



MP-SENet: A Speech Enhancement Model with Parallel Denoising of Magnitude and Phase Spectra

Ye-Xin Lu, Yang Ai, Zhen-Hua Ling*

National Engineering Research Center of Speech and Language Information Processing,
University of Science and Technology of China, Hefei, P. R. China

`yxlu0102@mail.ustc.edu.cn`, `{yangai, zhling}@ustc.edu.cn`

Abstract

This paper proposes MP-SENet, a novel Speech Enhancement Network which directly denoises Magnitude and Phase spectra in parallel. The proposed MP-SENet adopts a codec architecture in which the encoder and decoder are bridged by convolution-augmented transformers. The encoder aims to encode time-frequency representations from the input noisy magnitude and phase spectra. The decoder is composed of parallel magnitude mask decoder and phase decoder, directly recovering clean magnitude spectra and clean-wrapped phase spectra by incorporating learnable sigmoid activation and parallel phase estimation architecture, respectively. Multi-level losses defined on magnitude spectra, phase spectra, short-time complex spectra, and time-domain waveforms are used to train the MP-SENet model jointly. Experimental results show that our proposed MP-SENet achieves a PESQ of 3.50 on the public VoiceBank+DEMAND dataset and outperforms existing advanced speech enhancement methods.

Index Terms: speech enhancement, encoder-decoder, parallel denoising, magnitude spectra, phase spectra

1. Introduction

In real-life scenarios, speech waveforms captured by devices are inevitably degraded by noises, which immensely impacts real applications such as hearing aids and telecommunications. To alleviate the impact of noises and improve speech perceptual quality, many deep-learning-based speech enhancement (SE) methods have been proposed to recover clean waveforms from the degraded ones utilizing neural networks. Existing SE methods can be roughly divided into two categories, i.e., time-domain SE methods and time-frequency (TF) domain SE methods. The time-domain SE methods [1–5] adopted neural networks to learn the mapping from noisy waveforms to clean ones. Unfortunately, this category of methods still suffered from quality bottlenecks and showed inefficiency due to the direct generation of high-resolution waveforms. In comparison, the TF-domain SE methods exhibited superior performance.

The TF-domain SE methods aim to predict clean frame-level TF-domain representations and then reconstruct the enhanced waveforms. Generally, the phase is not included in the commonly used representations because it is a great challenge to directly enhance the phase spectra, given its wrapping and nonstructural properties. However, recent studies demonstrated that the phase information plays an essential role in the speech perceptual quality of SE methods, especially in low signal-to-noise (SNR) circumstances [6]. In earlier studies, researchers

only enhanced the magnitude spectra and reconstructed the waveforms using inverse short-time Fourier transform (ISTFT) from the enhanced magnitude and noisy phase spectra [7–10]. The absence of phase spectrum enhancement inevitably led to the degradation of the enhanced speech quality. To overcome the above issues, several approaches focused on the enhancement of short-time complex spectra, which implicitly recovered both clean magnitude and phase spectra [11–13]. More recent studies also proposed to enhance the magnitude followed by complex spectrum refinement [14, 15], which can alleviate the unbounded estimation problem [16] existing in the methods that only enhanced complex spectra. However, the compensation effect [17] between the magnitude and phase still existed, which led to imprecise phase estimation. Although PHASEN [18] proposed a two-stream network and acquired the ability to handle detailed phase patterns and utilize harmonic patterns, it still optimized the phase within the complex spectrum level. These methods were still unable to precisely and explicitly predict the clean phase spectra, leaving room for improvement in the enhanced speech quality. To this end, it is crucial to implement explicit prediction and optimization on the phase spectra for TF-domain SE methods.

