

BSDB-Net: Band-Split Dual-Branch Network with Selective State Spaces Mechanism for Monaural Speech Enhancement

Cunhang Fan¹, Enrui Liu¹, Andong Li^{2,3*}, Jianhua Tao⁴, Jian Zhou¹, Jiahao Li¹, Chengshi Zheng^{2,3}, Zhao Lv^{1*}

¹Anhui Province Key Laboratory of Multimodal Cognitive Computation, School of Computer Science and Technology, Anhui University, Hefei, China

²Key Laboratory of Noise and Vibration Research, Institute of Acoustics Chinese Academy of Sciences, Beijing, China

³University of Chinese Academy of Sciences, Beijing, China

⁴Department of Automation, Tsinghua University, Beijing, China

cunhang.fan@ahu.edu.cn, e23201090@stu.ahu.edu.cn, liandong@mail.ioa.ac.cn, jhtao@nlpr.ia.ac.cn, jzhou@ahu.edu.cn, e23301209@stu.ahu.edu.cn, cszheng@mail.ioa.ac.cn, kjlz@ahu.edu.cn

Abstract

Although the complex spectrum-based speech enhancement (SE) methods have achieved significant performance, coupling amplitude and phase can lead to a compensation effect, where amplitude information is sacrificed to compensate for the phase that is harmful to SE. In addition, to further improve the performance of SE, many modules are stacked onto SE, resulting in increased model complexity that limits the application of SE. To address these problems, we proposed a dual-path network based on compressed frequency using Mamba. First, we extract amplitude and phase information through parallel dual branches. This approach leverages structured complex spectra to implicitly capture phase information and solves the compensation effect by decoupling amplitude and phase, and the network incorporates an interaction module to suppress unnecessary parts and recover missing components from the other branch. Second, to reduce network complexity, the network introduces a band-split strategy to compress the frequency dimension. To further reduce complexity while maintaining good performance, we designed a Mamba-based module that models the time and frequency dimensions under linear complexity. Finally, compared to baselines, our model achieves an average 8.3 times reduction in computational complexity while maintaining superior performance. Furthermore, it achieves a 25 times reduction in complexity compared to transformer-based models.

Introduction

In the realm of audio signal processing, speech enhancement (SE) is regarded as a fundamental technique to recover the clean speech from noisy environments. The degradation of speech quality by background noise is not only perceptually bothersome but also significantly impairs the performance of automatic speech recognition (ASR) (Chan 2016; Liu et al. 2018). Also, SE is indispensable in smart devices, vehicular systems, and home automation. With the burgeoning prevalence of online conferencing, the demand for real-time SE solutions has surged, underscoring the necessity for

techniques that are both effective and computationally efficient (Pandey and Wang 2019).

Existing speech enhancement methods can be roughly categorized into two classes, namely in the time domain (Luo and Mesgarani 2019; Yin et al. 2020) and in the time-frequency (T-F) domain. This paper mainly focuses on the latter. Initially, SE methods focused on magnitude spectrum enhancement with the phase kept unaltered (Tan and Wang 2019). Subsequent researches have revealed the pivotal role of phase (Yin et al. 2020), which inspire many studies to model complex spectra to better recover phase information. Masking-based methods, such as the complex ratio mask (Williamson, Wang, and Wang 2015), have been recognized for their ability to modulate both the real (R) and imaginary (I) components of the noisy complex spectra, surpassing the performance of traditional masks like ideal binary mask (Divenyi 2004) and ideal ratio mask (Hummerstone, Stokes, and Brookes 2014). Besides, complex spectral mapping (CSM) has also been introduced to reconstruct the target RI components directly (Tan and Wang 2019). However, this direct mapping can lead to a compensation effect between phase and magnitude (Wang, Wichern, and Le Roux 2021). Recent advancements have seen the application of decoupling concepts in audio processing (Alex et al. 2024). Some multi-stage methods have been proposed to address SE problems. These methods decouple the original mapping problem into two stages: predicting the magnitude spectrum at first and refining the complex spectrum through residual learning in the second stage. This approach partially mitigates the implicit compensation effect between phase and magnitude (Li et al. 2021a). Nonetheless, the sequential nature of these methods can be limiting, as the performance of later stages is heavily dependent on the output of earlier stages (Li et al. 2022b).

Despite the promising results of previous works, their high computational complexity can impede the practical application of SE models. First, as a front-end task, the impact of SE on downstream tasks, such as ASR, should be considered. Moreover, effective deployment in edge or resource-constrained environments, like online meetings and real-time communications, necessitates extremely low-

*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

complexity SE methods. Strategies such as frequency band division (Yu and Luo 2023), have been introduced to reduce frequency modeling overhead. Additionally, the computational complexity of sequence modeling can be substantial, with some transformer-based models reaching hundreds of gigaflops (G/s) due to their quadratic attention mechanisms (Lu, Ai, and Ling 2023). Recently, State Space Models (Hamilton 1994), such as HIPPO (Gu et al. 2020) and the Structured State Space Model (Gu, Goel, and Ré 2021), have demonstrated promising performance with reduced complexity. The Selective SSM (Gu and Dao 2023), in particular, enables the establishment of long-range dependencies with linear computational complexity, making it suitable for sequence modeling.

In light of these challenges, we propose the Band-Split Dual-branch Network (BSDB-Net) for the monaural SE task. This approach introduces a decomposition strategy to enhance the magnitude and complex spectra in parallel. Specifically, we devised the Magnitude Enhancement Network (MEN) to suppress the noise components without interfering with phase information, and the Complex Spectral Enhancement Network (CEN) to restore phase information implicitly, complementing the magnitude features extracted by MEN and mitigating the compensation effect. Besides, to effectively decrease the overhead for frequency modeling, a frequency band-split strategy is adopted to compress the frequency axis. To better decrease the complexity for sequence modeling, the Mamba structure is introduced to reduce complexity in temporal and frequency sequential relations.

The major contributions of this paper are summarized as follows:

- We propose a dual-branch parallel speech enhancement network that implicitly extracts phase information while ensuring the independence of amplitude information.
- We adopt a frequency band-splitting strategy and Mamba-based sequence modeling modules, significantly reducing computational complexity.
- Comprehensive experiments on two public datasets demonstrate that BSDB-Net achieves an average **8.3** times reduction in computational complexity while maintaining comparable performance.

