



A Multi-dimensional Deep Structured State Space Approach to Speech Enhancement Using Small-footprint Models

Pin-Jui Ku¹, Chao-Han Huck Yang¹, Sabato Marco Siniscalchi^{1,2,3}, Chin-Hui Lee¹

¹Georgia Institute of Technology, U.S.A

²Kore University of Enna, Italy

³Norwegian University of Science and Technology, Norway

{pku9, chl}@gatech.edu

Abstract

We propose a multi-dimensional structured state space (S4) approach to speech enhancement. To better capture the spectral dependencies across the frequency axis, we focus on modifying the multi-dimensional S4 layer with whitening transformation to build new small-footprint models that also achieve good performance. We explore several S4-based deep architectures in time (T) and time-frequency (TF) domains. The 2-D S4 layer can be considered a particular convolutional layer with an infinite receptive field although it utilizes fewer parameters than a conventional convolutional layer. Evaluated on the VoiceBank-DEMAND data set, when compared with the conventional U-net model based on convolutional layers, the proposed TF-domain S4-based model is 78.6% smaller in size, yet it still achieves competitive results with a PESQ score of 3.15 with data augmentation. By increasing the model size, we can even reach a PESQ score of 3.18.

Index Terms: speech enhancement, deep structure state structure model, data augmentation, U-net structure, DCCRN

1. Introduction

Speech enhancement aims at reducing noise to improve the quality and intelligibility of noisy speech. It also serves as a front-end in robust speech recognition [1] and speaker recognition [2] in adverse environments. Conventional approaches to speech enhancement usually rely on using statistical models to predict speech and noise, e.g., spectral subtraction [3], Wiener filtering [4], and Kalman filtering [5, 6]. In recent years, deep neural network (DNN) based approaches have demonstrated superiority over conventional signal processing techniques, especially for alleviating problems caused by musical and non-stationary noise [7]. They can be roughly categorized according to the features used. In the time-frequency (TF) domain, two distinct types exist: (i) spectral mapping [8, 9], by establishing a regression function directly from mapping the noisy speech spectrogram to the clean one; and (ii) masking [10, 11] by estimating a 2-D ratio mask that describes the TF relationships between clean speech and background noise. More recently, generative models, aiming at learning a prior distribution over clean speech data, have been proposed, e.g., generative adversarial networks (GAN) [12, 13, 14], variational autoencoders (VAE) [15, 16], and diffusion models [17, 18, 19].

DNNs [9, 20, 21], recurrent neural networks (RNNs) [22, 23, 24], and convolutional neural networks (CNNs) [25, 26] have been proposed for speech enhancement. While RNNs are a natural choice for sequential inputs, their recursive nature makes them slow to train and suffer from optimization difficulties, such as the “vanishing gradient” problem [27], limiting their ability to handle long sequences. In contrast, CNNs

encode local context and enable parallel training, but they are fundamentally constrained by the size of their receptive field and may not achieve global coherence. Recently, deep state space models (SSM) [28], in particular S4-based ones [29], have achieved state-of-the-art results in modeling sequence data with extremely long-range dependencies. The S4 layer unifies the strengths of both RNN and CNN layers: (i) it allows parallel computing like CNN layers, and (ii) it can capture global information across the whole input sequence like RNNs. S4-based systems have been deployed for raw audio generation [30] and speech recognition [31].

In this study, we propose a structured state-space-based architecture in the multi-dimensional TF domain for speech enhancement. We explore 2-D extensions since the original S4-based layer was designed to handle only 1-D input signals. However, to adequately capture the spectral dependencies across the frequency axis, we first compute the covariance matrix of each frequency bin and apply a whitening technique to make the frequency bins statistically uncorrelated. Unfortunately, while whitening was effective under specific scenarios, the overall performance improvement was limited. We therefore utilize the multi-dimensional S4 layer proposed in [32] and adapt it to the TF domain. Experimental evidence shows that by injecting the multi-dimensional S4 layer inside a U-net-like architecture, we were able to deploy a novel deep enhancement model that attained better PESQ scores than the original time-domain model while reducing the model size by 78.6% (with only 0.75M parameters). Since S4 layer directly handles 1-D inputs, we have also built a sequence-to-sequence regression model in the time-domain. Furthermore, we have also evaluated the effect of data augmentation on speech enhancement.