Therefore, we propose MP-SENet, a TF-domain monaural SE model with parallel magnitude and phase spectra denoising. The MP-SENet forms a codec architecture and we bridge the encoder and decoder using two-stage convolution-augmented transformers (TS-Conformers) borrowed from CMGAN [15] to capture both local and global information. The encoder encodes the input noisy magnitude and phase spectra to compressed TF-domain representations for subsequent decoding. The parallel magnitude mask decoder and phase decoder decode out the clean magnitude and phase spectra, respectively, and finally the enhanced waveforms are reconstructed by ISTFT. Specifically, the magnitude mask decoder predicts boundary masks with a learnable sigmoid activation and then multiplies the masks with the noisy magnitude spectra to obtain the clean magnitude spectra. Inspired by our previous work [19], the phase decoder incorporates the parallel phase estimation architecture to directly predict the clean phase spectra. Experimental results show that our proposed MP-SENet outperforms state-of-the-art (SOTA) SE methods and alleviates the compensation effect between the magnitude and phase spectra by achieving explicit predictions and optimizations of them. To the best of our knowledge, we are the first to accomplish the direct enhancement of phase spectra.

2. Methodology

The overview of the model structure and training criteria of the proposed MP-SENet are illustrated in Fig. 1. The MP-SENet adopts a codec architecture to denoise the noisy speech wave-

* Corresponding author. This work was partially funded by the Fundamental Research Funds for the Central Universities.

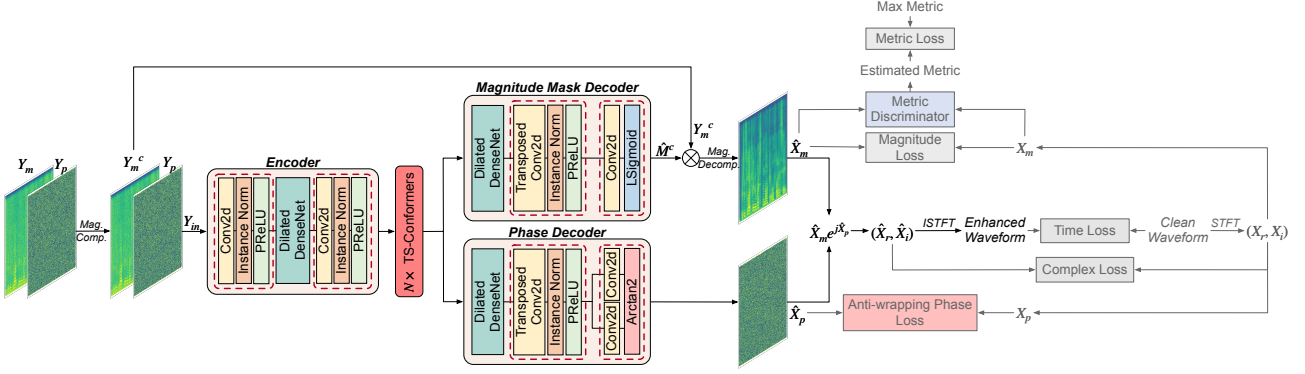


Figure 1: Overall structure and training principles of the proposed MP-SENet. The “Mag. Comp.” and “Mag. Decomp.” denote the magnitude compression and magnitude decompression operations, respectively. The dim parts only appear at the training stage.

form $\mathbf{y} \in \mathbb{R}^L$ and recover the clean speech waveform $\mathbf{x} \in \mathbb{R}^L$ in the TF domain, where L is the waveform length. Specifically, we first extract the magnitude spectrum $\mathbf{Y}_m \in \mathbb{R}^{T \times F}$ and the wrapped phase spectrum $\mathbf{Y}_p \in \mathbb{R}^{T \times F}$ from \mathbf{y} through STFT, where T and F denote the total number of frames and frequency bins, respectively. For more precise magnitude mask prediction, we apply the power-law compression on \mathbf{Y}_m and stack it with \mathbf{Y}_p to compose an input feature $\mathbf{Y}_{in} \in \mathbb{R}^{T \times F \times 2}$. Then the encoder encodes \mathbf{Y}_{in} into a compressed TF-domain representation, and subsequently, the TF-domain representation is processed by four TS-Conformers to capture time and frequency dependencies stage by stage. Finally, the parallel magnitude mask decoder and phase decoder predict the clean magnitude spectrum $\hat{\mathbf{X}}_m$ and clean phase spectrum $\hat{\mathbf{X}}_p$ from the TF-domain representation, respectively, and the final enhanced waveform $\hat{\mathbf{x}}$ is reconstructed by ISTFT. The details of the encoder, magnitude mask decoder, phase decoder, and training criteria are described as follows.

