}_c$  and  $\mathbf{Y}_c$  as two-channel real-valued tensors, and calculates MSE:

$$\mathcal{L}_{comp} = \|\mathbf{Y}_c - \hat{\mathbf{Y}}_c\|_2^2. \quad (3)$$

**Table 1.** Model configurations.

Size	$H$	$L$	$K$	$C_1$	$C_2$	$F$
Tiny (T)	256	2	2	24	20	16
Base (B)	256	2	3	48	36	24
Small (S)	256	3	3	64	48	36
Medium (M)	160	3	4	96	72	48
Large (L)	100	4	5	128	96	64

- **Consistency loss** is defined as:

$$\mathcal{L}_{con} = \|\mathbf{Y}_c - \text{STFT}_c(\hat{\mathbf{y}})\|_2^2, \quad (4)$$

where  $\text{STFT}_c(\mathbf{y}) = |\text{STFT}(\mathbf{y})|^c \cdot e^{j\angle \text{STFT}(\mathbf{y})}$ .

- **Waveform loss** is a mean absolute error in the waveform domain:

$$\mathcal{L}_{wav} = \|\mathbf{y} - \hat{\mathbf{y}}\|_1. \quad (5)$$

- **PESQ loss**  $\mathcal{L}_{pesq}$  is a differentiable PESQ loss [26]. Optimizing a model solely on the PESQ metric can degrade subjective quality [27]. To prevent this, we use a sufficiently small weighting factor ( $\lambda_5 = 0.001$ ). We observed that including  $\mathcal{L}_{pesq}$  improves not only PESQ but also other objective metrics.

We built *FastEnhancers* with five different sizes (Table 1) and compared them with other low-complexity models. VTCK-Demand dataset [14] was used for training and evaluation, with all audio files downsampled from 48kHz to 16kHz. During training, two-second-long segments were randomly extracted. We used the AdamP [28] optimizer with an initial learning rate of 0.002, a weight decay of 0.01, and a batch size of 64. We applied a cosine annealing learning rate scheduler with a 500-step warm-up [29].

We utilized a variety of objective metrics for evaluation. For speech quality, we measured DNSMOS [30] [31], SCOREQ (natural speech, full reference mode) [32], scale-invariant signal-to-distortion ratio (SISDR) [33], and PESQ (P.862.2 without Corrigendum 2) [34]. For speech intelligibility, we measured STOI [35], ESTOI [36], and word error rate (WER) using Whisper-large-v3-turbo [37]. We trained each model five times with five different random seeds and report the average scores.

We exported every model to ONNXRuntime [38] and measured RTFs on a single thread of a server CPU (Intel Xeon Gold 6248R) and a laptop CPU (Apple M1, MacBook Air) under a streaming condition. RTFs in Fig. 1 and Fig. 3 are measured on Xeon.

