

EXPLORING RESOLUTION-WISE SHARED ATTENTION IN HYBRID MAMBA-U-NETS FOR IMPROVED CROSS-CORPUS SPEECH ENHANCEMENT

Nikolai Lund Kühne*

Jesper Jensen^{*†}

Jan Østergaard*

Zheng-Hua Tan*

^{*}Department of Electronic Systems, Aalborg University, 9220, Denmark.

[†]Oticon A/S, 2765, Denmark

ABSTRACT

Recent advances in speech enhancement have shown that models combining Mamba and attention mechanisms yield superior cross-corpus generalization performance. At the same time, integrating Mamba in a U-Net structure has yielded state-of-the-art enhancement performance, while reducing both model size and computational complexity. Inspired by these insights, we propose RWSA-MambaUNet, a novel and efficient hybrid model combining Mamba and multi-head attention in a U-Net structure for improved cross-corpus performance. Resolution-wise shared attention (RWSA) refers to layerwise attention-sharing across corresponding time- and frequency resolutions. Our best-performing RWSA-MambaUNet model achieves state-of-the-art generalization performance on two out-of-domain test sets. Notably, our smallest model surpasses all baselines on the out-of-domain DNS 2020 test set in terms of PESQ, SSNR, and ESTOI, and on the out-of-domain EARS-WHAM_v2 test set in terms of SSNR, ESTOI, and SI-SDR, while using less than half the model parameters and a fraction of the FLOPs.

Index Terms— Speech Enhancement, Mamba, Attention, U-Net, Hybrid Model

1. INTRODUCTION

Speech enhancement aims at removing background noise from speech signals, thereby improving speech intelligibility and quality. It has a wide range of applications such as hearing assistive devices, mobile communication systems, and speaker verification.

In the past decade, research in deep-learning based speech enhancement has included a wide range of neural architectures [1]. The applied architectures include convolutional neural networks [2, 3], diffusion models [4, 5] and generative adversarial networks (GANs) [6, 7]. Recently, attention based neural architectures such as Transformers and Conformers have been the most prevalent, as these models have demonstrated state-of-the-art (SOTA) performance on multiple benchmarks [8, 9]. However, multi-head attention (MHA) based models scale quadratically with input size in terms of computational complexity [10]. This led to a newfound interest in recurrent models that instead scale linearly with respect to input size. Two of such models, Mamba [10] and Extended Long Short-Term Memory (xLSTM) [11], have already demonstrated SOTA in-domain speech enhancement performance [12–14]. Moreover, recent works such as Mamba-SEUNet [13] and MUSE [15] have demonstrated the effectiveness of U-Nets in speech enhancement, offering similar or improved enhancement performance with fewer model parameters and lower computational complexity. However, the efficacy of U-Nets for cross-corpus speech enhancement has not been explored yet.

Cross-corpus generalization performance is important for speech enhancement systems, as such systems may be expected to operate

across a diverse range of acoustic conditions, and it is infeasible to include all recording conditions, noise types, and speakers in the training dataset. Unfortunately, sequence models like LSTM, xLSTM, and Mamba have demonstrated worse generalization performance compared to purely attention based models [16–18]. As an alternative to purely attention based models, the hybrid MambAttention model [18] was recently proposed. By combining Mamba with a shared time- and frequency-MHA module, MambAttention achieved SOTA cross-corpus generalization performance on two out-of-domain test sets across all evaluation metrics used.

Based on the SOTA cross-corpus generalization of MambAttention, we hypothesize that explicitly aligning global time- and frequency relations is critical for robust cross-corpus speech enhancement. To realize this hypothesis, we propose resolution-wise shared attention (RWSA), which is shared layerwise time- and frequency-MHA modules across corresponding time- and frequency resolutions in Mamba-UNets. Our proposed RWSA-MambaUNet employs MambAttention blocks, which have demonstrated superior generalization performance [18]. By introducing RWSA, our best-performing RWSA-MambaUNet model achieves SOTA generalization performance on two out-of-domain test sets with different speakers, noise types, and recording conditions across PESQ, SSNR, ESTOI, and SI-SDR, at a significantly lower computational complexity than the baselines. Remarkably, even our smallest model outperforms all baselines on most metrics for cross-corpus generalization with less than half the model parameters. Code is publicly available.¹ Our major contributions are summarized as follows:

- We propose RWSA-MambaUNet, a novel and efficient hybrid model using resolution-wise shared attention in a U-Net architecture for improved cross-corpus generalization.
- We demonstrate that RWSA is essential for the SOTA cross-corpus performance of our RWSA-MambaUNet models.
- Our best-performing model surpasses existing SOTA baselines on two out-of-domain test sets across all evaluation metrics used, while requiring significantly fewer FLOPs.