2.1. Model structure

2.1.1. Encoder

As illustrated in Fig. 1, the encoder encodes the input feature \mathbf{Y}_{in} into a TF-domain representation with a lower sampling rate and higher dimensions. It is a cascade of a convolutional block, a dilated DenseNet [20], and another convolutional block. Each convolutional block consists of a 2D convolutional layer, an instance normalization [21], and a parametric rectified linear unit (PReLU) activation [22]. The first convolutional block increases the feature dimension by increasing the number of channels in the convolutional layer, while the second block down-samples the feature by expanding the stride in the convolutional layer. The dilated DenseNet utilizes four convolutional layers with dilation sizes of 1, 2, 4, 8 to extend the receptive field along the time axis, and applies dense connections to all the convolutional layers to avoid the vanishing gradient problem.

2.1.2. Magnitude mask decoder

As illustrated in Fig. 1, the magnitude mask decoder predicts a magnitude mask from the TF-domain representation and multiplies it with the noisy magnitude spectrum to obtain the clean magnitude spectrum. However, the commonly used magnitude mask $\mathbf{M} = \mathbf{X}_m / \mathbf{Y}_m \in \mathbb{R}^{T \times F}$ is unbounded, where \mathbf{X}_m denotes the magnitude spectrum of the clean waveform \mathbf{x} . Earlier studies usually employ the sigmoid function to limit the value range of the predicted mask $\hat{\mathbf{M}}$ to $(0, 1)$. There is an

insurmountable gap between the predicted mask and the real mask whose range is out of $(0, 1)$. Recent methods [14, 15] compensate the gap by introducing another stream of complex spectrum prediction. Nevertheless, the compensation effect between magnitude and phase spectra still exists, leading to discontinuity of the spectrogram and damage to its harmonic structure. Accordingly, we first apply power-law compression on \mathbf{X}_m and \mathbf{Y}_m with a compression factor c to narrow the scope of the mask for easier prediction. Hence, the prediction target of the magnitude mask decoder is the compressed mask $\hat{\mathbf{M}}^c = (\mathbf{X}_m / \mathbf{Y}_m)^c$ and we set c to 0.3 in the experiments. Subsequently, to achieve precise prediction, we further adopt the learnable sigmoid (LSigmoid) function from [23] to predict the compressed magnitude mask:

$$\text{LSigmoid}(t) = \frac{\beta}{1 + e^{1-\alpha t}}, \quad (1)$$

where β is set to 2.0 and $\alpha \in \mathbb{R}^F$ is a trainable parameter, which allows the model to adaptively change the shape of the activation function in different frequency bands.

Specifically, with a TF-domain representation processed by TS-Conformers as input, the magnitude decoder feeds it to a dilated DenseNet, a deconvolutional block, and a magnitude mask estimation architecture to get the estimated compressed mask $\hat{\mathbf{M}}^c$. The deconvolutional block is used for upsampling and is composed of a 2D transposed convolutional layer, an instance normalization, and a PReLU activation. The magnitude mask estimation architecture first uses a 2D convolutional layer to reduce the output channel numbers of the deconvolutional block to 1 and then outputs $\hat{\mathbf{M}}^c$ with the activation of the LSigmoid function. Finally, the enhanced magnitude spectrum $\hat{\mathbf{X}}_m$ can be obtained by mask decoding as follows:

$$\hat{\mathbf{X}}_m = (\mathbf{Y}_m^c \odot \hat{\mathbf{M}}^c)^{\frac{1}{c}}, \quad (2)$$

where \odot denotes the element-wise multiplication.