Related Work

Multi-stage Methods: Due to the lack of prior information, the performance of the single-stage SE pipeline is often heavily limited in complicated acoustic scenarios. In contrast, for the multi-stage pipeline, the original mapping problem is usually decomposed into several separate sub-tasks to enable the learning progressively (Li et al. 2022a). In DTLN (Westhausen and Meyer 2020), the authors proposed a stacked dual signal transformation network. In Full-SubNet (Hao et al. 2021), the idea of combing sub-band and full-band was proposed to restore the spectrum. CompNet (Fan et al. 2023) combined the time domain and T-F domain, using a cross-domain complementary approach to optimize the speech enhancement network. CTS-Net (Li et al. 2021a) utilized a two-stage paradigm to supplement phase

information on the basis of extracting the magnitude spectrum. TaylorNet (Li et al. 2022a) proposed an end-to-end framework to simulate the 0th-order and high-order items of Taylor-unfolding. FAF-Net (Yue et al. 2022) proposed reference-based speech enhancement via a feature alignment and fusion network.

Sequence Modeling: The exploration of long and short-term temporal dependencies within speech signals prompted the adoption of Recurrent Neural Networks (RNNs (Zaremba, Sutskever, and Vinyals 2014)) to better capture contextual relations. To address the issue of gradient explosion, Long Short-Term Memory (LSTM (Chen et al. 2016)) networks were introduced. The temporal convolutional module (TCM) (Bai, Kolter, and Koltun 2018) introduced in TCM was found to be more effective in time series modeling than LSTM. Recently, a compressed version of TCM called the squeezed temporal convolutional module (S-TCM) (Li et al. 2021a) was proposed. Transformer-based models (Xu, Tu, and Yang 2023) have emerged as a promising alternative due to their superior capability in modeling long-range dependencies using self-attention mechanisms (Yu et al. 2022). Due to its linear complexity, Mamba (Li and Chen 2024) is considered a promising alternative to transformer in sequence modeling.

Proposed Architecture

The received noisy speech in the short-time Fourier transform (STFT) domain can be presented as,

$$Y(t, f) = X(t, f) + N(t, f) \quad (1)$$

where $\{Y, X, N\}$ denote the mixture, clean and noise, respectively. $t \in \{1, \dots, T\}$ is the time frame index, and $f \in \{1, \dots, F\}$ is the frequency index.

The proposed BSDS framework is shown in Figure 1. The noisy magnitude and complex spectra are first decoupled. By separately band-splitting and feature encoding, they are converted into abstract representations for magnitude and phase, respectively. Then, stacked T-F Mamba-blocks are adopted for effectively modeling along the time and frequency axes. Subsequently, the representations from two streams are fused. The segmented frequency bands are then merged using a Mask-Decoder module to obtain the estimated complex spectrum.

Band-Split and Mask-Decoder

As shown in Figure 2, the complex and magnitude spectra are first compressed into lower-resolution bands using the Band-Split module (Yu and Luo 2023). To be specific, the noisy input spectrum X is segmented into a sequence of non-overlapping frequency bands $\{A_i\}_{i=1}^K$, each of which is individually projected to yield the embedding of dimension N . Subsequently, the K bands are stacked to obtain a 3-D tensor Z . The process can be formulated into

$$A_i \in \{A_1, A_2, \dots, A_K\}, A_i \in \mathbb{R}^{T \times F_i} \quad (2)$$

$$X_i = FC(LN(A_i)), X_i \in \mathbb{R}^{T \times N} \quad (3)$$

$$Z = \text{Concat}(X_1, X_2, \dots, X_K) \in \mathbb{R}^{K \times T \times N} \quad (4)$$

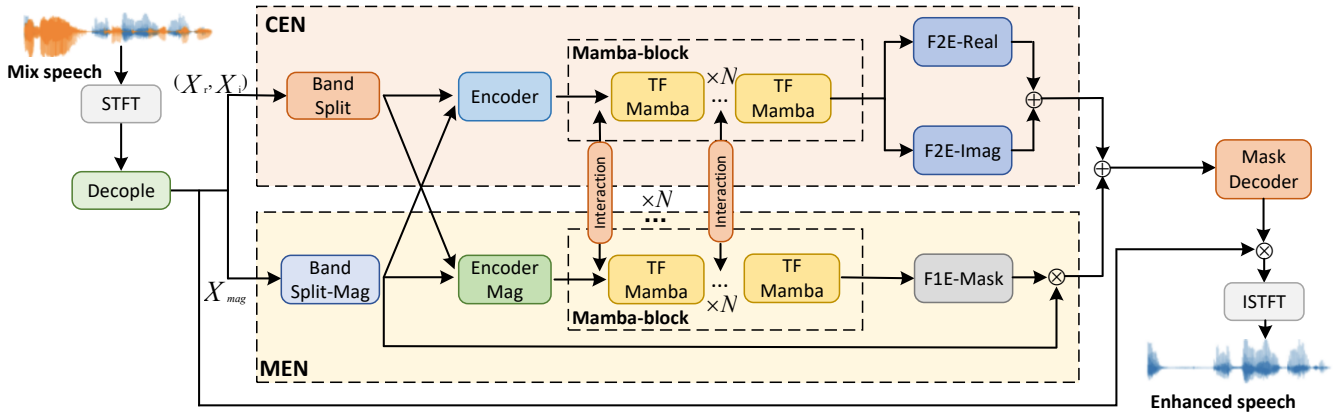


Figure 1: Overall architecture of the proposed BSDB-Net consists of three main components. The first part includes the Band-Split module for frequency band segmentation and the Mask-Decoder module for generating masks used in band synthesis. The second part features a dual-branch enhancement network: the MEN branch suppresses noise in the magnitude spectrum roughly, while the CEN branch primarily estimates complex spectra to capture phase characteristics. The third part involves the Mamba-block module designed for sequence modeling.

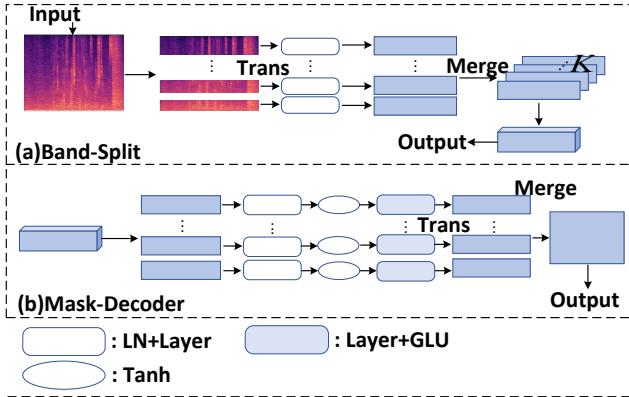


Figure 2: (a) The Band-Split module divides frequency bands for input into the modeling module. (b) The Mask-Decoder module synthesizes frequency bands post-modeling to generate masks.

where A_i represents the i frequency bands to be split, $\{FC, LN\}$ denote the linear layer and layer normalization, respectively. $Concat(\cdot)$ denotes the concatenation operation.