In the following sections, we will demonstrate three main contributions. We first conduct an initial investigation on how to best incorporate an S4 layer into speech enhancement models. To the best of our knowledge, this is the first attempt in this field. Next, we find that TF-domain models can be improved using a small hop-length to extract spectrograms with long time lengths through short-time Fourier transform (STFT) mainly because S4 can easily handle long time-series inputs. Finally, by unifying the strengths of RNN and CNN layers, we achieve competitive enhancement results even with a compact TF-domain S4-based U-Net architecture

2. Related Work

2.1. SSM - State Space Modeling

The S4 layer is based on the linear state space layer (LSSL) proposed in [28], where $u(t) \in \mathcal{C}$, and $v(t) \in \mathcal{C}$ are 1-dimensional input/output patterns in state space modeling, respectively, and $x(t) \in \mathcal{C}^N$ is an implicit N -dimensional state vector. LSSL de-

defines a 1-D mapping function $u(t) \rightarrow v(t)$ leveraging upon the following ordinary differential equations (ODEs):

$$\begin{aligned} \dot{x}(t) &= \mathbf{A}x(t) + \mathbf{B}u(t) \\ v(t) &= \mathbf{C}x(t) + \mathbf{D}u(t) \end{aligned} \quad (1)$$

Eq. (1) describes a state space model (SSM). The goal of LSSL is to treat the SSM as a black-box representation in a deep model, where the state-transition matrix $\mathbf{A} \in \mathcal{C}^{N \times N}$ and the projection matrices $\mathbf{B} \in \mathcal{C}^{1 \times N}$, $\mathbf{C} \in \mathcal{C}^{N \times 1}$, $\mathbf{D} \in \mathcal{C}^{1 \times 1}$ are parameters that could be learned by gradient descent. To use LSSL in a discrete sequence-to-sequence model, the bilinear method [33] can be used to discretize all continuous-time signals as follows:

$$\begin{aligned} x_k &= \bar{\mathbf{A}}x_{k-1} + \bar{\mathbf{B}}u_k \\ v_k &= \mathbf{C}x_k + \mathbf{D}u_k \\ \bar{\mathbf{A}} &= (\mathbf{I} - \Delta/2 \cdot \mathbf{A})^{-1}(\mathbf{I} + \Delta/2 \cdot \mathbf{A}) \\ \bar{\mathbf{B}} &= (\mathbf{I} - \Delta/2 \cdot \mathbf{A})^{-1}\Delta\mathbf{B} \end{aligned} \quad (2)$$

where Δ is a trainable time-step size parameter. Since Eq. (2) becomes a recurrence in the hidden state vector x_k , LSSL can now be considered a special RNN with linearity. By setting the initial state to be $x_{-1} = 0$ and unfolding Eq. refeq:continuous-SSM):

$$v_k = \mathbf{C}\bar{\mathbf{A}}^k\bar{\mathbf{B}}u_0 + \mathbf{C}\bar{\mathbf{A}}^{k-1}\bar{\mathbf{B}}u_1 + \dots + \mathbf{C}\bar{\mathbf{B}}u_k + \mathbf{D}u_k \quad (3)$$

That is, the output sequence $\mathbf{v} = (v_0, v_1, \dots, v_{L-1})$ with length L could be computed as a convolution $\mathbf{v} = \bar{\mathbf{K}} * \mathbf{u} + \mathbf{D}\mathbf{u}$, where $\mathbf{u} = (u_0, u_1, \dots, u_{L-1})$ and $\bar{\mathbf{K}}$ is called an SSM convolution kernel defined as a 1-D sequence:

$$\bar{\mathbf{K}} := (\bar{\mathbf{K}}_0, \bar{\mathbf{K}}_1, \dots, \bar{\mathbf{K}}_{L-1}) = (\mathbf{C}\bar{\mathbf{B}}, \mathbf{C}\bar{\mathbf{A}}\bar{\mathbf{B}}, \dots, \mathbf{C}\bar{\mathbf{A}}^{L-1}\bar{\mathbf{B}}) \quad (4)$$

Therefore, LSSL can also be viewed as a convolutional layer with a global receptive field [28, 29], and it can be computed very efficiently with fast Fourier transform once the SSM convolutional kernel, $\bar{\mathbf{K}}$, is known.