2. METHOD

2.1. MambAttention

Our RWSA-MambaUNet consists of multiple MambAttention blocks [18] at different time- and frequency resolutions. We employ these blocks, as they have demonstrated SOTA speech enhancement generalization performance [18]. MambAttention blocks comprise bidirectional Mamba blocks across time (T-Mamba) and frequency

¹<https://github.com/NikolaiKyhne/RWSAMamba-UNet>




Fig. 1: Overall structure of our proposed RWSA-MambaUNet. Resolution-wise shared attention (purple dashed lines) is layerwise sharing of MHA modules within MambaAttention blocks across corresponding resolutions between the upsampling and downsampling path. To simplify the figure, we have not depicted the residual connections between the output of the feature encoder, and the outputs of both refinement layers.

(F-Mamba), as well as shared time-MHA (T-MHA) and frequency-MHA (F-MHA) modules. Given an input $\mathbf{X} \in \mathbb{R}^{B \times C \times T \times F}$, where B is the batch size, C is the number of channels, and T and F denote the number of time frames and frequency bins, respectively, the forward pass of a MambaAttention block is given by [18]:

$$\mathbf{X}_{\text{Time}} = \text{reshape}(\mathbf{X}, [B \cdot F, T, C]), \quad (1)$$

$$\mathbf{X}_1 = \mathbf{X}_{\text{Time}} + \text{T-MHA}(\text{LN}(\mathbf{X}_{\text{Time}})), \quad (2)$$

$$\mathbf{X}_2 = \mathbf{X}_1 + \text{T-Mamba}(\mathbf{X}_1), \quad (3)$$

$$\mathbf{X}_{\text{Freq.}} = \text{reshape}(\mathbf{X}_2, [B \cdot T, F, C]), \quad (4)$$

$$\mathbf{X}_3 = \mathbf{X}_{\text{Freq.}} + \text{F-MHA}(\text{LN}(\mathbf{X}_{\text{Freq.}})), \quad (5)$$

$$\mathbf{X}_4 = \mathbf{X}_3 + \text{F-Mamba}(\mathbf{X}_3), \quad (6)$$

$$\mathbf{Y} = \text{reshape}(\mathbf{X}_4, [B, C, T, F]), \quad (7)$$

where $\text{reshape}(\text{input}, \text{size})$ reshapes the input to a desired size, and LN is Layer Normalization. MambaAttention utilizes the T- and F-Mamba blocks from SEMamba [12], and the output $\mathbf{X}_{\text{out}}$ of each T- and F-Mamba block is given by:

$$\mathbf{X}_{\text{out}} = \text{Conv1D}(\text{Concat}(\text{Mamba}(\mathbf{X}_{\text{in}}), \text{flip}(\text{Mamba}(\text{flip}(\mathbf{X}_{\text{in}}))))), \quad (8)$$

where $\mathbf{X}_{\text{in}}$ is the input to the T- and F-Mamba blocks, and $\text{Mamba}(\cdot)$, $\text{flip}(\cdot)$, $\text{Concat}(\cdot)$, and $\text{Conv1D}(\cdot)$ is the unidirectional Mamba, sequence flipping, concatenation, and 1D transposed convolution.

2.2. Model overview

Our RWSA-MambaUNet is portrayed in Figure 1.

Preprocessing and feature encoder: Before the noisy speech waveform $\mathbf{y} \in \mathbb{R}^L$ is processed by the feature encoder, a complex spectrogram is computed through a short-time Fourier transform (STFT). The input $\mathbf{Y}_{\text{in}} \in \mathbb{R}^{T \times F \times 2}$ to the feature encoder then becomes the compressed magnitude spectrum $(\mathbf{Y}_m)^c \in \mathbb{R}^{T \times F}$, extracted via power-law compression [19], concatenated with the wrapped phase spectrum $\mathbf{Y}_p \in \mathbb{R}^{T \times F}$. The feature encoder is identical to the one used in MP-SENet [8] and thus increases the number of input channels from 2 to C and halves the frequency dimension from F to $F' = F/2$. It consists of two convolution blocks, each comprising a 2D convolutional layer, an instance normalization, and a PReLU activation, sandwiching a dilated DenseNet [20].

