

IMPROVING DESIGN OF INPUT CONDITION INVARIANT SPEECH ENHANCEMENT

Wangyou Zhang^{1,2,*}, Jee-weon Jung^{2,*}, Yanmin Qian¹

¹Shanghai Jiao Tong University, China ²Carnegie Mellon University, USA

ABSTRACT

Building a single universal speech enhancement (SE) system that can handle arbitrary input is a demanded but underexplored research topic. Towards this ultimate goal, one direction is to build a single model that handles diverse audio duration, sampling frequencies, and microphone variations in noisy and reverberant scenarios, which we define here as “input condition invariant SE”. Such a model was recently proposed showing promising performance; however, its multi-channel performance degraded severely in real conditions. In this paper we propose novel architectures to improve the input condition invariant SE model so that performance in simulated conditions remains competitive while real condition degradation is much mitigated. For this purpose, we redesign the key components that comprise such a system. First, we identify that the channel-modeling module’s generalization to unseen scenarios can be sub-optimal and redesign this module. We further introduce a two-stage training strategy to enhance training efficiency. Second, we propose two novel dual-path time-frequency blocks, demonstrating superior performance with fewer parameters and computational costs compared to the existing method. All proposals combined, experiments on various public datasets validate the efficacy of the proposed model, with significantly improved performance on real conditions. Recipes with full model details will be released for reproducibility¹.

Index Terms— Universal speech enhancement, sampling-frequency-independent, microphone-number-invariant

1. INTRODUCTION

Speech enhancement (SE) aims to remove undesired signals in noisy and reverberant environments, thus enhancing desired speech quality [1]. While SE can be divided into many subtasks, this paper primarily focuses on denoising and dereverberation as they are most common in real world scenarios. The past decade has witnessed tremendous advancement of deep learning-based SE techniques with impressive performance in various scenarios. Spectral mapping [2], time-frequency (T-F) masking [3–5], time domain approaches with learnable convolution-based models [6–8], and sophisticated architectures [9–11] all brought significant advances in the SE literature. However, the majority of existing works mainly consider a single input condition, such as fixed single or multi-channel input, or input with a fixed sampling frequency. Thus, the designed SE models cannot be directly used in different scenarios, e.g., when switching

from one input condition to another. The mismatch between input conditions is also known to degrade the model performance.

Recently, an unconstrained speech enhancement and separation network (USES)² has been proposed to handle diverse input conditions with a single model [12]. With careful design and integration of several key components, the USES model becomes independent of 1) the sampling frequency; 2) the number (and geometry) of input microphone channels; and 3) the input signal length. The first independence is achieved leveraging the nature of short-time Fourier transform (STFT) and the adopted dual-path time-frequency model architecture [13]. The second is achieved by using the well-developed transform-average-concatenate (TAC) technique [14]. The last is achieved by introducing the memory tokens [12], which enable the model to perform segment-wise processing while maintaining the history information. In this paper, we term this type of SE model as “input condition invariant SE”, which is an important direction towards universal SE—our ultimate goal.

Although USES demonstrated promising performances on diverse datasets, we observe significant degradation in some real conditions, e.g., CHiME-4 multi-channel real recordings. We analyze that this can be partly due to the sub-optimal architecture. In particular, the adopted TAC modules perform channel modeling by concatenating the channel-averaged representation to each channel-specific representation. While this may be effective in many cases, it cannot well handle the microphone failure or channels with very different signal-to-noise ratios (SNRs). In addition, the TAC modules are used for processing both single- and multi-channel signals. Such a coupled processing makes it difficult to tune the single- and multi-channel performance separately. Hence, we argue that USES can be further developed to generalize better across various conditions.