For the Band-Merge module, let us denote the input as $B \in \mathbb{R}^{K \times T \times N}$. Similarly, for the frequency band feature $B_i \in \mathbb{R}^{T \times N}$, where $i \in \{1, \dots, K\}$ denotes the band index, it is processed through separate layers to obtain the decoded target feature. After all the bands are processed, they are concatenated along the frequency axis to form the output M , whose formulation can be given by

$$M_i = GLU(\text{Tanh}(FC(LN(B_i)))) \quad (5)$$

$$M = \text{Merge}(M_1, M_2, \dots, M_K), M \in \mathbb{R}^{F \times T} \quad (6)$$

where B_i denotes the i th band feature, $\{\text{Tanh}, GLU\}$ denote the Tanh activation function and gated linear unit, re-

spectively. $\text{Merge}(\cdot)$ is the concatenation operation along the frequency axis.

Dual-Branch

After the band split module, the magnitude-oriented and complex-oriented features are fed into the MEN and CEN branches for modeling, respectively.

For both branches, the data first pass through the encoder layer. Instead of using typical encoder layers with downsampling operations, here we coarsely model the features without frequency downsampling to mitigate the possible information loss:

$$C_m = \text{PReLU}(\text{LN}(\text{Conv2d}(\text{Inter}(Z_{mag}, Z_{ri})))) \quad (7)$$

$$C_{ri} = \text{PReLU}(\text{LN}(\text{Conv2d}(\text{Inter}(Z_{ri}, Z_{mag})))) \quad (8)$$

where Z_{mag}, Z_{ri} represent the magnitude and complex features after Band-Split module, C_m, C_{ri} denote the output of magnitude and complex encoders, respectively. $\text{Inter}(a, b)$ denotes the interaction module between the features a and b .

As shown in Figure 1, the feature processed by the encoder is passed to the sequence modeling using the Mamba-Block. After that, a mask is generated via the F1E-Mask layer. The F1E Mask refers to the extraction of features that include amplitude and the generation of a mask. It includes a dilated block layer to expand the model's receptive field, followed by the PReLU activation function and layer normalization. Finally, the input data to the MEN and the estimated mask are multiplied to obtain the filtered feature of the MEN branch. The whole process can be given by

$$D = \text{PReLU}(\text{LN}(\text{DConv}(M_{am}(C_{mag})))) \quad (9)$$

$$D_1 = \text{Sigmoid}(\text{Conv}(D)) \quad (10)$$

$$D_2 = \text{Tanh}(\text{Conv}(D)) \quad (11)$$

$$D_{meg} = D_1 \otimes D_2, MEN_{out} = A_{mag} \otimes D_{meg} \quad (12)$$

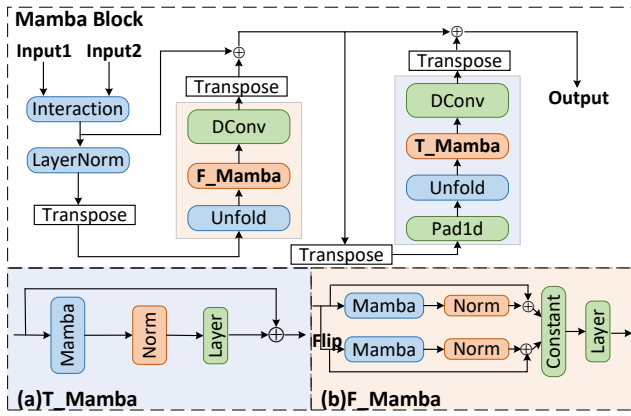


Figure 3: The Mamba-Block: It is primarily divided into temporal modeling and frequency modeling. (a) The proposed unidirectional Mamba module. (b) The proposed bidirectional Mamba module.

where \otimes denotes the elementary multiplication operation. $\{Mam(\cdot), DConv(\cdot)\}$ represent the Mamba-Block and the Dilated-Block, respectively. MEN_{out} represents the output of MEN branch. The CEN branch processes data in a similar way to the MEN branch, with the only difference lying in the final module. The F2E-Real and F2E-Imag modules refer to the implicit extraction of phase information by obtaining complex spectral features:

$$E = Mam(C_{ri}) \quad (13)$$

$$E_r = PReLU(LN(DConv(E))) \quad (14)$$

$$E_i = PReLU(LN(DConv(E))) \quad (15)$$

$$CEN_{out} = E_r \oplus E_i \quad (16)$$

where \oplus denotes the elementary sum operation. CEN_{out} represents the output of CEN branch, and E_r, E_i represent the real and imaginary parts outputted by different decoder modules, respectively.

Mamba-Block

As shown in Figure 3, the Mamba-Block module is used for sequence modeling in BSDB-Net. The corresponding equation is as follows:

$$h_n = \bar{\mathbf{A}}h_{n-1} + \bar{\mathbf{B}}x_n \quad (17)$$

$$y_n = \mathbf{C}h_n \quad (18)$$

where $\bar{\mathbf{A}}$ and $\bar{\mathbf{B}}$ represents the parameters of the discretization matrix. The process of discretization converts continuous parameters ($\Delta, \mathbf{A}, \mathbf{B}$) into discrete ones ($\bar{\mathbf{A}}, \bar{\mathbf{B}}$), enabling the model to handle discrete data effectively. Mamba incorporates Selective SSMS into an H3 structure.

First, the Mamba-Block receives data from both the MEN and CEN branches enters the interaction layer for fusion. After passing through the LN layer, the dimension is transposed from $Z \in \mathbb{R}^{B \times T \times N \times F}$ to $Z \in \mathbb{R}^{(B \times T) \times N \times F}$ to better model the frequency dimension. Subsequently, the data is sent into the F_Mamba module for sequence modeling, as

illustrated in Figure 3(b). This module achieves bidirectional modeling, the forward and backward data are sent separately into the Mamba module:

$$X_{in} = Inter(X_{ri}, X_{mag}), F_{in} = Tran(LN(X_{in})) \quad (19)$$

$$F_{out} = X_{in} \oplus Deconv(F_Ma(Unfold(F_{in}))) \quad (20)$$

where X_{ri}, X_{mag} represent the input of Mamba-Block in the Dual-branch. Inter represents interaction layer. F_Ma as shown in Figure 3(b) represents sequence modeling.