U-Net architecture: The output of the feature encoder is processed by multiple MambaAttention blocks with skip connections, patch embedding layers, as well as convolutional downsampling and upsampling blocks, following a U-Net-style architecture. The patch embeddings, originally proposed in [21], consist of depthwise separable and deformable convolutions, which facilitates learning intricate fine-grained acoustic details. The MambaAttention blocks focus on capturing time and frequency dependencies across acoustic features at different time- and frequency resolutions. The final magnitude and phase refinement layers comprise a patch embedding, N stacks of TF-Mamba blocks from [12], and a 3×3 convolution.

Resolution-wise shared attention: We believe attention-sharing across corresponding resolutions aids reconstruction in the U-Net, since global time- and frequency relationships are aligned at matching resolution levels across both the down- and upsampling paths. RWSA (purple dashed lines in Figure 1) leverages the fact that distinct MambAttention blocks are used at corresponding time- and frequency resolutions in both paths of the U-Net. By sharing the T- and F-MHA modules not only within each individual MambAttention layer but also across layers in both the downsampling and upsampling paths, the model jointly aligns global temporal- and spectral dependencies across multiple resolution scales. As we will demonstrate, RWSA improves generalization performance, while minimizing the memory cost of the attention blocks.

Magnitude mask and wrapped phase decoder: Finally, after the magnitude and phase refinement layers, the clean magnitude and phase spectra are estimated through the magnitude mask and wrapped phase decoder, respectively. Following [9], both the magnitude mask and wrapped phase decoder consist of a dilated DenseNet, followed by a sub-pixel convolution [22], an instance normalization, and a PReLU activation. In the magnitude mask decoder, this is followed by a deconvolution block reducing the output channels from C to 1. A learnable sigmoid function with $\beta = 2$ is used to estimate the magnitude mask, as in [8]. In the wrapped phase decoder, the sub-pixel convolution is followed by two parallel 2D convolutional layers yielding the pseudo-real and pseudo-imaginary part components. The clean wrapped phase spectrum is estimated using the two-argument arctangent function [8]. The final enhanced waveform is recovered by applying an inverse STFT to the estimated clean magnitude spectrum and estimated wrapped phase spectrum.

We follow MambAttention [18] and use a linear combination of loss functions, including a PESQ-based GAN discriminator, along with time, magnitude, complex, phase, and consistency losses.

3. EXPERIMENTS

3.1. Datasets

We train and evaluate our models on the VB-DemandEx dataset [18]. The dataset contains 10,840 noisy-clean pairs of audio clips for training, 730 for validation, and 840 for testing. The clean speech originates from the VoiceBank corpus [23], where 26 distinct speakers are used for training, 2 distinct speakers are used for validation, and 2 distinct speakers are used for testing. The noisy audio clips are created by mixing clean samples with noise from the DEMAND database [24] as well as babble and speech-shaped noise at 7 segmental SNRs (SSNRs) ($-10, -5, 0, 5, 10, 15, 20$ dB).

In addition, we train and evaluate our models on the large-scale Deep Noise Suppression Challenge 2020 dataset (DNS 2020) [25]. DNS 2020 contains 500 hours of clean speech from 2,150 speakers and more than 180 hours of noise clips. Since DNS 2020 has no validation set, we use the validation set generated in [18]. Using the official script provided in [25], we generate 3,000 hours of noisy-clean pairs of audio clips for training with SSNRs uniformly sampled between -5 dB and 15 dB. This yields 1.08 M 10-second audio clips. We use the DNS 2020 test set without reverberation for evaluating our models. The test set contains 150 noisy-clean pairs, generated from audio clips spoken by 20 distinct speakers.

Finally, we also evaluate cross-corpus performance on the 16 kHz version of the EARS-WHAM.v2 test set [26, 27]. EARS-WHAM.v2 contains clean speech, recorded in an anechoic chamber, from 107 distinct speakers. The clean speech covers reading tasks in 7 reading styles, emotional reading, conversational speech, and

freeform speech. Using the script provided in [26], we mix the clean speech from speakers $p102$ to $p107$ with noise recordings from the WHAM! dataset [27] at SNRs randomly sampled in the interval $[-2.5, 17.5]$ dB. This results in 886 noisy-clean pairs for testing.