2.1.3. Phase decoder

As illustrated in Fig. 1, the phase decoder directly predicts the clean phase spectrum from the TF-domain representation. In order to overcome the difficulties caused by the nonstructural and wrapping characteristics of the phase, we follow our previous work [19], cascade a parallel phase estimation architecture after a dilated DenseNet and a deconvolutional block in the phase decoder. The parallel phase estimation architecture first adopts

two parallel 2D convolutional layers to output the pseudo-real part component $\hat{\mathbf{X}}_p^{(r)}$ and pseudo-imaginary part component $\hat{\mathbf{X}}_p^{(i)}$, and then activates these two components to predict the clean wrapped phase spectrum $\hat{\mathbf{X}}_p$ using the two-argument arctangent (Arctan2) function, i.e.,

$$\hat{\mathbf{X}}_p = \arctan\left(\frac{\hat{\mathbf{X}}_p^{(i)}}{\hat{\mathbf{X}}_p^{(r)}}\right) - \frac{\pi}{2} \cdot \text{Sgn}^*(\hat{\mathbf{X}}_p^{(i)}) \cdot [\text{Sgn}^*(\hat{\mathbf{X}}_p^{(r)}) - 1], \quad (3)$$

where $\text{Sgn}^*(t)$ is a redefined function which equals to 1 when $t \geq 0$, and equals to -1 when $t < 0$.

2.2. Training criteria

We define multi-level loss functions for training the proposed MP-SENet. In keeping with [15], we use the time loss $\mathcal{L}_{\text{Time}}$, magnitude loss $\mathcal{L}_{\text{Mag.}}$, and complex loss $\mathcal{L}_{\text{Com.}}$, i.e.,

$$\mathcal{L}_{\text{Time}} = \mathbb{E}_{\mathbf{x}, \hat{\mathbf{x}}} [\|\mathbf{x} - \hat{\mathbf{x}}\|_1], \quad (4)$$

$$\mathcal{L}_{\text{Mag.}} = \mathbb{E}_{\mathbf{X}_m, \hat{\mathbf{X}}_m} [\|\mathbf{X}_m - \hat{\mathbf{X}}_m\|_2^2], \quad (5)$$

$$\mathcal{L}_{\text{Com.}} = \mathbb{E}_{\mathbf{X}_r, \hat{\mathbf{X}}_r} [\|\mathbf{X}_r - \hat{\mathbf{X}}_r\|_2^2] + \mathbb{E}_{\mathbf{X}_i, \hat{\mathbf{X}}_i} [\|\mathbf{X}_i - \hat{\mathbf{X}}_i\|_2^2], \quad (6)$$

where $(\mathbf{X}_r, \mathbf{X}_i), (\hat{\mathbf{X}}_r, \hat{\mathbf{X}}_i)$ denote the real and imaginary parts of the clean and enhanced complex spectra. Additionally, inspired by [24], a metric discriminator is also adopted for generative adversarial training, and the perceptual evaluation of speech quality (PESQ) is used as the target metric. We rescale the PESQ score to (0, 1), so the discriminator is trained to output the scaled PESQ score with pairs of clean and estimated magnitude spectra as inputs. The generator aims to output the magnitude spectrum whose metric can be discriminated to approach 1. The discriminator loss \mathcal{L}_D and the corresponding generated metric loss $\mathcal{L}_{\text{Metric}}$ are described as follows:

$$\mathcal{L}_D = \mathbb{E}_{\mathbf{X}_m} [\|D(\mathbf{X}_m, \mathbf{X}_m) - 1\|_2^2] + \mathbb{E}_{\mathbf{X}_m, \hat{\mathbf{X}}_m} [\|D(\mathbf{X}_m, \hat{\mathbf{X}}_m) - Q_{\text{PESQ}}\|_2^2], \quad (7)$$

$$\mathcal{L}_{\text{Metric}} = \mathbb{E}_{\mathbf{X}_m, \hat{\mathbf{X}}_m} [\|D(\mathbf{X}_m, \hat{\mathbf{X}}_m) - 1\|_2^2], \quad (8)$$

where D denotes the discriminator and $Q_{\text{PESQ}} \in [0, 1]$ denotes the scaled PESQ score.