**Table 2.** Performance on Voicebank-Demand.

Model	Para. (K)	MACs	RTF (Xeon)	RTF (M1)	DNSMOS (P.808)	DNSMOS (P.835)			SCOREQ	SISDR	PESQ	STOI	ESTOI	WER
						SIG	BAK	OVL						
GTCRN <sup>a</sup> [9]	24	40M	0.060	0.042	3.43	3.36	4.02	3.08	0.330	18.8	2.87	0.940	0.848	3.6
LiSenNet <sup>b</sup> [11]	37	56M	-	-	3.34	3.30	3.90	2.98	0.425	13.5	3.08	0.938	0.842	3.7
LiSenNet <sup>c</sup> [11]	37	56M	0.034	0.028	3.42	3.34	<b>4.03</b>	3.07	0.335	18.5	2.98	0.941	0.851	3.4
FSPEN <sup>d</sup> [10]	79	64M	0.046	0.038	3.40	3.33	4.00	3.05	0.324	18.4	3.00	0.942	0.850	3.6
BSRNN <sup>d</sup> [12]	334	245M	0.059	0.062	3.44	3.36	4.00	3.07	0.303	18.9	3.06	0.942	0.855	3.4
<i>FastEnhancer</i> -B	92	262M	<b>0.022</b>	<b>0.026</b>	<b>3.47</b>	<b>3.38</b>	4.02	<b>3.10</b>	<b>0.285</b>	<b>19.0</b>	<b>3.13</b>	<b>0.945</b>	<b>0.861</b>	<b>3.2</b>
<i>FastEnhancer</i> -T	22	55M	<b>0.012</b>	<b>0.013</b>	3.42	3.34	4.01	3.06	0.334	18.6	2.99	0.940	0.850	3.6
<i>FastEnhancer</i> -B	92	262M	0.022	0.026	3.47	3.38	4.02	3.10	0.285	19.0	3.13	0.945	0.861	3.2
<i>FastEnhancer</i> -S	195	664M	0.034	0.048	3.49	3.40	4.03	3.12	0.265	19.2	3.19	0.947	0.866	3.2
<i>FastEnhancer</i> -M	492	2.9G	0.101	0.173	3.48	3.39	4.02	3.11	0.243	19.4	3.24	0.950	0.873	<b>2.8</b>
<i>FastEnhancer</i> -L	1105	11G	0.313	0.632	<b>3.53</b>	<b>3.44</b>	<b>4.04</b>	<b>3.16</b>	<b>0.239</b>	<b>19.6</b>	<b>3.26</b>	<b>0.952</b>	<b>0.877</b>	3.1

<sup>a</sup>Evaluated using the official checkpoint.

<sup>b</sup>Trained using the official training code. Not streamable because of input normalization and griffin-lim. Thus, RTFs are not reported.

<sup>c</sup>To make the model streamable, input normalization and griffin-lim are removed. Trained following the experimental setup in Section 4.

<sup>d</sup>Re-implemented and trained following the experimental setup in Section 4 for a fair comparison.

**Table 3.** Ablation study for convolutions and normalization layers.

Ablation	Para. (K)	RTF (Xeon)	RTF (M1)	SISDR	STOI
<i>FastEnhancer</i> -B	92	<b>0.022</b>	<b>0.026</b>	<b>19.0</b>	<b>94.5</b>
<b>Convolution</b>					
Time-axis kernel size 3	187	0.028	0.037	<b>19.0</b>	<b>94.5</b>
<b>Replace BatchNorm to LayerNorm</b>	92	0.028	0.029	18.9	<b>94.5</b>

## 5. RESULTS

### 5.1. Performance Comparisons


In Table 2, *FastEnhancer*-B outperforms all baseline models on almost every metric. While having relatively large number of parameters and MACs, its speed-optimized architecture results in the lowest RTFs. Our smallest model, *FastEnhancer*-T, achieves remarkably low RTFs of 0.012 on Xeon and 0.013 on M1 while maintaining comparable performance to baseline models. As the model size increases, overall performance improves at the cost of increased latency. As a side observation, compared to *FastEnhancer*-S, *FastEnhancer*-M shows slightly degraded DNSMOS scores despite improved scores on every other metric. This indicates that relying on a limited set of objective metrics can lead to misleading conclusions.

### 5.2. Ablation Studies

We performed ablation studies to validate the effectiveness of our architectural design. First, we explored the impact of convolution kernel size. Starting with *FastEnhancer*-B, we modified every convolution whose frequency-axis kernel size is 3 in encoder and decoder blocks, enlarging their time-axis kernel size from 1 to 3. This necessitated caching operations and significantly increased RTFs, but no improvement in quality or intelligibility was observed (Table 3).

Second, to investigate the effect of normalization layers, we replaced all BNs with LNs. While this maintained speech quality and intelligibility, it increased RTFs, despite having the same number of parameters and MACs. This result confirms that using BNs and fusing them into adjacent layers is the most efficient option.