The modeling for the time dimension is similar to that for the frequency dimension. However, the key difference is that unidirectional Mamba is employed for time modeling since this model operates under a causal framework. Firstly, the data will be Transpose $Z \in \mathbb{R}^{(B \times K) \times N \times T}$. Subsequently, it will go through the T-Mamba module as shown in Figure 3(a) for sequence modeling. Once the modeling is completed, the data will pass through a deconvolution layer and be Transposed into $Z \in \mathbb{R}^{B \times T \times N \times F}$:

$$T_{in} = Unfold(Pad(Tran(F_{out}))) \quad (21)$$

$$T_{out} = Tran(Deconv(T - Ma(T_{in}))) \quad (22)$$

$$Mamba_{out} = F_{out} \oplus T_{out} \quad (23)$$

where T_Ma as shown in Figure 3(a) represents sequence modeling, $Mamba_{out}$ represents the output of Mamba-Block. Trans represents the tensor dimension reshaping.

Interaction Module

We achieve information interaction through the interaction module. This module first concatenates $Input_1$ and $Input_2$ using a Cat layer, then passes them through a Conv2d layer and a LayerNorm layer, followed by generating a mask using a Sigmoid activation function. After that, it multiplies the result with $Input_1$ to output the final data:

$$Input = Cat(Input_1, Input_2) \quad (24)$$

$$Mask = Sigmoid(LN(Conv2d(Input))) \quad (25)$$

$$Output = Input_1 + (Input_2 \otimes Mask) \quad (26)$$

Experimental Setup

Datasets

We selected the WSJ0-SI84 and the DNS-Challenge noise dataset to create synthetic data to evaluate our model and conduct ablation experiments. Subsequently, we compared our model with others using a widely-used dataset, Voice-Bank+Demand.

WSJ0-SI84+DNS-Challenge: WSJ0-SI84 (Paul and Baker 1992) consists of 7138 clean speech samples from 83 speakers. From these, 5428 and 957 utterances from 77 speakers are randomly selected for the training and validation sets. To construct "noisy-clean" training pairs, approximately 20,000 types of noise from the DNS-Challenge (Reddy et al. 2020) dataset's noise library are randomly selected and concatenated, resulting in a total duration of approximately 55 hours. The training set and validation set of the dataset are synthesized by the author

| Metrics | | Feat. | PESQ | | | | ESTOI(%) | | | | SI-SDR(dB) | | | |
|-------------------|-------------|-------------|-------------|-------------|-------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| SNR(dB) | | | -5 | 0 | 5 | Avg. | -5 | 0 | 5 | Avg. | -5 | 0 | 5 | Avg. |
| Set-A | Noisy | - | 1.54 | 1.86 | 2.17 | 1.85 | 29.25 | 43.11 | 57.53 | 43.30 | -5.00 | 0.00 | 5.00 | 0.00 |
| | ConvTasNet | Wave | 2.11 | 2.54 | 2.88 | 2.52 | 60.06 | 73.80 | 82.90 | 72.25 | 6.56 | 10.43 | 13.63 | 10.21 |
| | DPRNN | Wave | 2.17 | 2.60 | 2.96 | 2.57 | 61.74 | 74.74 | 83.53 | 73.34 | 6.88 | 10.60 | 13.82 | 10.43 |
| | DDAEC | Wave | 2.27 | 2.79 | 3.16 | 2.74 | 63.12 | 76.65 | 84.73 | 74.83 | 7.22 | 11.23 | 14.15 | 10.87 |
| | LSTM | Mag | 1.97 | 2.37 | 2.67 | 2.34 | 49.33 | 64.14 | 74.98 | 62.82 | 2.49 | 6.58 | 9.54 | 6.20 |
| | CRN | Mag | 1.97 | 2.45 | 2.79 | 2.41 | 50.52 | 66.21 | 77.24 | 64.66 | 2.66 | 7.23 | 10.79 | 6.89 |
| | GCRN | RI | 2.02 | 2.55 | 2.92 | 2.50 | 56.44 | 72.83 | 82.08 | 70.45 | 5.36 | 9.72 | 12.67 | 9.25 |
| | DCCRN | RI | 1.90 | 2.46 | 2.84 | 2.40 | 50.98 | 68.06 | 78.73 | 65.92 | 4.17 | 8.61 | 11.74 | 8.17 |
| | FullSubNet | RI | 2.20 | 2.64 | 2.97 | 2.60 | 50.44 | 67.34 | 78.88 | 65.56 | 4.34 | 9.01 | 12.81 | 8.72 |
| | CTSNet | M+RI | 2.32 | 2.79 | 3.14 | 2.75 | 62.92 | 76.20 | 84.35 | 74.49 | 6.75 | 10.84 | 13.96 | 10.52 |
| | GaGNet | M+RI | 2.36 | 2.85 | 3.22 | 2.81 | 65.84 | 78.13 | 85.79 | 76.59 | 7.36 | 11.23 | 14.31 | 10.97 |
| BSDB(ours) | M+RI | 2.43 | 2.92 | 2.26 | 2.87 | 66.27 | 78.19 | 86.66 | 77.04 | 7.42 | 11.34 | 14.39 | 11.05 | |
| Set-B | Noisy | - | 1.74 | 2.04 | 2.41 | 2.06 | 44.59 | 57.38 | 69.45 | 57.14 | -5.00 | 0.00 | 5.00 | 0.00 |
| | ConvTasNet | Wave | 2.57 | 2.90 | 3.21 | 2.89 | 73.00 | 81.79 | 87.90 | 80.90 | 10.25 | 13.18 | 16.07 | 13.17 |
| | DPRNN | Wave | 2.66 | 2.98 | 3.27 | 2.97 | 74.23 | 82.44 | 88.32 | 81.66 | 10.49 | 13.37 | 16.17 | 13.34 |
| | DDAEC | Wave | 2.83 | 3.17 | 3.43 | 3.15 | 75.57 | 83.65 | 89.03 | 82.75 | 10.91 | 13.67 | 16.16 | 13.58 |
| | LSTM | Mag | 2.37 | 2.69 | 2.94 | 2.67 | 64.27 | 74.51 | 81.76 | 73.51 | 6.30 | 9.13 | 11.08 | 8.84 |
| | CRN | Mag | 2.48 | 2.83 | 3.16 | 2.82 | 66.46 | 77.30 | 85.01 | 76.26 | 7.04 | 10.43 | 13.41 | 10.29 |
| | GCRN | RI | 2.55 | 2.94 | 3.21 | 2.90 | 70.31 | 80.82 | 86.91 | 79.35 | 9.01 | 12.27 | 14.69 | 11.99 |
| | DCCRN | RI | 2.47 | 2.89 | 3.17 | 2.84 | 66.43 | 77.92 | 84.66 | 76.34 | 8.40 | 11.74 | 13.82 | 11.32 |
| | FullSubNet | RI | 2.66 | 3.00 | 3.30 | 2.99 | 67.08 | 78.20 | 85.81 | 77.03 | 8.87 | 12.30 | 15.60 | 12.26 |
| | CTSNet | M+RI | 2.86 | 3.19 | 3.45 | 3.17 | 76.60 | 84.12 | 89.28 | 83.33 | 10.99 | 13.84 | 16.44 | 13.75 |
| | GaGNet | M+RI | 2.94 | 3.28 | 3.56 | 3.26 | 78.24 | 85.30 | 90.10 | 84.55 | 11.23 | 14.00 | 16.56 | 13.93 |
| BSDB(ours) | M+RI | 2.92 | 3.29 | 3.58 | 3.26 | 77.80 | 85.64 | 90.54 | 84.66 | 11.47 | 14.50 | 17.21 | 14.39 | |