3.2. Implementation details

Unless otherwise stated, all experimental details and training configurations match those presented in MambAttention [18]. We train on 30,600 point audio segments, and use an FFT order of 510, a Hann window size of 510, and a hop size of 120 for all STFTs. Moreover, we use a magnitude spectrum compression factor of $c = 0.3$. In the MambAttention blocks, we use $h = 8$ attention heads for the bottleneck layers and $h = 4$ attention heads anywhere else. Checkpoints are saved every 250 steps, and for evaluation we select the checkpoint that obtains the highest PESQ score on the validation set. Models trained on VB-DemandEx and DNS 2020 are trained for 550 k and 950 k steps respectively, with a batch size $B = 8$ on four NVIDIA L40S GPUs. Table 3 provides important hyperparameters for our proposed RWSA-MambaUNet models.

Table 3: Model hyperparameters for the proposed RWSA-MambaUNet models.

Model	# Channels C	# Blocks N	Params
RWSA-MambaUNet-XS	16	2	1.02M
RWSA-MambaUNet-S	16	4	1.95M
RWSA-MambaUNet-M	24	4	3.91M

3.3. Evaluation metrics

We apply wide-band PESQ [28] to evaluate the speech quality of the enhanced speech. Moreover, we report the waveform-matching-based evaluation metrics SSNR [29] and scale-invariant signal-to-distortion ratio (SI-SDR) [30]. The intelligibility of the enhanced speech is predicted using extended short-time objective intelligibility (ESTOI) [31]. Across these measures, higher values indicate better performance. Finally, we report FLOPs, which are calculated based on processing a single audio sample on one GPU. We train all models with 5 different seeds, and report the mean and standard deviation.

4. RESULTS

4.1. Generalization performance

We evaluate in-domain performance on the VB-DemandEx dataset. For assessing cross-corpus generalization performance, we evaluate on two out-of-domain test sets with different noise, speaker, and recording conditions from DNS 2020 [25] and EARS-WHAM.v2 [26, 27]. For simplicity, we rename the LSTM baseline from [14] to LSTM-SENet.

In Table 1, we report in- and out-of-domain speech enhancement performance. From Table 1, it is clear that our RWSA-MambaUNet-XS outperforms all the LSTM-SENet, xLSTM-SENet, SEMamba, MP-SENet, and the MambAttention baselines on the out-of-domain DNS 2020 test set across PESQ, SSNR, and ESTOI, and on the EARS-WHAM.v2 test set across SSNR, ESTOI, and SI-SDR, with only 1.02 M parameters and 9.22 G FLOPs. By doubling the number of layers from 2 to 4, our RWSA-MambaUNet-S further improves both in- and out-of-domain enhancement performance across

Table 1: In-domain and out-of-domain speech enhancement performance. Models are trained on VB-DemandEx. All baselines are trained using their originally provided code. Best reported mean is marked bold.