One major bottleneck of LSSL is that treating \mathbf{A} as trainable parameters means computing $\bar{\mathbf{K}}$ as many times as it is necessary in the training stage. However, it is challenging to compute $\bar{\mathbf{A}}^i$ efficiently enough for practical usage without making the whole model unstable. The S4 layer overcomes this critical issue by decomposing the state transition matrix \mathbf{A} into a sum of a low-rank matrix [34] and a skew-symmetric matrix [35]. In practice, S4 re-parameterizes the state-transition matrices \mathbf{A} as $\mathbf{A} = \mathbf{\Lambda} - \mathbf{P}\mathbf{P}^*$, where $\mathbf{\Lambda} \in \mathcal{C}^N$ is a diagonal matrix and $\mathbf{P} \in \mathcal{C}^N$. In S4, \mathbf{D} is set equal to $\mathbf{0}$ since it could be replaced by a residual connection. Therefore, an S4 layer is comprised of 4N trainable matrix parameters: $\mathbf{\Lambda}$, \mathbf{P} , \mathbf{B} , and \mathbf{C} . The interested reader is referred to [29, 30] for more details on S4.

2.2. SSM with Multi-dimensional Patterns

The S4 layer was developed for 1-D inputs, which limited its applicability. In [28, 29, 30], that limitation was overcome by (i) running H independent copies of the S4 layer on a 2-D input features with a shape of (H, L) , and (ii) mixing all output features by a position-wise feedforward layer. In practice, the 2-D input is assumed to consist of H independent signals, analogous to a 1-D CNN layer with H channels. However, most of the 2-D inputs are correlated with each other in both axes. For example, a 2-D spectrogram is not only time-dependent but also has strong spectral dependencies across the frequency axis. In [32],

the conventional S4 layer was extended to multi-dimensional signals by turning the standard SSM (1-D ODEs) into multi-dimensional partial differential equations (PDEs) governed by an independent SSM in each dimension. To make it clear, let $\mathbf{u} = u(t_1, t_2) \in \mathcal{R}^2 \rightarrow \mathcal{C}$ and $\mathbf{v} = v(t_1, t_2) \in \mathcal{R}^2 \rightarrow \mathcal{C}$ be the input and output signals, and $x(t_1, t_2)$ be the SSM state with dimension $N_1 \times N_2$, which equals to the outer product of $x_1(t_1, t_2) \in \mathcal{C}^{N_1}$ and $x_2(t_1, t_2) \in \mathcal{C}^{N_2}$. A 2-D SSM thus can be represented by the following two-variable PDEs:

$$\begin{aligned} \frac{\partial}{\partial t_1}x(t_1, t_2) &= \mathbf{A}_1x_1(t_1, t_2) \otimes x_2(t_1, t_2) + \mathbf{B}_1u(t_1, t_2) \\ \frac{\partial}{\partial t_2}x(t_1, t_2) &= x_1(t_1, t_2) \otimes \mathbf{A}_2x_2(t_1, t_2) + \mathbf{B}_2u(t_1, t_2) \\ v(t_1, t_2) &= \langle \mathbf{C}, x(t_1, t_2) \rangle \end{aligned} \quad (5)$$

where \otimes is the outer-product operator, $\langle \cdot, \cdot \rangle$ is the point-wise inner-product operators for two matrices. Factoring that matrix \mathbf{C} as a low-rank tensor and using the standard 1-D S4 layer as a black box, a 2D-version S4 layer is equivalent to a multi-dimensional convolution with an infinite receptive field. The multi-dimensional S4 layer is referred to as S4ND [32]. The same bilinear method mentioned in Sec. 2.1 can be applied when handling 2D discrete-time inputs.

3. Proposed Deep Structured State Space Modeling for Speech Enhancement

3.1. Time-domain S4-based Model

The original S4-based model in [30] can be used as a time-domain regression SE model, consisting of repeated S4-block combined with a U-net architecture. This model is referred to as *Time-domain S4 U-Net*. A 1-D convolutional layer with kernel size = 1 will first turn the single-channel noisy speech into a hidden feature with 64 channels before it goes through the whole U-net path. Another single-kernel 1-D convolutional layer will instead shrink the channel size back to 1, which represents the enhanced speech. More details can be found [30].