In previous TF-domain SE works, phase spectra are optimized within the complex spectra or by directly calculating the absolute L^p distance between the natural clean and enhanced phase spectra \mathbf{X}_p and $\hat{\mathbf{X}}_p$. However, due to the phase wrapping property, the absolute distance between two phases may not be their actual distance, revealing the inappropriateness of conventional losses (e.g., absolute L^p distance) for phase optimization. Consistent with the anti-wrapping losses we proposed in [19], we respectively define the instantaneous phase loss, group delay loss, and instantaneous angular frequency loss between \mathbf{X}_p and $\hat{\mathbf{X}}_p$ as follows:

$$\mathcal{L}_{\text{IP}} = \mathbb{E}_{\mathbf{X}_p, \hat{\mathbf{X}}_p} [\|f_{\text{AW}}(\mathbf{X}_p - \hat{\mathbf{X}}_p)\|_1], \quad (9)$$

$$\mathcal{L}_{\text{GD}} = \mathbb{E}_{\Delta_{\text{DF}}(\mathbf{X}_p, \hat{\mathbf{X}}_p)} [\|f_{\text{AW}}(\Delta_{\text{DF}}(\mathbf{X}_p - \hat{\mathbf{X}}_p))\|_1], \quad (10)$$

$$\mathcal{L}_{\text{IAF}} = \mathbb{E}_{\Delta_{\text{DT}}(\mathbf{X}_p, \hat{\mathbf{X}}_p)} [\|f_{\text{AW}}(\Delta_{\text{DT}}(\mathbf{X}_p - \hat{\mathbf{X}}_p))\|_1], \quad (11)$$

where $f_{\text{AW}}(t) = |t - 2\pi \cdot \text{round}(\frac{t}{2\pi})|$, $t \in \mathbb{R}$ is an anti-wrapping function, which is used to avoid the error expansion issue caused by phase wrapping. Δ_{DF} and Δ_{DT} represent the differential operators along the frequency axis and time axis, respectively. The anti-wrapping phase loss is defined as:

$$\mathcal{L}_{\text{Pha.}} = \mathcal{L}_{\text{IP}} + \mathcal{L}_{\text{GD}} + \mathcal{L}_{\text{IAF}}. \quad (12)$$

The final generator loss \mathcal{L}_G is the linear combination of the losses mentioned above:

$$\mathcal{L}_G = \gamma_1 \mathcal{L}_{\text{Time}} + \gamma_2 \mathcal{L}_{\text{Mag.}} + \gamma_3 \mathcal{L}_{\text{Com.}} + \gamma_4 \mathcal{L}_{\text{Metric}} + \gamma_5 \mathcal{L}_{\text{Pha.}} \quad (13)$$

where $\gamma_1, \gamma_2, \dots, \gamma_5$ are hyperparameters and are set to 0.2, 0.9, 0.1, 0.05, and 0.3 with empirical trials, respectively. The training criteria of the MP-SENet is to minimize \mathcal{L}_G and \mathcal{L}_D simultaneously.

3. Experiments

3.1. Datasets and experimental setup

We used the publicly available VoiceBank+DEMAND dataset [25] for our experiments, which includes pairs of clean and noisy audio clips with a sampling rate of 48 kHz. The clean audio set is selected from the Voice Bank corpus [26], which consists of 11,572 audio clips from 28 speakers for training and 872 audio clips from 2 unseen speakers for testing. The clean audio clips are mixed with 10 types of noise (8 types from the DEMAND database [27] and 2 artificial types) at SNRs of 0dB, 5dB, 10dB, and 15 dB for the training set and 5 types of unseen noise from the DEMAND database at SNRs of 2.5 dB, 7.5dB, 12.5 dB, and 17.5 dB for the test set.

We resampled all the audio clips to 16 kHz in the experiments. To extract input features from raw waveforms using STFT, the FFT point number, Hanning window size, and hop size were set to 400, 400 (25 ms), and 100 (6.25 ms), respectively. All the models were trained using the AdamW optimizer [28] until 100 epochs. The learning rate was set initially to 0.0005 and halved every 30 epochs¹.