Dataset			In-Domain				Out-Of-Domain								
			VB-DemandEx				DNS 2020				EARS-WHAM_v2				
			PESQ	SSNR	ESTOI	SI-SDR	PESQ	SSNR	ESTOI	SI-SDR	PESQ	SSNR	ESTOI	SI-SDR	
Model	Params	FLOPs													
Noisy	-	-	1.625	-1.068	0.630	4.976	1.582	6.218	0.810	9.071	1.235	-0.803	0.640	5.359	
xLSTM-SENet [14]	2.20M	80.71G	2.973 $\pm$ 0.051	7.933 $\pm$ 0.133	0.795 $\pm$ 0.008	16.414 $\pm$ 0.317	1.724 $\pm$ 0.368	3.246 $\pm$ 1.332	0.686 $\pm$ 0.097	3.412 $\pm$ 3.482	1.505 $\pm$ 0.151	0.446 $\pm$ 0.566	0.559 $\pm$ 0.053	1.396 $\pm$ 2.141	
LSTM-SENet [14]	2.34M	88.59G	3.002 $\pm$ 0.026	7.981 $\pm$ 0.210	0.802 $\pm$ 0.003	16.637 $\pm$ 0.123	1.984 $\pm$ 0.454	4.901 $\pm$ 1.656	0.724 $\pm$ 0.117	4.749 $\pm$ 3.346	1.570 $\pm$ 0.179	0.854 $\pm$ 0.773	0.566 $\pm$ 0.083	1.916 $\pm$ 2.894	
SEMamba [12]	2.25M	65.46G	3.002 $\pm$ 0.022	7.590 $\pm$ 0.177	0.800 $\pm$ 0.003	16.593 $\pm$ 0.159	2.281 $\pm$ 0.134	5.837 $\pm$ 1.033	0.820 $\pm$ 0.028	9.298 $\pm$ 1.576	1.631 $\pm$ 0.053	0.921 $\pm$ 0.508	0.603 $\pm$ 0.026	2.809 $\pm$ 0.523	
MP-SENet [8]	2.05M	74.29G	2.935 $\pm$ 0.065	7.641 $\pm$ 0.283	0.787 $\pm$ 0.010	16.202 $\pm$ 0.318	2.666 $\pm$ 0.010	7.369 $\pm$ 0.382	0.875 $\pm$ 0.009	13.665 $\pm$ 0.892	1.862 $\pm$ 0.097	2.107 $\pm$ 0.270	0.677 $\pm$ 0.029	6.090 $\pm$ 0.672	
MambaAttention [18]	2.33M	65.52G	3.026 $\pm$ 0.007	7.674 $\pm$ 0.411	0.801 $\pm$ 0.002	16.684 $\pm$ 0.095	2.919 $\pm$ 0.118	8.133 $\pm$ 0.733	0.911 $\pm$ 0.009	15.169 $\pm$ 1.363	2.010 $\pm$ 0.053	2.505 $\pm$ 0.224	0.725 $\pm$ 0.020	7.348 $\pm$ 0.445	
RWSA-MambaUNet-XS	1.02M	9.22G	2.893 $\pm$ 0.009	7.041 $\pm$ 0.073	0.780 $\pm$ 0.002	15.212 $\pm$ 0.064	2.940 $\pm$ 0.019	9.421 $\pm$ 0.132	0.922 $\pm$ 0.002	14.722 $\pm$ 0.120	1.987 $\pm$ 0.023	3.106 $\pm$ 0.188	0.729 $\pm$ 0.006	8.541 $\pm$ 0.347	
RWSA-MambaUNet-S	1.95M	14.91G	2.936 $\pm$ 0.006	7.350 $\pm$ 0.013	0.789 $\pm$ 0.002	15.453 $\pm$ 0.065	3.042 $\pm$ 0.020	9.670 $\pm$ 0.024	0.930 $\pm$ 0.001	15.047 $\pm$ 0.079	2.033 $\pm$ 0.030	3.334 $\pm$ 0.069	0.740 $\pm$ 0.008	8.946 $\pm$ 0.297	
RWSA-MambaUNet-M	3.91M	28.47G	3.001 $\pm$ 0.006	7.490 $\pm$ 0.113	0.800 $\pm$ 0.002	16.017 $\pm$ 0.085	3.126 $\pm$ 0.011	10.019 $\pm$ 0.074	0.936 $\pm$ 0.001	15.600 $\pm$ 0.065	2.101 $\pm$ 0.0011	3.690 $\pm$ 0.054	0.763 $\pm$ 0.005	9.198 $\pm$ 0.250	

Table 2: Ablation study. Default configurations for our RWSA-MambaUNet is with RWSA, and with T- and F-MHA modules in the MambaAttention blocks. Models are trained on VB-DemandEx. Best reported mean is marked bold.