In this paper, we aim to improve the model design of USES and enhance both performance and generalization in diverse conditions. To this end, we propose to decouple the single-channel/multi-channel processing and redesign key components including new time-frequency modeling and channel modeling modules. Based on the two proposed new architectures, USES2-Comp and USES2-Swin, we also adjust the technique combination accordingly to achieve the capability of handling diverse input conditions. Additionally, we propose a two-stage training strategy to improve the training efficiency of the model, which allows improving the single- and multi-channel processing capabilities separately. This also enables us to avoid data balancing between single- and multi-channel data. Following [12], we experiment with the combination of five commonly-used public corpora (VoiceBank+DEMAND [15], DNS1 [16], CHiME-4 [17], REVERB [18], and WHAMR! [19]) to demonstrate the model’s capability. Extensive results across datasets underscore the proposed single model’s substantial improvement of SE performance and generalization, outperforming the conventional USES model in diverse conditions.

²“Unconstrained” here means that the model is not constrained to be only used in a specific input condition.

*Equal contribution.

The authors would like to thank Shinji Watanabe for the helpful discussions during paper preparation. The experiments were done using the PI supercomputer at Shanghai Jiao Tong University and the Bridges2 system at PSC and Delta systems at NCSA through allocation CIS210014 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296.

¹<https://github.com/espnet/espnet>

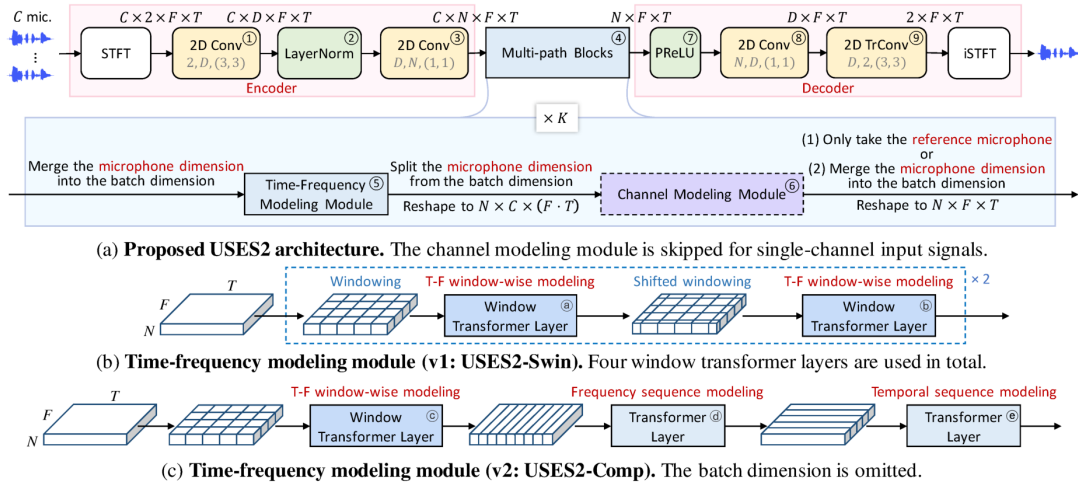


Fig. 1: The proposed input condition invariant SE model, USES2. Kernel size and feature maps of convolution layers are in gray.

2.1. Overview 2. PROPOSED METHODS

The proposed USES2 adopts the time-frequency dual-path modeling structure. This makes the model *inherently* sampling-frequency-independent (SFI), identical to [12]. As shown in Fig. 1 (a), the proposed model primarily comprises an STFT-based encoder ①–③, a few multi-path blocks ④ for spectral and spatial modeling, and an iSTFT-based decoder ⑦–⑨. Without loss of generality, the input signal has C microphone channels ($C \geq 1$). The C -channel signal is first converted to a complex-valued spectrum with shape $C \times 2 \times F \times T$, where F and T denote the frequency and temporal dimensions, respectively. Each complex T-F bin in the spectrum is then projected to an N -dimensional embedding by applying two 2D-convolution layers ① ③ and a layer normalization ②.