Table 1: The objective is to compare the effects of different models on PESQ, ESTOI, and SI-SDR metrics for Set-A and Set-B in an unseen speaker test set.

| Modle | Cau. | PESQ | MACs |
|------------|------|-------------|-----------------|
| LSTM | ✓ | 2.37 | 3.69 G/s |
| CRN | ✓ | 2.45 | 2.54 G/s |
| GCRN | ✓ | 2.55 | 2.40 G/s |
| FullSubNet | ✓ | 2.64 | 29.83 G/s |
| CTSNet | ✓ | 2.79 | 5.48 G/s |
| ConvTasNet | ✓ | 2.54 | 5.22 G/s |
| DPRNN | ✓ | 2.60 | 8.47 G/s |
| DDAEC | ✓ | 2.79 | 36.85 G/s |
| DBTNet | ✗ | 3.18 | 42.64 G/s |
| GaGNet | ✓ | 2.85 | 2.81 G/s |
| 64-4(our) | ✓ | 2.62 | 0.88 G/s |
| 64-6(our) | ✓ | 2.70 | 0.98 G/s |
| 128-4(our) | ✓ | 2.78 | 1.34 G/s |
| 128-6(our) | ✓ | 2.92 | 1.68 G/s |
| 256-2(our) | ✓ | 2.74 | 1.87 G/s |
| 256-4(our) | ✓ | 2.83 | 3.06 G/s |
| 256-6(our) | ✓ | 2.95 | 4.26 G/S |

Table 2: The aim is to compare different models and the impact of varying input dimensions and module stacking times in Mamba on PESQ and Computational Complexity. (Here, “64-4” represents an input dimension of 64 with modules stacked four times, and so forth.)

Li (Li et al. 2021a), and the Set-A and Set-B are consistent with the literature (Li et al. 2022b).

VoiceBank+Demand: VoiceBank (Veaux, Yamagishi, and King 2013) consists of 30 speakers, with 28 speakers used for the training set and the remaining two speakers used for testing. The training (Botinhao et al. 2016) set includes 11,572 “noisy-clean” pairs, mixed with 10 types of noise (8 from the Demand (Thiemann, Ito, and Vincent 2013) noise database and two from artificial noise), this is a public dataset in SE, and this paper uses the same dataset as mentioned above.

Implementation Setup

All speech signals in the training set are sampled at 16 kHz. The speech signals are framed using a 20 ms Hann window with 50% overlap between frames. Transforming these framed signals into the time-frequency domain involves using a 320-point FFT. Following findings from the literature (Li et al. 2021b), we use a power spectral density compression strategy, where the compression coefficient is set to 0.5, denoted as $|X|_{0.5}$, $|S|_{0.5}$. The Adam optimizer is utilized with parameters $\beta = 0.9$ and $\beta = 0.999$. The learning rate is initialized to $5e-4$, and if the validation loss does not decrease for two consecutive evaluations, the learning rate is halved.

Baseline Models

On the WSJ0-SI84 dataset, a total of 9 baseline methods were selected for comparison with the proposed model. ConvTasNet, DPRNN (Luo, Chen, and Yoshioka 2020), and

DDAEC (Pandey and Wang 2020) are all time-domain SE models. LSTM (Chen et al. 2016), CRN (Tan and Wang 2018), GCRN (Tan and Wang 2019), DCCRN (Hu et al. 2020), FullSubNet (Hao et al. 2021), CTSNet (Li et al. 2021a), and GaGNet (Li et al. 2022b) are models in the time-frequency domain. Among them, GCRN and DCCRN were developed based on CRN. FullSubNet introduces full-band modeling and sub-band modeling. CTSNet and GaGNet are respectively parallel and serial two-stage amplitude-phase decoupling models.

On the VoiceBank+Demand dataset, a total of 13 baseline methods were selected for comparison with the proposed model. SEGAN (Pascual, Bonafonte, and Serra 2017), MMSEGAN (Soni, Shah, and Patil 2018), MetricGAN (Fu et al. 2019), and SRTNET (Qiu et al. 2023) are generative models of speech enhancement (SE). Wavenet (Oord et al. 2016) operates as a time-domain model. PHASEN (Yin et al. 2020), MHSAPK (Koizumi et al. 2020), DCRNN, TSTNN (Wang, He, and Zhu 2021), S4NDUNet (P-J et al. 2023), FDFnet (Zhang, Zou, and Zhu 2024), CSTnet, and GaGnet are all models in the time-frequency domain, with the latter five being multi-stage SE models. CompNet (Fan et al. 2023) is a model that spans both the time-domain and time-frequency domain.

Loss Function

Our BSDB-Net employs a dual-branch approach for speech enhancement, utilizing the “RI+Mag” loss function to supervise the optimization of both phase and magnitude components simultaneously:

$$\mathcal{L}_{RI} = \left\| \tilde{S}_r - S_r \right\|_F^2 + \left\| \tilde{S}_i - S_i \right\|_F^2 \quad (27)$$

$$\mathcal{L}_{Mag} = \left\| \sqrt{|\tilde{S}_r|^2 + |\tilde{S}_i|^2} - \sqrt{|S_r|^2 + |S_i|^2} \right\|_F^2 \quad (28)$$

$$\mathcal{L} = \beta \mathcal{L}_{RI} + (1 - \beta) \mathcal{L}_{Mag} \quad (29)$$

where $\|\cdot\|_F$ represents Frobenius norm, and β is empirically set to 0.5.