Dataset			In-Domain				Out-Of-Domain							
			VB-DemandEx				DNS 2020				EARS-WHAM_v2			
Model	Params	FLOPs	PESQ	SSNR	ESTOI	SI-SDR	PESQ	SSNR	ESTOI	SI-SDR	PESQ	SSNR	ESTOI	SI-SDR
Noisy	-	-	1.625	-1.068	0.630	4.976	1.582	6.218	0.810	9.071	1.235	-0.803	0.640	5.359
RWSA-MambaUNet-S	1.95M	14.91G	2.936 _{-0.006}	7.350 _{+0.013}	0.789 _{-0.002}	15.453 _{+0.065}	3.042 _{-0.020}	9.670 _{-0.024}	0.930 _{-0.001}	15.047 _{-0.079}	2.033 _{-0.030}	3.334 _{-0.069}	0.740 _{-0.008}	8.946 _{+0.297}
w/o RWSA	1.98M	14.91G	2.906 _{+0.017}	7.119 _{+0.004}	0.782 _{-0.002}	15.275 _{+0.124}	2.956 _{+0.026}	9.461 _{+0.030}	0.924 _{+0.001}	14.838 _{+0.210}	1.957 _{+0.031}	3.010 _{+0.097}	0.731 _{+0.003}	8.448 _{+0.161}
w/o MHA modules [13]	1.88M	14.45G	2.915 _{+0.021}	7.162 _{+0.091}	0.786 _{+0.003}	15.456 _{+0.116}	2.932 _{+0.009}	9.308 _{+0.109}	0.922 _{+0.001}	14.757 _{+0.035}	1.922 _{+0.024}	3.096 _{+0.069}	0.714 _{+0.010}	7.901 _{+0.280}

all metrics. Finally, as shown in Table 1, by increasing the number of channels from 16 to 24, our RWSA-MambaUNet-M outperforms all baselines across all used evaluations metrics on both out-of-domain test sets. While bigger in parameter count, RWSA-MambaUNet-M still requires significantly less FLOPs compared to the baselines. We observed no performance gains by further increasing the model size.

Interestingly, compared to the baselines, we only observe consistent SI-SDR improvements on the out-of-domain EARS-WHAM_v2 test set, which is the only dataset used, where the clean references are recorded in an anechoic chamber [26]. In comparison, our RWSA-MambaUNet models slightly underperform across the SI-SDR metric on the in-domain VB-DemandEx and out-of-domain DNS 2020 test sets. We attribute the performance differences across datasets to the characteristics of the reference signals. This behaviour aligns with the findings of [32], indicating that our RWSA-MambaUNet models primarily learn to reconstruct clean speech.

4.2. Ablation study

To understand the performance impact of key aspects of our RWSA-MambaUNet models, we conduct an ablation study on the RWSA and the shared T- and F-MHA modules in the MambaAttention blocks, which are used inside our RWSA-MambaUNet models. Since ablations in [18] already demonstrated the positive performance impact of sharing the parameters of the T- and F-MHA modules in the MambaAttention blocks, we omit this ablation.

The ablation study in Table 2 reveals that removing RWSA from the RWSA-MambaUNet-S model decreases both cross-corpus generalization performance and in-domain performance while slightly increasing model size. Removing the MHA modules from the MambaAttention blocks reduces our RWSA-MambaUNet model to Mamba-SEUNet [13]. From Table 2, we observe that removing the MHA modules negatively affects generalization performance, as all metrics across both out-of-domain test sets decrease.

4.3. Results on DNS 2020

To investigate the scalability of our proposed RWSA-MambaUNet models with respect to training dataset size and diversity, we train them on the large-scale DNS 2020 dataset.

Table 4 reveals that our RWSA-MambaUNet-XS matches or outperforms the xLSTM-SENet, SEMamba, and MP-SENet baselines on the ESTOI metric, at less than half the parameter count. Moreover, our RWSA-MambaUNet-S slightly outperforms all baselines except the SOTA MambaAttention model on the PESQ and ESTOI metric, while delivering a similar SSNR score to SEMamba. Finally, our RWSA-MambaUNet-M matches or outperforms all baselines across SSNR and ESTOI, with a lower computational complexity. MambaAttention remains slightly superior for in-domain speech enhancement performance as shown in Table 1 and Table 4.

Table 4: Speech Enhancement performance on DNS 2020. All baselines are trained using their originally provided code. Best reported mean is marked bold.