The stacked K multi-path blocks ④ perform the spectral and spatial modeling. In the first K_s multi-path blocks, the feature is processed by both time-frequency modeling module ⑤ and channel modeling module ⑥. In the remaining $K - K_s$ blocks, for the sake of efficiency, only the time-frequency modeling module ⑤ is used for spectral refining. The processed multi-channel feature is finally converted to single-channel by taking the reference channel and discarding the rest. Note that we propose to always skip the channel modeling modules if the input is a single-channel signal, different to the design in [12]. This new **decoupled design** allows separate optimization of single-channel and multi-channel processing capabilities, which cannot be achieved in [12]. The decoder with a parametric rectification linear unit (PReLU) ⑦, a point-wise convolution layer ⑧, and a 2D transposed convolution layer (TrConv, ⑨) maps the feature to the complex-valued spectrum, which is further transformed to the waveform via iSTFT.

2.2. Time-frequency modeling module ⑤

The original transformer-based frequency-path and time-path modeling in [12, 13] exhibits strong performance. Nevertheless, they separately model the temporal and frequency dimensions sequentially, whereas joint modeling has the potential to be further beneficial [20, 21]. Indeed, the short-time stationary nature of speech signals indicates that adjacent frames and frequency bins are highly correlated. Therefore, exploiting the information in a local T-F region should also benefit the spectral modeling. Inspired by the success of the Swin transformer in computer vision [22] and music classification [23], we incorporate the idea of transformer-based local window modeling into SE. Two novel architectures are proposed: (v1:

USES2-Swin) purely Swin transformer-based joint T-F modeling; (v2: USES2-Comp) complementary T-F modeling.

v1, Fig. 1 (b). USES2-Swin solely relies on the window-wise transformer layers for spectral modeling. Similarly to the Swin transformer design [22, 24], the windowing and shifted windowing operations are respectively applied before the odd ④ and even ⑥ transformer layers. In the windowing operation, the T-F representation is evenly partitioned into non-overlapping windows of size $W_F \times W_T$, where W_F and W_T are the numbers of frequency bins and frames in each window, respectively. Zero-paddings are added before windowing (and later removed) so that the numbers of frequency bins and frames are exactly divisible by W_F and W_T , respectively. The window transformer layer is applied to each window independently for local T-F modeling. Following [22], a standard transformer structure [25] with layer normalization is used, except that a learnable 2D relative positional bias is added to the attention matrix in each multi-head self-attention module. Thanks to the use of a fixed-duration STFT window, we have constant T-F resolution in the spectrum regardless of the sampling frequency (SF) [12]³. Thus, the learned relative positional bias can be directly used without finetuning for handling signals of different SFs. Subsequently, the shifted windowing operation is introduced to expand the receptive field of each window in deeper layers. It shifts the window partition by $\lfloor \frac{W_F}{2} \rfloor$ and $\lfloor \frac{W_T}{2} \rfloor$ along the frequency and time dimensions, respectively. By stacking multiple transformer layers with windowing and shifted windowing, the T-F modeling module can effectively exploit information from different T-F regions. (Since this new structure can naturally handle arbitrarily long signals thanks to the window-based processing, we remove the memory tokens proposed in [12] in this architecture.

v2, Fig. 1 (c). USES2-Comp regards the window-wise modeling as a complementary design to the existing frequency-path and time-path modeling. We insert a window transformer layer ④ at the beginning of each T-F modeling module to jointly capture the local T-F information. Then, following [12], a group of learnable memory tokens [mem] of shape $1 \times N \times 1 \times G$ is concatenated before the processed feature along the time dimension for long sequence modeling, where G is the group size. Two improved transformer layers [13, 26] are subsequently used for frequency-path modeling ④ and time-path modeling ⑥.

³The T-F resolution here means the time/frequency difference between adjacent time frames/frequency bins.