Evaluation Metrics

Multiple objective metrics are adopted, including narrow-band (NB) and wide-band (WB) perceptual evaluation speech quality (PESQ) (Rix et al. 2001) for speech quality, short-time objective intelligibility (STOI) and its extended version ESTOI (Jensen and Taal 2016) for intelligibility, SISDR (Hu and Loizou 2007) for speech distortion, and MOS (CSIG, CBAK, COVL) (Hu and Loizou 2007) for speech quality.

Results and Analysis

Ablation Study

We conducted ablation experiments on the WSJ0 dataset, which cover the following three aspects: (1) Whether the dual-path structure is effective; (2) How many layers of Mamba-based modules should be stacked and how many hidden layers should be fed into Mamba to achieve the best

| Model | PESQ | ESTOI | SI-SDR(dB) |
|----------------|-------------|--------------|--------------|
| BSDB-MEN | 2.67 | 71.25 | 9.84 |
| BSDB-CEN | 2.76 | 74.32 | 10.37 |
| BSDB-DB | 2.92 | 78.19 | 11.34 |

Table 3: Compare the effects of enhancing magnitude spectrum only(BSDB-MEN), enhancing complex spectrum only(BSDB-CEN), and enhancing both magnitude and complex spectra in parallel(BSDB-DB) on PESQ, ESTOI, and SDR.

effect; (3) Whether using Mamba as a sequence modeling model is effective.

Effect of Dual-Branch Model: Here we are primarily investigating the effectiveness of the dual-branch structure. Initially, we removed the CEN branch to train the MEN single branch, and subsequently removed MEN while retaining the CEN branch for training. As shown in Table 3, the dual-path structure outperforms the single-path structure across all metrics. This means that the coordinated efforts of CEN and MEN can improve the quality of the target speech. CEN filters out primary noise for rough estimation, while MEN continuously supplements speech information, thereby enhancing the overall performance of the system.

The Number of Layers and Hidden Layers of Mamba-Block: The dimensions of the input data to Mamba have an impact on the model’s performance. Additionally, stacking Mamba-Block layers can further enhance model performance but also increases complexity. As shown in Table 2, we selected Mamba input dimensions of 64, 128, and 256, and stacked modules 2, 4, and 6 times in different combinations. From our findings, when we stack sequence modeling modules and increase the number of Mamba hidden layers, performance will improve and the increase in depth has a more significant effect than the increase in breadth, using a combination like 128-6 for Mamba already achieves a balance between performance and complexity. Additionally, the current SE models are all at the complexity level of Gb/s, while our model further compresses the complexity to the level of Mb/s without a significant decrease in performance.

Effect of Mamba-Block: Our model aims to ensure performance while compressing model complexity, benefiting from advancements in Selective State Spaces. To demonstrate the effectiveness of Mamba-Block, we replaced it with LSTM and Transformer while keeping other factors constant, as shown in Table 5 Our results indicate that our model maintains optimal performance with significantly reduced complexity, validating the applicability and feasibility of Mamba for our task. Although the performance improvement over the Transformer is relatively minor, the significant reduction in complexity is the reason why our model opts for Mamba for sequence modeling.

Model Complexity Comparison

As shown in Table 2, we evaluated the complexity of our proposed model and other baseline models on the WSJ0-SI84 dataset. It is worth noting that all input samples were

| Modle | Year | PESQ-WB | STOI% | CSIG | CBAK | COVL |
|---------------------|------|-------------|-----------|-------------|-------------|-------------|
| Noisy | - | 1.97 | 92.1 | 3.35 | 2.44 | 2.63 |
| SEGAN | 2017 | 2.16 | 92.5 | 3.48 | 2.94 | 2.80 |
| MMSEGAN | 2018 | 2.53 | 93 | 3.8 | 3.12 | 3.14 |
| Wavenet | 2018 | - | - | 3.62 | 3.32 | 2.98 |
| MetricGAN | 2019 | 2.86 | - | 3.99 | 3.18 | 3.42 |
| DCCRN | 2020 | 2.68 | 93.7 | 3.88 | 3.18 | 3.27 |
| PHASEN | 2020 | 2.99 | - | 4.21 | 3.55 | 3.62 |
| MHSA-SPK | 2020 | 2.99 | - | 4.15 | 3.42 | 3.53 |
| TSTNN | 2021 | 2.96 | 95 | 4.17 | 3.53 | 3.49 |
| CTS-Net | 2022 | 2.92 | - | 4.25 | 3.46 | 3.59 |
| GaGnet | 2022 | 2.94 | 94.7 | 4.26 | 3.45 | 3.59 |
| SRTNET | 2023 | 2.69 | - | 4.12 | 3.19 | 3.39 |
| CompNet | 2023 | 2.90 | - | 4.16 | 3.37 | 3.53 |
| FDNet | 2024 | 3.05 | - | 4.23 | 3.55 | 3.65 |
| S4DSE | 2024 | 2.55 | - | 3.94 | 3.00 | 3.32 |
| BSDBNet(256) | 2024 | 3.11 | 95 | 4.33 | 3.58 | 3.73 |
| BSDBNet(128) | 2024 | 3.07 | 94.8 | 4.32 | 3.58 | 3.71 |

Table 4: Comparison was conducted with other state-of-the-art methods, including both time-domain and time-frequency domain approaches. “-” indicates where results were not provided in the original text.

set to one second of audio, ensuring fairness in our experiments. The model selected from the previous ablation experiments has an average computational complexity about 8 times lower than the baseline and has good performance. BSDB-Net exhibits lower complexity along with excellent performance. In terms of performance, our model is only slightly behind DBTNet. The reasons for this can be analyzed as follows: First, our model is a causal model while DBTNet is a non-causal model. Second, DBTNet has been continuously increasing its model complexity to achieve SOTA performance. As can be seen from Table 2, its model complexity is about 20 times that of our model, yet the performance has not been significantly improved. When we further compress the complexity to the level of M/s, proving the effectiveness of our network structure.