Model	Params	FLOPs	PESQ	SSNR	ESTOI	SI-SDR
Noisy	-	-	1.582	6.218	0.810	9.071
xLSTM-SENet [14]	2.20M	80.71G	3.588 $\pm$ 0.017	14.526 $\pm$ 0.482	0.954 $\pm$ 0.001	20.854 $\pm$ 0.226
LSTM-SENet [14]	2.34M	88.59G	3.598 $\pm$ 0.031	15.021 $\pm$ 0.168	0.956 $\pm$ 0.002	21.003 $\pm$ 0.215
SEMamba [12]	2.25M	65.46G	3.594 $\pm$ 0.012	14.830 $\pm$ 0.473	0.955 $\pm$ 0.001	21.035 $\pm$ 0.123
MP-SENet [8]	2.05M	74.29G	3.605 $\pm$ 0.021	14.967 $\pm$ 0.044	0.954 $\pm$ 0.000	20.919 $\pm$ 0.021
MambaAttention [18]	2.33M	65.52G	3.671 $\pm$ 0.008	15.116 $\pm$ 0.049	0.959 $\pm$ 0.000	21.234 $\pm$ 0.033
RWSA-MambaUNet-XS	1.02M	9.22G	3.563 $\pm$ 0.002	14.685 $\pm$ 0.039	0.955 $\pm$ 0.000	20.457 $\pm$ 0.016
RWSA-MambaUNet-S	1.95M	14.91G	3.614 $\pm$ 0.009	14.869 $\pm$ 0.091	0.957 $\pm$ 0.000	20.798 $\pm$ 0.049
RWSA-MambaUNet-M	3.91M	28.47G	3.649 $\pm$ 0.017	15.119 $\pm$ 0.069	0.959 $\pm$ 0.000	21.119 $\pm$ 0.094

5. CONCLUSION

In this paper, we proposed a novel and efficient hybrid RWSA-MambaUNet model for improved cross-corpus speech enhancement performance. Experiments revealed that our best-performing RWSA-MambaUNet model significantly outperforms existing baselines on two very different out-of-domain test sets across all evaluation metrics. Notably, even our smallest RWSA-MambaUNet model outperforms existing state-of-the-art models across most metrics on two out-of-domain corpora, while using significantly fewer parameters and FLOPs. Finally, we demonstrated that resolution-wise shared attention contributes to the superior cross-corpus enhancement performance of our RWSA-MambaUNet models.