The time-domain S4 U-Net is optimized by the loss function proposed in [26], which is the L1 loss over the waveform together with a multi-resolution short-time Fourier transform (STFT) loss over the spectrogram magnitude [36]:

$$\begin{aligned} \text{loss}_{STFT}(y, \tilde{y}) &= \frac{||STFT(y)| - |STFT(\tilde{y})||_F}{||STFT(y)||_F} \\ &+ \frac{1}{T} ||\log(|STFT(y)|) - \log(|STFT(\tilde{y})||_1 \end{aligned} \quad (6)$$

$$\text{loss}(y, \tilde{y}) = \frac{1}{T} ||y - \tilde{y}||_1 + \frac{1}{M} \sum_{i=1}^M \text{loss}_{STFT}^{(i)}(y, \tilde{y}) \quad (7)$$

where STFT stands for short-time Fourier transform, $|\cdot|_F$ is the Frobenius norm, y, \tilde{y} are the clean and estimated speech, respectively, and M means the total number of different STFT settings when computing $\text{loss}_{STFT}(y, \tilde{y})$. We use the same three STFT settings to compute $\text{loss}_{STFT}^{(i)}(y, \tilde{y})$ in [26].

3.2. TF-domain S4-based Model

We deploy two S4-based SE models in the TF-domain, namely (i) only the magnitude of the spectrogram is enhanced and then the noisy phase is used to reconstruct the speech waveform, or (ii) the complex noisy spectrogram is used to directly estimate both the magnitude and phase information simultaneously. In

the magnitude-only scenario, two variants are put forth: a *regression* model, which maps noisy spectrograms to clean spectrograms, and a *masking* model which predicts a spectrogram mask. In the complex scenario, we adopt the masking method for the complex spectrogram proposed in [37, 38] since the regression-based SE model led to an unstable training process in the complex case. In sum, there are three training scenarios: *mag-regression*, *mag-masking*, and *complex-masking*.

L1 loss is employed to optimize our TF-domain deep architecture, as suggested in [39]. More specifically, the following loss is used when only the magnitude information is taken into account (i.e., *mag-regression* and *mag-masking*):

$$\text{Loss}_{mag} = \frac{1}{TF} \sum_{t=0}^{T-1} \sum_{f=0}^{F-1} |S(t, f) - \tilde{S}(t, f)| \quad (8)$$

where S and \tilde{S} are the magnitude spectrograms of the clean speech and enhanced speech, respectively, and T and F are the numbers of frames and frequency bins, respectively.

When the complex spectrogram is used as the input/output, the loss function becomes:

$$\begin{aligned} \text{Loss}_{complex} = & \frac{1}{TF} \sum_{t=0}^{T-1} \sum_{f=0}^{F-1} |S_r(t, f) - \tilde{S}_r(t, f)| \\ & + |S_i(t, f) - \tilde{S}_i(t, f)| + \text{Loss}_{mag} \end{aligned} \quad (9)$$

where r , and i denote the real and imaginary parts of the complex-valued spectrogram, respectively.

3.2.1. 1-D S4-based Model with Whitening Transform

As discussed in Section 2.2, we could directly adopt the 1-D S4-based layer to build a deep SE model working on 2-D spectrograms by treating frequency bins as a group of independent 1-D signals. In doing so, the deep SE model can not exploit dependencies across the frequency axis. However, we introduced a whitening transformation [40] to the input spectrogram so that the frequency bins can be considered statistically uncorrelated to one another. This system is called *TF-domain S4 U-Net*.

3.2.2. Multi-dimensional S4-based Model

Although the whitening transformation can theoretically alleviate the frequency-dependent problem, it doesn't really help the model capture the 2-D information as well as a 2-D CNN layer does. Our experiments (see Sec. 4.2) show that while adding the whitening transformation is helpful, it doesn't perform better than the time-domain S4-based model but with more model parameters. For that reason, we build a new SE model based on the S4ND layer, which is named *S4ND U-Net*.

Fig. 1 illustrates how the proposed *S4ND U-Net* is built. *S4ND U-Net* is similar to the Deep Complex Convolutional Recurrent Network (DCCRN) model [38], which is a U-net-like architecture with several down-layers and up-layers, where we have however replaced all complex CNN/RNN layers with S4ND layers. Furthermore, we have reduced the model complexity by shrinking the number of down-layers and up-layers to 2 (the original DCCRN model had 6 layers), since it is unnecessary to stack too many S4ND layers to increase the receptive field. With complex spectrograms, we stack the real and imaginary parts as a real-value spectrogram with two channels.