3.2. Comparison with advanced SE methods

Several representative time-domain SE methods including SEGAN [1], DEMUCS [3] and SE-Conformer [4], and TF-domain SE methods including MetricGAN [24], PHASEN [18], MetricGAN+ [23], and four SOTA methods (i.e., DPT-FSNet [12], TridentSE [13], DB-AIAT [14], and CMGAN [15]) were selected to compare with MP-SENet. Six commonly used objective evaluation metrics were chosen to evaluate the enhanced speech quality, including PESQ, segmental signal-to-noise ratio (SSNR), short-time objective intelligibility (STOI), and three composite measures (CSIG, CBAK, and COVL) which predict the mean opinion score (MOS) on signal distortion, background noise intrusiveness, and overall effect, respectively. For all the metrics, higher values indicate better performance.

The objective results are presented in Table. 1. Obviously, our proposed MP-SENet achieved a satisfactory PESQ score of 3.50 and outperformed all other SE methods on most metrics, reflecting our model's strong denoising ability. From Table. 1, we can see that time-domain methods were always inferior to TF-domain methods. Inside the TF-domain methods, only our proposed MP-SENet and PHASEN used both magnitude and phase spectra as input conditions. In spite of that, our proposed MP-SENet had 0.51, 0.52, 0.40, 0.60, and 0.46 improvements on the PESQ, CSIG, CBAK, COVL, and SSNR scores compared to PHASEN. The improvement was significant, proving that the magnitude mask decoder and phase decoder can recover more precise magnitude and phase spectra. When compared to the four TF-domain SOTA approaches, MP-SENet performed the best on the PESQ and three MOS-based metrics but slightly worse than DB-AIAT and CMGAN on the SSNR. For more ev-

¹Audio samples and source codes of the MP-SENet are available at <https://github.com/yxlu-0102/MP-SENet>.

Table 1: Comparison with other methods on VoiceBank+DEMAND dataset. “-” denotes the result is not provided in the original paper.

Method	Year	Input	#Param.	PESQ	CSIG	CBAK	COVL	SSNR	STOI
Noisy	-	-	-	1.97	3.35	2.44	2.63	1.68	0.91
SEGAN [1]	2017	Waveform	43.18M	2.16	3.48	2.94	2.80	7.73	0.92
DEMUCS [3]	2021	Waveform	33.53M	3.07	4.31	3.40	3.63	-	0.95
SE-Conformer [4]	2021	Waveform	-	3.13	4.45	3.55	3.82	-	0.95
MetricGAN [24]	2019	Magnitude	-	2.86	3.99	3.18	3.42	-	-
MetricGAN+ [23]	2021	Magnitude	-	3.15	4.14	3.16	3.64	-	-
DPT-FSNet [12]	2021	Complex	0.88M	3.33	4.58	3.72	4.00	-	0.96
TridentSE [13]	2023	Complex	3.03M	3.47	4.70	3.81	4.10	-	0.96
DB-AIAT [14]	2021	Magnitude+Complex	2.81M	3.31	4.61	3.75	3.96	10.79	-
CMGAN [15]	2022	Magnitude+Complex	1.83M	3.41	4.63	3.94	4.12	11.10	0.96
PHASEN [18]	2020	Magnitude+Phase	-	2.99	4.21	3.55	3.62	10.18	-
MP-SENet	2023	Magnitude+Phase	2.05M	3.50	4.73	3.95	4.22	10.64	0.96

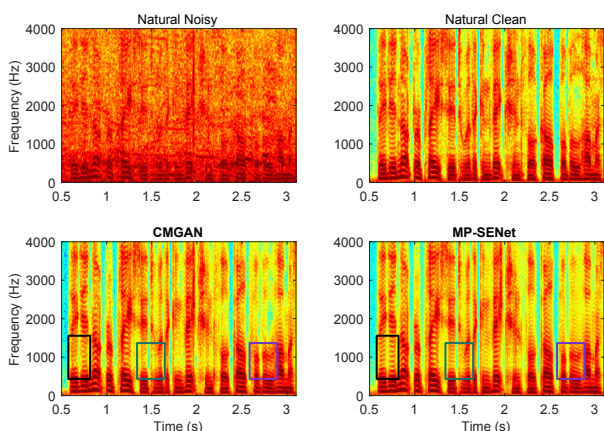


Figure 2: Spectrogram visualization of the natural noisy speech, natural clean speech, and speeches enhanced by the CMGAN and our proposed MP-SENet. For a more intuitive comparison, we only visualised the low-frequency areas.