6. REFERENCES

- [1] Y. Xie and Z.-H. Tan, “A survey of deep learning for complex speech spectrograms,” *arXiv preprint arXiv:2505.08694*, 2025.
- [2] S.-W. Fu, Y. Tsao, and X. Lu, “Snr-aware convolutional neural network modeling for speech enhancement,” in *Proc. Inter-speech*, 2016, pp. 3768–3772.
- [3] M. Kolbæk, Z.-H. Tan, S. H. Jensen, and J. Jensen, “On loss functions for supervised monaural time-domain speech enhancement,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 825–838, 2020.
- [4] Y.-J. Lu, Z.-Q. Wang, S. Watanabe, A. Richard, C. Yu, and Y. Tsao, “Conditional diffusion probabilistic model for speech enhancement,” in *Proc. IEEE ICASSP*, 2022, pp. 7402–7406.
- [5] J. Richter, S. Welker, J.-M. Lemerrier, B. Lay, and T. Gerkmann, “Speech enhancement and dereverberation with diffusion-based generative models,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 31, pp. 2351–2364, 2023.
- [6] D. Michelsanti and Z.-H. Tan, “Conditional generative adversarial networks for speech enhancement and noise-robust speaker verification,” in *Proc. Interspeech*, 2017, pp. 2008–2012.
- [7] S.-W. Fu, C.-F. Liao, Y. Tsao, and S.-D. Lin, “Metricgan: Generative adversarial networks based black-box metric scores optimization for speech enhancement,” in *Proc. ICML*, 2019, pp. 2031–2041.
- [8] Y.-X. Lu, Y. Ai, and Z.-H. Ling, “Mp-senet: A speech enhancement model with parallel denoising of magnitude and phase spectra,” in *Proc. Interspeech*, 2023, pp. 3834–3838.
- [9] Y.-X. Lu, Y. Ai, and Z.-H. Ling, “Explicit estimation of magnitude and phase spectra in parallel for high-quality speech enhancement,” *Neural Networks*, vol. 189, pp. 107562, 2025.
- [10] A. Gu and T. Dao, “Mamba: Linear-time sequence modeling with selective state spaces,” in *Proc. COLM*, 2024.
- [11] M. Beck, K. Pöppel, M. Spanring, A. Auer, O. Prudnikova, M. K. Kopp, G. Klambauer, J. Brandstetter, and S. Hochreiter, “xLSTM: Extended long short-term memory,” in *Proc. NeurIPS*, 2024.
- [12] R. Chao et al., “An investigation of incorporating mamba for speech enhancement,” in *Proc. IEEE SLT*, 2024, pp. 302–308.
- [13] J. Wang, Z. Lin, T. Wang, M. Ge, L. Wang, and J. Dang, “Mamba-seunet: Mamba unet for monaural speech enhancement,” in *Proc. IEEE ICASSP*, 2025, pp. 1–5.
- [14] 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.
- [15] Z. Lin, X. Chen, and J. Wang, “Muse: Flexible voiceprint receptive fields and multi-path fusion enhanced taylor transformer for u-net-based speech enhancement,” in *Proc. Inter-speech*, 2024, pp. 672–676.
- [16] A. Pandey and D. Wang, “On cross-corpus generalization of deep learning based speech enhancement,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 2489–2499, 2020.
- [17] A. Pandey and D. Wang, “Self-attending rnn for speech enhancement to improve cross-corpus generalization,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 1374–1385, 2022.
- [18] N. L. Kühne, J. Jensen, J. Østergaard, and Z.-H. Tan, “MambaAttention: Mamba with Multi-Head Attention for Generalizable Single-Channel Speech Enhancement,” *arXiv preprint arXiv:2507.00966*, 2025.
- [19] S. Wisdom et al., “Differentiable consistency constraints for improved deep speech enhancement,” in *Proc. IEEE ICASSP*, 2019, pp. 900–904.
- [20] A. Pandey and D. Wang, “Densely connected neural network with dilated convolutions for real-time speech enhancement in the time domain,” in *Proc. IEEE ICASSP*, 2020, pp. 6629–6633.
- [21] Y. Qiu, K. Zhang, C. Wang, W. Luo, H. Li, and Z. Jin, “Mbtaylorformer: Multi-branch efficient transformer expanded by taylor formula for image dehazing,” in *Proc. IEEE/CVF ICCV*, 2023, pp. 12802–12813.
- [22] W. Shi et al., “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network,” in *Proc. CVPR*, 2016, pp. 1874–1883.
- [23] C. Veaux, J. Yamagishi, and S. King, “The voice bank corpus: Design, collection and data analysis of a large regional accent speech database,” in *Proc. IEEE O-COCOSDA/CASLRE*, 2013, pp. 1–4.
- [24] J. Thiemann, N. Ito, and E. Vincent, “The diverse environments multi-channel acoustic noise database (demand): A database of multichannel environmental noise recordings,” in *Proceedings of Meetings on Acoustics*. AIP Publishing, 2013, vol. 19.
- [25] C. K.A. Reddy, V. Gopal, R. Cutler, E. Beyrami, R. Cheng, H. Dubey, S. Matushevych, R. Aichner, A. Aazami, S. Braun, P. Rana, S. Srinivasan, and J. Gehrke, “The interspeech 2020 deep noise suppression challenge: Datasets, subjective testing framework, and challenge results,” in *Proc. Interspeech*, 2020, pp. 2492–2496.
- [26] J. Richter, Y.-C. Wu, S. Krenn, S. Welker, B. Lay, S. Watanabe, A. Richard, and T. Gerkmann, “EARS: An anechoic fullband speech dataset benchmarked for speech enhancement and dereverberation,” in *Proc. Interspeech*, 2024, pp. 4873–4877.
- [27] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. Le Roux, “WHAM!: Extending speech separation to noisy environments,” in *Proc. Interspeech*, 2019, pp. 1368–1372.
- [28] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, “Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs,” in *Proc. IEEE ICASSP*, 2001, vol. 2, pp. 749–752.
- [29] J. H. L. Hansen and B. L. Pellom, “An effective quality evaluation protocol for speech enhancement algorithms,” in *Proc. ICSLP*. Citeseer, 1998, vol. 7, pp. 2819–2822.
- [30] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “Sdr – half-baked or well done?,” in *Proc. IEEE ICASSP*, 2019, pp. 626–630.
- [31] J. Jensen and C. H. Taal, “An algorithm for predicting the intelligibility of speech masked by modulated noise maskers,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 11, pp. 2009–2022, 2016.
- [32] S. D. Jepsen, M. G. Christensen, and J. R. Jensen, “A study of the scale invariant signal to distortion ratio in speech separation with noisy references,” *Accepted to IEEE ASRU*, 2025.