2.3. Channel modeling module ⑥

In addition to the improved T-F modeling module, we further improve the channel modeling module ⑥ in Fig. 1 (a) to boost the spatial processing capability. The vanilla TAC module has limited capability in handling scenarios where different microphone channels have significantly different signal-to-noise ratios (SNRs) as observed in [12]. We conjecture that the design choice of the TAC module incurs this limitation because it simply averages the representations from different channels. The channel-wise attention [27] also handles different microphone numbers and permutations, achieving impressive performance in matched conditions. However, this technique suffers from overfitting according to our preliminary experiments, violating our goal to build an input condition invariant SE model with good generalization. (c.f. rows 5–6 in Table 2)

We propose a novel channel modeling module by combining the best of both approaches, named transform-attention-concatenate (TA_nC). It maintains the channel independence and generalization while being capable of filtering noisy channels with attention. Given an input feature $\mathbf{X} \in \mathbb{R}^{N \times C \times (FT)}$, the channel modeling process of the proposed TA_nC module is formulated as follows:

$$\mathbf{Y} = \text{PReLU}(\text{FC}(\mathbf{X})) \in \mathbb{R}^{H \times C \times (FT)}, \quad (1)$$

$$\tilde{\mathbf{Y}} = \text{PReLU}(\text{FC}(\text{Attn}(\mathbf{Y}))) \in \mathbb{R}^{H \times C \times (FT)}, \quad (2)$$

$$\hat{\mathbf{X}} = \text{LN}(\text{PReLU}(\text{FC}([\mathbf{Y}, \tilde{\mathbf{Y}}]))) \in \mathbb{R}^{N \times C \times (FT)}, \quad (3)$$

where $[\mathbf{Y}, \tilde{\mathbf{Y}}] \in \mathbb{R}^{2H \times C \times (FT)}$ denotes concatenating two features along the embedding dimension. H is the projected embedding dimension. $\text{FC}(\cdot)$ and $\text{LN}(\cdot)$ respectively denote the linear projection and layer normalization, both operating along the embedding dimension. $\text{Attn}(\cdot)$ denote the channel-wise attention as defined below:

$$\mathbf{Q} = \text{Reshape}(\text{LN}(\text{ReLU}(\text{FC}(\mathbf{Y})))) \in \mathbb{R}^{1 \times C \times (HFT)}, \quad (4)$$

$$\mathbf{K} = \text{Reshape}(\text{LN}(\text{ReLU}(\text{FC}(\mathbf{Y})))) \in \mathbb{R}^{1 \times C \times (HFT)}, \quad (5)$$

$$\mathbf{V} = \text{LN}(\text{ReLU}(\text{FC}(\mathbf{Y}))) \in \mathbb{R}^{H \times C \times (FT)}, \quad (6)$$

$$\mathbf{A} = \text{LN}(\text{ReLU}(\text{FC}(\text{Softmax}(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{HT^2}})\mathbf{V}))), \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{H \times C \times (FT)}$ is the output of the Attn module. Note that the Attn module merges the time and frequency dimensions into the embedding dimension when calculating the attention map $\text{Softmax}(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{HT^2}}) \in \mathbb{R}^{1 \times C \times C}$ in Eq. (7), which makes it independent of the size of F and T .

2.4. Two-stage training strategy

We propose a two-stage training strategy to improve the training efficiency of the proposed SE model. Since we have decoupled the single- and multi-channel processing in Section 2.1, it is now possible to separate the optimization of these two capabilities with single- and multi-channel data, respectively. Specifically, in the first stage, we update all parameters except for the TA_nC modules ⑥ in Fig. 1 (a) on diverse single-channel data. This allows the model to focus on improving the single-channel SE performance. In the second stage, we train the TA_nC modules ⑥ while freezing all other parameters on diverse multi-channel data. This ensures to improve the multi-channel SE performance on top of the well-trained single-channel SE function with better generalization. The benefits of the proposed two-stage training strategy include: 1) it is more flexible as we can use single- and multi-channel data separately for training without caring about the data balance issue; 2) we can optimize multi-channel processing independently without affecting single-channel SE capabilities, which is preferable in some practical applications.