idence, we visualized the spectrograms of speeches enhanced by MP-SENet and CMGAN as shown in Fig. 2. By comparing the contents of boxes with the same color, we can see that the MP-SENet alleviated the damage of low-frequency harmonic structures that occurred in CMGAN. This result indicated that predicting the phase directly instead of using complex spectrum refinement can alleviate the compensation effect between magnitude and phase and improve the enhancement precision of the phase spectra. In general, our proposed MP-SENet achieved a new SOTA performance among the objective metrics with a moderate model size of 2.05 M parameters.

3.3. Ablation study

To verify the role of each key component in the MP-SENet, we performed ablation studies and the results are presented in Table. 2. Obviously, the objective metrics collapsed when removing the magnitude compression operation (“w/o Mag. comp.”) and apparently degraded when replacing the LSigmoid activation with PReLU (“w/o LSigmoid”), which demonstrated the magnitude compression operation and LSigmoid were both effective for precise magnitude prediction. When we removed the phase decoder and combined the enhanced magnitude spectrum with the noisy phase spectrum to generate a waveform (“w/o

Table 2: Results of the ablation studies.

Method	PESQ	CSIG	CBAK	COVL	SSNR
MP-SENet	3.50	4.73	3.95	4.22	10.64
w/o Mag. comp.	2.97	4.06	3.61	3.56	9.39
w/o LSigmoid	3.40	4.67	3.87	4.14	10.10
w/o Pha. dec.	3.31	4.61	3.82	4.06	10.05
w/o Pha. loss	3.39	4.65	3.87	4.13	10.19
w/o Com. loss	3.44	4.72	3.89	4.19	10.21
w/o Metric disc.	3.39	4.69	3.89	4.15	10.48

Pha. dec.”), all the metrics greatly degraded which proved that phase prediction was indispensable. To further investigate the effects of phase optimization approaches, we conducted ablation studies on the phase loss (“w/o Pha. loss”) and complex loss (“w/o Com. loss”), which explicitly and implicitly optimized the phase, respectively. Results demonstrated that although both of them contributed to the overall performance, the explicit phase optimization was quite pivotal to SE tasks. Lastly, we ablated the metric discriminator (“w/o Metric disc.”) to assess the generator’s ability and improve the training efficiency. Surprisingly, our proposed MP-SENet without discriminator was still capable of high-quality SE, and its performance was even comparable to that of CMGAN.

4. Conclusions

In this paper, we proposed an SE model called MP-SENet, which denoised the magnitude and phase spectra in parallel. The overall structure of the MP-SENet was a Conformer-embedded codec architecture. The encoder encoded noisy magnitude and phase spectra, and the parallel magnitude mask decoder and phase decoder decoded out the clean magnitude and phase spectra, respectively. The major breakthrough of the MP-SENet lay in the direct enhancement of phase spectra. Experimental results show that our proposed MP-SENet achieved a SOTA performance on the VoiceBank+DEMAND dataset compared with other advanced SE methods. Moreover, ablation studies verified the effectiveness of each component and optimization method in the MP-SENet. Applying the parallel magnitude and phase enhancement method to other SE tasks (e.g., speech dereverberation, speech separation, and speech super-resolution) will be the focus of our future work.