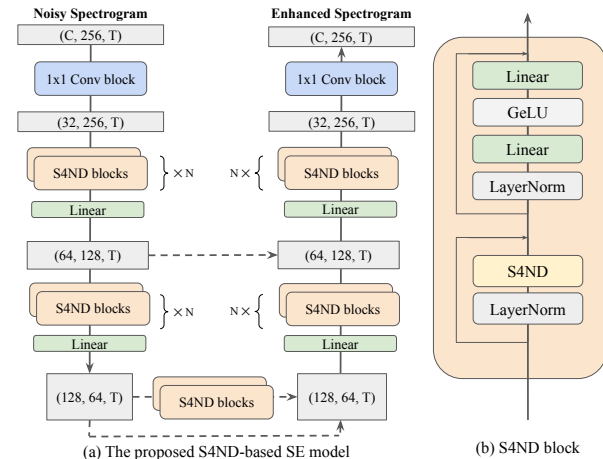


Figure 1: (a) The proposed *S4ND U-Net*. C is either 1 or 2, depending on whether the input is a magnitude spectrogram ($C = 1$) or a complex spectrogram ($C = 2$), and N is the number of the S4 blocks stacked per down/up layer. $N=4$ is used in all experiments. (b) The detailed structure of an *S4ND* block. The original S4 block's design [30] was used but S4 was replaced by *S4ND*.

4. Experiments and Result Analyses

4.1. Experimental Setup

4.1.1. Dataset

We choose the VoiceBank-DEMAND dataset [41], a common SE benchmark data set providing recordings from 30 speakers with 10 types of noise, to assess our models. The data set is split into a training and a testing set with 28 and 2 speakers, respectively. Four types of signal-to-noise ratios (SNRs) are used to mix clean samples with noise samples in the dataset, [0, 5, 10, 15] dB for training and [2.5, 7.5, 12.5, 17.5] dB for testing. We follow the approach used by Lu et al. [17] to form the validation set by excerpting two speakers from the training set. This resulted in 10,802 utterances for training and 770 for validation. The testing set included a total of 824 utterances. All recordings were downsampled from 48 kHz to 16 kHz. Our evaluation metrics include the wide-band perceptual evaluation of speech quality (PESQ) [42], prediction of the signal distortion (CSIG), prediction of the background intrusiveness (CBAK), and prediction of the overall speech quality (COVL).

4.1.2. STFT feature transformation and data augmentation

During training, an amplitude transformation as in [18] is used to normalize the amplitude of all complex coefficients of both noisy and clean spectrograms. Two out of four data augmentation methods in [26], namely, *Remix* and *BandMask* are used. *Remix* shuffles the noises within one batch to form new noisy mixtures. *BandMask* is a band-stop filter that removes 20% of the frequencies starting from f_0 , a random frequency sampled uniformly in the mel scale.

4.2. Experimental Results

In Table 1, we first assess our baseline *time-domain S4 U-Net* discussed in Sec. 3.1 and compare it with other time-domain techniques. We build *time-domain S4 U-Net* according to the default hyper-parameter setting in [29]. A visual inspection of

Table 1: Evaluation results of different time-domain U-Net models. Note that we utilize only parts of the data augmentation methods in DEMUCS [26].

Model	Params (Million)	PESQ	CSIG	CBAK	COVL
Noisy	n/a	1.97	3.35	2.44	2.63
Wave-U-Net [43]	10.0	2.40	3.52	3.24	2.96
Attention Wave-U-Net [25]	-	2.62	3.91	3.35	3.27
DEMUCS (small) [26]	18.9	2.93	4.22	3.25	3.52
DEMUCS (large) [26]	33.5	3.07	4.31	3.4	3.63
Time-domain S4 U-Net + data augmentation	7.14	2.97 3.02	4.36 4.45	3.49 3.56	3.65 3.75

Table 2: A comparison of different TF-domain models.

Model	Params (Million)	PESQ	CSIG	CBAK	COVL
Noisy	n/a	1.97	3.35	2.44	2.63
DCCRN [38, 44]	3.7	2.54	3.74	3.13	2.75
S-DCCRN [44]	2.34	2.84	4.03	3.43	2.97
DCU-Net-10 [37]	1.4	2.72	3.74	3.60	3.22
DCU-Net-16 [37]	2.3	2.93	4.10	3.77	3.52
DCU-Net-20 [37]	3.5	3.13	4.24	4.00	3.69
Metric GAN+ [14]	2.6	3.15	4.14	3.16	3.64
TF-domain S4 U-Net + data augmentation + whitening	17.85	2.96 3.05 3.07	4.24 4.34 4.35	3.49 3.54 3.56	3.60 3.69 3.72
S4ND U-Net + data augmentation	0.75	2.99 3.15	4.37 4.52	3.53 3.62	3.70 3.85

Table 1 shows that when compared with other models, time-domain S4 U-Net achieves a good PESQ (with PESQ=3.02 in the bottom row) and the best CSIG, CBAK, and COVL scores. Although DEMUCS (with PESQ=3.07 in the fourth row) has a better PESQ score than *time-domain S4 U-Net*, its performance degrades to 2.93 when its model size shrinks from 33.5 million to 18.9 million. This result verifies that the S4 layers can replace CNN/RNN layers in a DNN-based model to reduce the model sizes while maintaining competitive performances.