Third, we created Dual-Path RNNs by replacing the frequency-axis MHSA blocks with frequency-axis GRU blocks. While STOI scores were similar, DPRNNs clearly degraded SISDR compared to

**Fig. 3.** Ablation study for RNNFormer.

the proposed RNNFormers (Fig. 3). This demonstrates the superiority of self-attention over a recurrent approach for frequency modeling.

Finally, we constructed Dual-Path Transformers by replacing the time-axis GRU blocks with time-axis MHSA blocks with a look-behind of 31 frames (corresponding to 0.5 seconds). Although DPTs have fewer parameters than RNNFormers, they require key-value caching, resulting in significantly increased RTFs. This result suggests that a recurrent approach is a more efficient choice than self-attention for temporal modeling in a streaming scenario.

## 6. CONCLUSION

In this paper, we introduced *FastEnhancer*, a streaming neural speech enhancement model designed for low latency. We made careful architectural choices, including convolutional layers with time-axis kernel size of 1, normalization layers that can be fused into other layers at inference, time RNNs, and frequency transformers. Our results demonstrate that *FastEnhancer* significantly outperforms prior works in terms of speech quality and intelligibility while achieving the fastest inference speed on a single CPU thread. This makes our model a highly practical and fast solution for real-world, on-device deployment.