5. References

- [1] S. Pascual, A. Bonafonte, and J. Serrà, “SEGAN: Speech enhancement generative adversarial network,” in *Proc. Interspeech*, 2017, pp. 3642–3646.
- [2] A. Pandey and D. Wang, “TCNN: Temporal convolutional neural network for real-time speech enhancement in the time domain,” in *Proc. ICASSP*, 2019, pp. 6875–6879.
- [3] A. Défossez, G. Synnaeve, and Y. Adi, “Real time speech enhancement in the waveform domain,” in *Proc. Interspeech*, 2020, pp. 3291–3295.
- [4] E. Kim and H. Seo, “SE-Conformer: Time-domain speech enhancement using conformer,” in *Proc. Interspeech*, 2021, pp. 2736–2740.
- [5] Z. Kong, W. Ping, A. Dantrey, and B. Catanzaro, “Speech denoising in the waveform domain with self-attention,” in *Proc. ICASSP*, 2022, pp. 7867–7871.
- [6] K. Paliwal, K. Wójcicki, and B. Shannon, “The importance of phase in speech enhancement,” *Speech Communication*, vol. 53, no. 4, pp. 465–494, 2011.
- [7] C. Valentini-Botinhao and J. Yamagishi, “Speech enhancement of noisy and reverberant speech for text-to-speech,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 8, pp. 1420–1433, 2018.
- [8] Y. Ai, J.-X. Zhang, L. Chen, and Z.-H. Ling, “DNN-based spectral enhancement for neural waveform generators with low-bit quantization,” in *Proc. ICASSP*, 2019, pp. 7025–7029.
- [9] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, “A regression approach to speech enhancement based on deep neural networks,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 1, pp. 7–19, 2014.
- [10] J. Kim, M. El-Khamy, and J. Lee, “T-GSA: Transformer with Gaussian-weighted self-attention for speech enhancement,” in *Proc. ICASSP*, 2020, pp. 6649–6653.
- [11] K. Tan and D. Wang, “Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 380–390, 2019.
- [12] F. Dang, H. Chen, and P. Zhang, “DPT-FSNet: Dual-path transformer based full-band and sub-band fusion network for speech enhancement,” in *Proc. ICASSP*, 2022, pp. 6857–6861.
- [13] D. Yin, Z. Zhao, C. Tang, Z. Xiong, and C. Luo, “TridentSE: Guiding speech enhancement with 32 global tokens,” *arXiv preprint arXiv:2210.12995*, 2022.
- [14] G. Yu, A. Li, C. Zheng, Y. Guo, Y. Wang, and H. Wang, “Dual-branch attention-in-attention transformer for single-channel speech enhancement,” in *Proc. ICASSP*, 2022, pp. 7847–7851.
- [15] R. Cao, S. Abdulatif, and B. Yang, “CMGAN: Conformer-based Metric GAN for Speech Enhancement,” in *Proc. Interspeech*, 2022, pp. 936–940.
- [16] D. S. Williamson, Y. Wang, and D. Wang, “Complex ratio masking for monaural speech separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 3, pp. 483–492, 2015.
- [17] Z.-Q. Wang, G. Wichern, and J. Le Roux, “On the compensation between magnitude and phase in speech separation,” *IEEE Signal Processing Letters*, vol. 28, pp. 2018–2022, 2021.
- [18] D. Yin, C. Luo, Z. Xiong, and W. Zeng, “PHASEN: A phase-and-harmonics-aware speech enhancement network,” in *Proc. AAAI*, vol. 34, no. 05, 2020, pp. 9458–9465.
- [19] Y. Ai and Z.-H. Ling, “Neural speech phase prediction based on parallel estimation architecture and anti-wrapping losses,” in *Proc. ICASSP*, 2023.
- [20] A. Pandey and D. Wang, “Densely connected neural network with dilated convolutions for real-time speech enhancement in the time domain,” in *Proc. ICASSP*, 2020, pp. 6629–6633.
- [21] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Instance normalization: The missing ingredient for fast stylization,” *arXiv preprint arXiv:1607.08022*, 2016.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proc. ICCV*, 2015, pp. 1026–1034.
- [23] 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, pp. 201–205.
- [24] 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.
- [25] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, “Investigating RNN-based speech enhancement methods for noise-robust text-to-speech,” in *Proc. SSW*, 2016, pp. 146–152.
- [26] 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. O-COCOSDA/CASLRE*, 2013, pp. 1–4.
- [27] J. Thiemann, N. Ito, and E. Vincent, “The diverse environments multi-channel acoustic noise database (DEMAND): A database of multichannel environmental noise recordings,” in *Proc. ICA*, vol. 19, no. 1, 2013, p. 035081.
- [28] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” *arXiv preprint arXiv:1711.05101*, 2017.