Comparisons with Baselines on WSJ0-SI84 Corpus

As shown in Table 1, the objective metrics results for the proposed method and baseline models on the WSJ0-SI84 dataset include PESQ, ESTOI, and SI-SDR. From the table, we can draw the following conclusions: Firstly, it is evident that models based on complex spectra generally outperform those based solely on magnitude spectra. For instance, models like GCRN and DCCRN consistently outperform CRN and LSTM across all metrics. This indicates that incorporating phase information in addition to amplitude recovery can significantly enhance both speech quality and intelligibility. Secondly, multi-stage models that simultaneously consider magnitude and complex spectra outperform models that focus solely on a single spectrum type or operate in the time domain. Models like CSTNet and GaGNet show superior performance compared to others, suggesting that parallel optimization of amplitude and complex spectra can effectively leverage phase information to generate

| Model | PESQ | MACs | Parameters |
|-------------------------|-------------|----------------|--------------|
| Ours-LSTM | 2.82 | 11.05G/s | 14.71M |
| Ours-Transformer | 2.90 | 28.06G/s | 14.06M |
| Ours-Mamba(ours) | 2.92 | 1.68G/s | 9.78M |

Table 5: Compare the impact of replacing the sequence modeling module with LSTM, Transformer, and Mamba on PESQ, Computational Complexity, and Parameters.

higher-quality speech. Lastly, our model achieves state-of-the-art (SOTA) results across all metrics. The reasons for this can be analyzed as follows: BSDB-NET employs a parallel two-stage model, which is better at decoupling amplitude and phase information to address compensation effects compared to other models; A sequence modeling module based on Mamba is constructed, demonstrating good performance at a linear complexity.

Comparisons with Baselines on VoiceBank + Demand

In addition to the WSJ0-SI84 corpus, we also conducted experiments on another public benchmark, VoiceBank + Demand. BSDB-Net was compared with other baselines, and our model achieved superior results across all metrics. As evident from Table 4, the baseline models achieved average improvements of 0.32, 1.5%, 0.31, 0.17, and 0.35 in PESQ, STOI, CSIG, CBAK, and COVL, respectively. This indicates that the methodologies proposed by our model adeptly address the issues presented in the background. The main purpose of experimenting on this public dataset is to fairly demonstrate that our model can maintain competitive performance while significantly reducing complexity, thereby better proving that the network has indeed effectively solved the proposed problem.

Conclusion

In this paper, we propose a Band-Split Dual-branch Network based on Selective State Spaces. This network reduces complexity while estimating speech spectra through a complementary mechanism. Specifically, we divide the network into a magnitude enhancement net (MEN) and a complex spectral enhancement net (CEN), which jointly filter noise components while continuously refining and supplementing speech spectrum information. To further reduce model complexity, we introduce a band-splitting strategy in each branch. Additionally, we incorporate Mamba for sequence modeling, ensuring model performance while reducing complexity. In future work, we intend to further enhance model performance and decrease the computational complexity. Besides, we intend to apply the proposed method to more tasks, like multi-channel speech enhancement, dereverberation, and target extraction. Mamba is still far from mature in deployment, which renders it necessary for further optimization.

Acknowledgments

This work is supported by the STI 2030—Major Projects (No. 2021ZD0201500), the National Natural Science Foundation of China (NSFC) (No.62201002, 6247077204), Excellent Youth Foundation of Anhui Scientific Committee (No. 2408085Y034), Distinguished Youth Foundation of Anhui Scientific Committee (No. 2208085J05), Special Fund for Key Program of Science and Technology of Anhui Province (No. 202203a07020008), Cloud Ginger XR-1.

References

- Alex, T.; Ahmed, S.; Mustafa, A.; Awais, M.; and Jackson, P. J. 2024. DTF-AT: Decoupled Time-Frequency Audio Transformer for Event Classification. In *AAAI 2024*, volume 38, 17647–17655.
- Bai, S.; Kolter, J. Z.; and Koltun, V. 2018. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*.
- Botinhao, C. V.; Wang, X.; Takaki, S.; and Yamagishi, J. 2016. Investigating RNN-based speech enhancement methods for noise-robust Text-to-Speech. In *9th ISCA Speech Synthesis Workshop*, 159–165.
- Chan, J. N. L. Q., W. 2016. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *speech and signal processing*, 4960–4964. Detroit, Michigan: IEEE international conference on acoustics.
- Chen, J.; Wang, Y.; Yoho, S. E.; Wang, D.; and Healy, E. W. 2016. Large-scale training to increase speech intelligibility for hearing-impaired listeners in novel noises. *The Journal of the Acoustical Society of America*, 139(5): 2604–2612.
- Divenyi, P. 2004. *Speech separation by humans and machines*. Springer Science & Business Media.
- Fan, C.; Zhang, H.; Li, A.; Xiang, W.; Zheng, C.; Lv, Z.; and Wu, X. 2023. CompNet: Complementary network for single-channel speech enhancement. *Neural Networks*, 168: 508–517.
- Fu, S.-W.; Liao, C.-F.; Tsao, Y.; and Lin, S.-D. 2019. Metricgan: Generative adversarial networks based black-box metric scores optimization for speech enhancement. In *International Conference on Machine Learning*, 2031–2041. PmLR.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Gu, A.; Dao, T.; Ermon, S.; Rudra, A.; and Ré, C. 2020. Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems*, 33: 1474–1487.
- Gu, A.; Goel, K.; and Ré, C. 2021. Efficiently modeling long sequences with structured state spaces. *International Conference on Learning Representations*.
- Hamilton, J. D. 1994. State-space models. *Handbook of econometrics*, 4: 3039–3080.
- Hao, X.; Su, X.; Horaud, R.; and Li, X. 2021. Fullsubnet: A full-band and sub-band fusion model for real-time single-channel speech enhancement. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing*, 6633–6637. IEEE.
- Hu, Y.; Liu, Y.; Lv, S.; Xing, M.; Zhang, S.; Fu, Y.; Wu, J.; Zhang, B.; and Xie, L. 2020. DCCRN: Deep complex convolution recurrent network for phase-aware speech enhancement. In *INTER_SPEECH*, volume 2020. International Speech Communication Association.
- Hu, Y.; and Loizou, P. C. 2007. Evaluation of objective quality measures for speech enhancement. *IEEE Transactions on audio, speech, and language processing*, 16(1): 229–238.
- Hummerson, C.; Stokes, T.; and Brookes, T. 2014. On the ideal ratio mask as the goal of computational auditory scene analysis. In *Blind source separation: advances in theory, algorithms and applications*, 349–368. Springer.
- Jensen, J.; and Taal, C. H. 2016. An algorithm for predicting the intelligibility of speech masked by modulated noise maskers. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(11): 2009–2022.
- Koizumi, Y.; Yatabe, K.; Delcroix, M.; Masuyama, Y.; and Takeuchi, D. 2020. Speech enhancement using self-adaptation and multi-head self-attention. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing*, 181–185. IEEE.
- Li, A.; Liu, W.; Zheng, C.; Fan, C.; and Li, X. 2021a. Two heads are better than one: A two-stage complex spectral mapping approach for monaural speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29: 1829–1843.
- Li, A.; You, S.; Yu, G.; Zheng, C.; and Li, X. 2022a. Taylor, Can You Hear Me Now? A Taylor-Unfolding Framework for Monaural Speech Enhancement. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 4193–4200. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Li, A.; Zheng, C.; Peng, R.; and Li, X. 2021b. On the importance of power compression and phase estimation in monaural speech dereverberation. *JASA express letters*, 1(1).
- Li, A.; Zheng, C.; Zhang, L.; and Li, X. 2022b. Glance and gaze: A collaborative learning framework for single-channel speech enhancement. *Applied Acoustics*, 187: 108499.
- Li, K.; and Chen, G. 2024. Spmamba: State-space model is all you need in speech separation. *arXiv preprint arXiv:2404.02063*.
- Liu, G.; Shi, J.; Chen, X.; Xu, J.; and Xu, B. 2018. Improving Speech Separation with Adversarial Network and Reinforcement Learning. In *2018 IJCNN*, 1–7.
- Lu, Y.-X.; Ai, Y.; and Ling, Z.-H. 2023. MP-SENet: A Speech Enhancement Model with Parallel Denoising of Magnitude and Phase Spectra. In *Proc. INTER_SPEECH 2023*, 3834–3838.
- Luo, Y.; Chen, Z.; and Yoshioka, T. 2020. Dual-path rnn: efficient long sequence modeling for time-domain single-channel speech separation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing*, 46–50. IEEE.