## 7. REFERENCES

- [1] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Trans. ASSP*, vol. 27, no. 2, pp. 113–120, 1979.
- [2] V. C. Shields, "Separation of added speech signals by digital comb filtering," M.S. thesis, Dep. Elec. Eng., Massachusetts Inst. Tech., Cambridge, 1970.
- [3] R. Frazier, S. Samsam, L. Braida, and A. Oppenheim, "Enhancement of speech by adaptive filtering," in *Proc. IEEE ICASSP*, 1976, vol. 1, pp. 251–253.
- [4] J.S. Lim and A.V. Oppenheim, "Enhancement and bandwidth compression of noisy speech," *Proc. IEEE*, vol. 67, no. 12, pp. 1586–1604, 1979.
- [5] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Trans. ASLP*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [6] C. Li et al., "Dual-path RNN for long recording speech separation," in *Proc. IEEE SLT Workshop*, 2021, pp. 865–872.
- [7] X. Le, H. Chen, K. Chen, and J. Lu, "DPCRN: Dual-path convolution recurrent network for single channel speech enhancement," in *Proc. Interspeech*, 2021, pp. 2811–2815.
- [8] H.-S. Choi, S. Park, J. H. Lee, H. Heo, D. Jeon, and K. Lee, "Real-time denoising and dereverberation with tiny recurrent U-Net," in *Proc. IEEE ICASSP*, 2021, pp. 5789–5793.
- [9] X. Rong, T. Sun, X. Zhang, Y. Hu, C. Zhu, and J. Lu, "GTCRN: A speech enhancement model requiring ultralow computational resources," in *Proc. IEEE ICASSP*, 2024, pp. 971–975.
- [10] L. Yang, W. Liu, R. Meng, G. Lee, S. Baek, and H.-G. Moon, "Fspen: an ultra-lightweight network for real time speech enhancement," in *Proc. IEEE ICASSP*, 2024, pp. 10671–10675.
- [11] H. Yan, J. Zhang, C. Fan, Y. Zhou, and P. Liu, "LiSenNet: Lightweight sub-band and dual-path modeling for real-time speech enhancement," in *Proc. IEEE ICASSP*, 2025, pp. 1–5.
- [12] J. Yu, H. Chen, Y. Luo, R. Gu, and C. Weng, "High fidelity speech enhancement with band-split RNN," in *Proc. Interspeech*, 2023, pp. 2483–2487.
- [13] Yi Luo and Jianwei Yu, "Music source separation with band-split RNN," *IEEE/ACM Trans. ASLP*, vol. 31, pp. 1893–1901, 2023.
- [14] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Investigating RNN-based speech enhancement methods for noise-robust text-to-speech," in *Proc. ISCA Speech Synth. Workshop*, 2016, pp. 146–152.
- [15] R. Chao et al., "An investigation of incorporating Mamba for speech enhancement," in *Proc. IEEE SLT Workshop*, 2024, pp. 302–308.
- [16] N. L. Kühne, J. Østergaard, J. Jensen, and Z. -H. Tan, "xLSTM-SENet: xLSTM for single-channel speech enhancement," in *Proc. Interspeech*, 2025, pp. 5148–5152.
- [17] J. Chen, Q. Mao, and D. Liu, "Dual-path transformer network: Direct context-aware modeling for end-to-end monaural speech separation," in *Proc. Interspeech*, 2020, pp. 2642–2646.
- [18] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, "Attention is all you need in speech separation," in *Proc. IEEE ICASSP*, 2021, pp. 21–25.
- [19] S. H. Bae, S. W. Chae, Y. Kim, K. Lee, H. Lim, and L.-H. Kim, "Streaming dual-path transformer for speech enhancement," in *Proc. Interspeech*, 2023, pp. 824–828.
- [20] T. Salimans and D. P. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Adv. Neural Inf. Process. Syst.*, 2016, pp. 901–909.
- [21] D. Hendrycks and K. Gimpel, "Gaussian error linear units (GELUs)," *arXiv:1606.08415*, 2016.
- [22] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [23] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv:1607.06450*, 2016.
- [24] K. Cho et al., "Learning phrase representations using RNN encoder–decoder for statistical machine translation," in *Proc. Conf. Empir. Methods Nat. Lang. Process.*, 2014, pp. 1724–1734.
- [25] A. Vaswani et al., "Attention is all you need," in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [26] L. Schmidt, N. Werner, and N. Peters, "torch-pesq," [Online]. Available: <https://github.com/audiolabs/torch-pesq>. [Accessed: Jul. 6, 2025].
- [27] D. de Oliveira, S. Welker, J. Richter, and T. Gerkmann, "The PESQetarian: On the relevance of Goodhart's law for speech enhancement," in *Proc. Interspeech*, 2024, pp. 3854–3858.
- [28] B. Heo et al., "AdamP: Slowing down the slowdown for momentum optimizers on scale-invariant weights," in *Int. Conf. Learn. Representations*, 2021.
- [29] I. Loshchilov and F. Hutter, "SGDR: Stochastic gradient descent with warm restarts," in *Int. Conf. on Learn. Representations*, 2017.
- [30] C. K. A. Reddy, V. Gopal, and R. Cutler, "Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *Proc. IEEE ICASSP*, 2021, pp. 6493–6497.
- [31] C. K. A. Reddy, V. Gopal, and R. Cutler, "Dnsmos P.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *Proc. IEEE ICASSP*, 2022, pp. 886–890.
- [32] A. Ragano, J. Skoglund, and A. Hines, "SCOREQ: Speech quality assessment with contrastive regression," in *Adv. Neural Inf. Process. Syst.*, 2024, pp. 105702–105729.
- [33] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR – half-baked or well done?," in *Proc. IEEE ICASSP*, 2019, pp. 626–630.
- [34] ITU-T, "Wideband extension to Recommendation P.862 for the assessment of wideband telephone networks and speech codecs," Rec. P.862.2, International Telecommunication Union, Geneva, Switzerland, 2007.
- [35] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time–frequency weighted noisy speech," *IEEE Trans. ASLP*, vol. 19, no. 7, pp. 2125–2136, Sept. 2011.
- [36] J. Jensen and C. H. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Trans. ASLP*, vol. 24, no. 11, pp. 2009–2022, Nov. 2016.
- [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. Int. Conf. Mach. Learn.*, 2023, pp. 28492–28518.
- [38] ONNX Runtime developers, "ONNX runtime," [Online]. Available: <https://onnxruntime.ai>. [Accessed: Jul. 6, 2025].