Next in Table 2, we compare the proposed *TF-domain S4 U-Net* and *S4ND U-Net* with previously proposed TF-domain models. All networks use a U-Net structure and the same complex-masking approach [37] except Metric GAN+. From Table 2, we can see that S4ND U-Net attains the best results among all models on PESQ/CSIG/COVL metrics with relatively good CBAK scores. It is worth noting that Metric GAN+ optimizes the PESQ score directly (with the best PESQ=3.15 in the sixth row) at the cost of getting worse results on the other metrics. Additionally, DCU-Net-20 attains the best CBAK of 4.00 with PESQ=3.13 in the fifth row. However, its evaluation results significantly drop when halving the number of layers. In contrast, S4ND U-Net uses only 0.75M parameters, corresponding to a 78.6% reduction in model size compared to DCU-Net-20. It also achieves a PESQ score of 3.15 (equal to the best GAN+) as shown in the bottom row.

To have a fair comparison, we retrain DCCRN from scratch with the same data augmentation and amplitude transformation (see Sec. 4.1.2) used to build S4ND U-Net. Two STFT settings (i.e., subscripts "1" and "2" in Table 3) are adopted to investigate how DCCRN and S4ND U-Net react when the input STFT spectrograms have different parameter settings. Furthermore, DCCRN parameters were limited to 1.21M by halving the number of down/up layers from 6 to 3. As shown in Table 3, S4ND U-Net outperforms both the small and large DCCRN under the first STFT setting, with PESQ scores of 3.15 and 3.18 in the fifth

Table 3: Evaluation results of models with two sizes (S & M) and two STFT settings (subscripts "1" & "2"). "1" means we adopt 510/400/100 for n-fft/win-length/hop-length, while "2" means 510/255/255 was adopted. Note that we set n-fft=510 to make the frequency dimension a ratio of two.

Model (size)	Params (Million)	PESQ	CSIG	CBAK	COVL
DCCRN ₁ (S)	1.19	2.75	3.62	2.99	3.17
DCCRN ₁ (M)	3.48	2.87	3.93	2.59	3.39
DCCRN ₂ (S)	1.19	2.69	3.62	3.35	3.14
DCCRN ₂ (M)	3.48	2.83	3.91	2.51	3.36
S4ND U-Net ₁ (S)	0.75	3.15	4.52	3.62	3.85
S4ND U-Net ₁ (M)	4.44	3.18	4.49	3.63	3.85
S4ND U-Net ₂ (S)	0.75	3.09	4.44	3.60	3.78
S4ND U-Net ₂ (M)	4.44	3.11	4.46	3.59	3.79

Table 4: A comparison of three TF-domain S4ND U-Nets.

Scenario	PESQ	CSIG	CBAK	COVL
<i>mag-regression</i>	3.05	4.44	3.47	3.76
<i>mag-masking</i>	3.12	4.51	3.58	3.83
<i>complex-masking</i>	3.15	4.52	3.62	3.85

and sixth rows, respectively. When a bigger hop-length=255 is chosen and the time-length of input spectrograms decreases correspondingly (i.e., the second STFT setting), the S4ND U-Net's performance drops a little bit (with PESQ=3.09 and 3.11 in the bottom two rows) but remains superior to DCCRN. This suggests our S4ND U-Net works particularly well when dealing with long input sequences. With a large S4ND U-Net with 4.4M parameters by increasing the number of down/up layers from 2 to 5, We only improve PESQ slightly from 3.15 to 3.18. This confirms there is no need to stack too many S4ND layers to increase the receptive field. Even without data augmentation, S4ND U-Net (with PESQ=2.99 shown in the second last row of Table 2) still outperforms both DCCRN and DCU-Net-16.