- Luo, Y.; and Mesgarani, N. 2019. Conv-tasnet: Surpassing ideal time–frequency magnitude masking for speech separation. *IEEE/ACM transactions on audio, speech, and language processing*, 27(8): 1256–1266.
- Oord, A. v. d.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; and Kavukcuoglu, K. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
- P-J, K.; Yang, C.; Siniscalchi, S.; et al. 2023. A Multi-dimensional Deep Structured State Space Approach to Speech Enhancement Using Small-footprint Models. In *INTERSPEECH*, volume 2023, 2453–2457. International Speech Communication Association.
- Pandey, A.; and Wang, D. 2019. TCNN: Temporal convolutional neural network for real-time speech enhancement in the time domain. In *ICASSP 2019*, 6875–6879. IEEE.
- Pandey, A.; and Wang, D. 2020. Densely connected neural network with dilated convolutions for real-time speech enhancement in the time domain. In *ICASSP 2020*, 6629–6633. IEEE.
- Pascual, S.; Bonafonte, A.; and Serra, J. 2017. SEGAN: Speech enhancement generative adversarial network. *arXiv preprint arXiv:1703.09452*.
- Paul, D. B.; and Baker, J. 1992. The design for the Wall Street Journal-based CSR corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Harriman, New York, February 23-26, 1992*.
- Qiu, Z.; Fu, M.; Yu, Y.; Yin, L.; Sun, F.; and Huang, H. 2023. SRTNET: Time Domain Speech Enhancement via Stochastic Refinement. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Reddy, C. K.; Gopal, V.; Cutler, R.; Beyrami, E.; Cheng, R.; Dubey, H.; Matushevych, S.; Aichner, R.; Aazami, A.; Braun, S.; et al. 2020. The interspeech 2020 deep noise suppression challenge: Datasets, subjective testing framework, and challenge results. *Interspeech 2020*.
- Rix, A. W.; Beerends, J. G.; Hollier, M. P.; and Hekstra, A. P. 2001. Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs. In *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, volume 2, 749–752. IEEE.
- Soni, M. H.; Shah, N.; and Patil, H. A. 2018. Time-frequency masking-based speech enhancement using generative adversarial network. In *ICASSP 2018 - 2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 5039–5043. IEEE.
- Tan, K.; and Wang, D. 2018. A convolutional recurrent neural network for real-time speech enhancement. In *Interspeech*, volume 2018, 3229–3233.
- Tan, K.; and Wang, D. 2019. Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28: 380–390.
- Thiemann, J.; Ito, N.; and Vincent, E. 2013. The diverse environments multi-channel acoustic noise database (demand): A database of multichannel environmental noise recordings. In *Proceedings of Meetings on Acoustics*, volume 19. AIP Publishing.
- Veaux, C.; Yamagishi, J.; and King, S. 2013. The voice bank corpus: Design, collection and data analysis of a large regional accent speech database. In *2013 international conference oriental COCOSDA held jointly with 2013 conference on Asian spoken language research and evaluation (O-COCOSDA/CASLRE)*, 1–4. IEEE.
- Wang, K.; He, B.; and Zhu, W.-P. 2021. TSTNN: Two-stage transformer based neural network for speech enhancement in the time domain. In *ICASSP 2021 IEEE International Conference on Acoustics, Speech and Signal Processing*, 7098–7102. IEEE.
- Wang, Z.-Q.; Wichern, G.; and Le Roux, J. 2021. On the compensation between magnitude and phase in speech separation. *IEEE Signal Processing Letters*, 28: 2018–2022.
- Westhausen, N. L.; and Meyer, B. T. 2020. Dual-Signal Transformation LSTM Network for Real-Time Noise Suppression. *Interspeech 2020*.
- Williamson, D. S.; Wang, Y.; and Wang, D. 2015. Complex ratio masking for monaural speech separation. *IEEE/ACM transactions on audio, speech, and language processing*, 24(3): 483–492.
- Xu, X.; Tu, W.; and Yang, Y. 2023. Selector-enhancer: learning dynamic selection of local and non-local attention operation for speech enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 13853–13860.
- Yin, D.; Luo, C.; Xiong, Z.; and Zeng, W. 2020. Phasen: A phase-and-harmonics-aware speech enhancement network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 9458–9465.
- Yu, G.; Li, A.; Wang, H.; Wang, Y.; Ke, Y.; and Zheng, C. 2022. DBT-Net: Dual-branch federative magnitude and phase estimation with attention-in-attention transformer for monaural speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30: 2629–2644.
- Yu, J.; and Luo, Y. 2023. Efficient Monaural Speech Enhancement with Universal Sample Rate Band-Split RNN. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Yue, H.; Duo, W.; Peng, X.; and Yang, J. 2022. Reference-based speech enhancement via feature alignment and fusion network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 11648–11656.
- Zaremba, W.; Sutskever, I.; and Vinyals, O. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*.
- Zhang, Y.; Zou, H.; and Zhu, J. 2024. A Two-Stage Framework in Cross-Spectrum Domain for Real-Time Speech Enhancement. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 12587–12591. IEEE.