Finally, we assess S4ND U-Net under three different TF-domain scenarios, namely *mag-regression*, *mag-masking*, and *complex-masking*. Results in Table 4 show that *complex-masking* attains a better PESQ score of 3.15 as shown in the bottom row when compared to 3.05 in the top and 3.12 in the middle rows. It is noted that although complex masking manages to achieve the best result among the three TF-domain scenarios compared here, models with only enhanced magnitudes plus noisy phase information for waveform reconstruction also perform well using the proposed S4ND U-net architecture.

5. Conclusion

We have conducted a series of experiments to investigate new uses of S4 layers for speech enhancement. We first develop an S4-based deep SE neural model to enhance speech in the time domain and attain promising results. To extend the approach to multi-dimensional inputs, we propose two techniques: a whitening transformation and an S4ND U-Net architecture. We found that our proposed S4ND U-Net not only outperforms the original time-domain S4 U-Net but also attains comparable or better scores in four evaluation metrics when contrasted with other time-domain and TF-domain U-Net models. Our results also empirically verify that fewer S4ND layers can be used for compact model design, resulting in a robust TF-domain SE model with a limited size. Codes used in this work are released at <https://github.com/Kuray107/S4ND-U-Net-speech-enhancement>

6. References

- [1] B. Wu, K. Li, F. Ge, Z. Huang, M. Yang, S. M. Siniscalchi, and C.-H. Lee, "An end-to-end deep learning approach to simultaneous speech dereverberation and acoustic modeling for robust speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, pp. 1289–1300, 2017.
- [2] J. H. Hansen and T. Hasan, "Speaker recognition by machines and humans: A tutorial review," *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 74–99, 2015.
- [3] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. 27, no. 2, pp. 113–120, 1979.
- [4] J. Lim and A. Oppenheim, "Enhancement and bandwidth compression of noisy speech," *Proceedings of the IEEE*, vol. 67, no. 12, pp. 1586–1604, 1979.
- [5] K. Paliwal and A. Basu, "A speech enhancement method based on kalman filtering," in *Proc. ICASSP*, 1987.
- [6] V. Krishnan, S. Siniscalchi, D. Anderson, and M. Clements, "Noise robust aurora-2 speech recognition employing a codebook-constrained kalman filter preprocessor," in *Proc. ICASSP*, 2006.
- [7] M. Berouti, R. Schwartz, and J. Makhoul, "Enhancement of speech corrupted by acoustic noise," in *Proc. ICASSP*, 1979.
- [8] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, "An experimental study on speech enhancement based on deep neural networks," *IEEE Signal Processing Letters*, vol. 21, no. 1, pp. 65–68, 2014.
- [9] —, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Trans. Audio, Speech, Language Process*, vol. 23, no. 1, pp. 7–19, 2015.
- [10] D. Wang, "Time-frequency masking for speech separation and its potential for hearing aid design," *Trends in amplification*, vol. 12, no. 4, pp. 332–353, 2008.
- [11] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Trans. Audio, Speech, Language Process*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [12] S. Pascual, A. Bonafonte, and J. Serrà, "Segan: Speech enhancement generative adversarial network," in *Proc. Interspeech*, 2017.
- [13] 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.
- [14] S.-W. Fu, C. Yu, T.-A. Hsieh, P. Plantinga, M. Ravanelli, X. Lu, and Y. Tsao, "MetricGAN+: An Improved Version of MetricGAN for Speech Enhancement," in *Proc. Interspeech*, 2021.
- [15] Y. Bando, K. Sekiguchi, and K. Yoshii, "Adaptive Neural Speech Enhancement with a Denoising Variational Autoencoder," in *Proc. Interspeech*, 2020.
- [16] H. Fang, G. Carbajal, S. Wermter, and T. Gerkmann, "Variational autoencoder for speech enhancement with a noise-aware encoder," in *Proc. ICASSP*, 2021.
- [17] Y.-J. Lu, Z.-Q. Wang, S. Watanabe, A. Richard, C. Yu, and Y. Tsao, "Conditional diffusion probabilistic model for speech enhancement," in *Proc. ICASSP*, 2022.
- [18] J. Richter, S. Welker, J.-M. Lemerrier, B. Lay, and T. Gerkmann, "Speech enhancement and dereverberation with diffusion-based generative models," *arXiv preprint arXiv:2208.05830*, 2022.
- [19] H. Yen, F. G. Germain, G. Wichern, and J. L. Roux, "Cold diffusion for speech enhancement," *arXiv preprint arXiv:2211.02527*, 2022.
- [20] E. M. Grais, M. U. Sen, and H. Erdogan, "Deep neural networks for single channel source separation," in *Proc. ICASSP*, 2014.
- [21] Y. Wang, A. Narayanan, and D. Wang, "On training targets for supervised speech separation," *IEEE/ACM Trans. Audio, Speech, Language Process*, vol. 22, no. 12, pp. 1849–1858, 2014.
- [22] T. Gao, J. Du, L.-R. Dai, and C.-H. Lee, "Densely connected progressive learning for lstm-based speech enhancement," in *Proc. ICASSP*, 2018.
- [23] P.-S. Huang, M. Kim, M. Hasegawa-Johnson, and P. Smaragdis, "Deep learning for monaural speech separation," in *Proc. ICASSP*, 2014.
- [24] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. L. Roux, J. R. Hershey, and B. Schuller, "Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR," in *Proc. LVA/ICA*, 2015.
- [25] R. Giri, U. Isik, and A. Krishnaswamy, "Attention wave-u-net for speech enhancement," in *Proc. WASPAA*, 2019.
- [26] A. Défossez, G. Synnaeve, and Y. Adi, "Real Time Speech Enhancement in the Waveform Domain," in *Proc. Interspeech*, 2020.
- [27] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proc. ICML*, 2013.
- [28] A. Gu, I. Johnson, K. Goel, K. K. Saab, T. Dao, A. Rudra, and C. Re, "Combining recurrent, convolutional, and continuous-time models with linear state space layers," in *Proc. NeuIPS*, 2021.
- [29] A. Gu, K. Goel, and C. Re, "Efficiently modeling long sequences with structured state spaces," in *Proc. ICLR*, 2022.
- [30] K. Goel, A. Gu, C. Donahue, and C. Re, "It's raw! Audio generation with state-space models," in *Proc. ICML*, 2022.
- [31] K. Miyazaki, M. Murata, and T. Koriyama, "Structured state space decoder for speech recognition and synthesis," *arXiv preprint arXiv:2210.17098*, 2022.
- [32] E. Nguyen, K. Goel, A. Gu, G. Downs, P. Shah, T. Dao, S. Baccus, and C. Ré, "S4ND: Modeling images and videos as multidimensional signals with state spaces," in *Proc. NeuIPS*, 2022.
- [33] A. Tustin, "A method of analysing the behaviour of linear systems in terms of time series," *Journal of the Institution of Electrical Engineers-Part IIA: Automatic Regulators and Servo Mechanisms*, vol. 94, no. 1, pp. 130–142, 1947.
- [34] I. Markovskiy, "Structured low-rank approximation and its applications," *Automatica*, vol. 44, no. 4, pp. 891–909, 2008.
- [35] J. R. Bunch, "A note on the stable decomposition of skew-symmetric matrices," *Mathematics of Computation*, vol. 38, no. 158, pp. 475–479, 1982.
- [36] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. ICASSP*, 2020.
- [37] H.-S. Choi, J.-H. Kim, J. Huh, A. Kim, J.-W. Ha, and K. Lee, "Phase-aware speech enhancement with deep complex u-net," in *Proc. ICLR*, 2019.
- [38] Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, and L. Xie, "DCCRN: Deep Complex Convolution Recurrent Network for Phase-Aware Speech Enhancement," in *Proc. Interspeech*, 2020.
- [39] J. Qi, J. Du, S. M. Siniscalchi, X. Ma, and C.-H. Lee, "On mean absolute error for deep neural network based vector-to-vector regression," *IEEE Signal Processing Letters*, vol. 27, pp. 1485–1489, 2020.
- [40] A. Kessy, A. Lewin, and K. Strimmer, "Optimal whitening and decorrelation," *The American Statistician*, vol. 72, 2015.
- [41] C. Valentini-Botinhao, "Noisy speech database for training speech enhancement algorithms and tts models," 2017.
- [42] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *Proc. ICASSP*, 2001.
- [43] C. Macartney and T. Weyde, "Improved speech enhancement with the wave-u-net," *arXiv preprint arXiv:1811.11307*, 2018.
- [44] S. Lv, Y. Fu, M. Xing, J. Sun, L. Xie, J. Huang, Y. Wang, and T. Yu, "S-dccrn: Super wide band dccrn with learnable complex feature for speech enhancement," in *Proc. ICASSP*, 2022.