Table 1: Information of the datasets. “#Ch” denotes the number of microphone channels. “(Sim)” and “(Real)” denote the synthetic and recorded data. “(A)” and “(R)” represent anechoic and reverberant.

Dataset	Hours (train/dev/test)	SF	#Ch
VoiceBank+DEMAND [15]	8.8/0.6/0.6	48 kHz	1
DNS1 (v1) [16]	(A)90/(A)10/(R)0.42 & (A)0.42	16 kHz	1
CHiME-4 [17]	(Sim)14.7/(Sim)2.9/(Sim)2.3 & (Real)2.2	16 kHz	5
REVERB [18]	(Sim)15.5/(Sim)3.2/(Sim)4.8 & (Real)0.7	16 kHz	8
WHAMR! [19]	(R)58.0/(R)14.7/(R)9.0 (A)58.0/(A)14.7/(A)9.0	16 kHz	2

3. EXPERIMENTS

3.1. Datasets

We configure a global training set spanning ~245 hours to train our USES2 combining five well recognized datasets, identical to [12]. The combined dataset covers diverse input conditions as shown in Table 1, and more information can be found in [12]. Performances are individually reported for each dataset’s evaluation partition.

3.2. Configurations

In all experiments, the USES2-Swin / USES2-Comp model comprises $K = 3 / K = 4$ multi-path blocks in Fig. 1 (b) / (c), where the first $K_s = 2$ blocks include the channel modeling module ⑥ in Fig. 1 (a). There are totally 12 transformer layers in both models. STFT / iSTFT window and hop sizes are respectively set to 32 and 16 ms, regardless of the sampling frequency. Existing sub-modules’ hyperparameters, including embedding dimension, bottleneck dimension, and transformer layer architecture, are identical to [12]. By default, we take the first channel as the reference channel when processing multi-channel data. All experiments were conducted using the ESPnet-SE [28] toolkit, where the Adam optimizer is used for training. The learning rate increases linearly to $4e-4$ in the first 4000 steps and then decreases by half when the validation performance does not improve for two consecutive epochs. Four-second chunks are used during training for efficiency. We use a batch size of four. For multi-channel data training, we shuffle the channels of each sample and randomly select up to four channels. The scale-invariant multi-resolution L_1 loss in the frequency domain with a time-domain L_1 loss term is adopted [29]. We set the STFT window sizes of the multi-resolution L_1 loss to {256, 512, 768, 1024} and the time-domain loss weight to 0.5.

Metrics. We report performances using five metrics below: wide-band PESQ (PESQ-WB) [30], STOI [31], signal-to-distortion ratio (SDR) [32], DNSMOS (OVRL)⁴ [33], and word error rate (WER). The OpenAI Whisper model⁵ [?] is used for WER evaluation.

3.3. Results analysis

Architecture exploration with a single dataset. Table 2 first describes our incremental exploration with the proposed techniques on the CHiME-4 dataset. The objective here is to find a model that performs strongly in both simulated and real conditions. Rows 1 and 2 serve as baselines. Conventional USES [12] shows severely degraded performance for the real test condition with WER surging to 78.1%. We first explore the proposed decoupled channel processing, 2-stage training, and channel modeling module through rows 3 to 8. It is worth noting that the cascaded structure of channel-wise attention and TAC (Rows 5–6) cannot generalize well to real conditions despite its strong performance on simulation data, as mentioned in Section 2.3. Applying decoupled processing and TA_nC for

⁴https://github.com/microsoft/DNS-Challenge/blob/master/DNSMOS/DNSMOS/sig_bak_ovr.onnx

⁵<https://huggingface.co/openai/whisper-large-v2>

Table 2: Experiments on the proposed methods trained and evaluated using the CHiME-4 (5ch) dataset. Models are trained only on 8 kHz data, and tested on 16 kHz data. The values within parentheses denote single-channel performance while all others are multi-channel performance. The best and second best results are made bold and underlined. All models except for Nos. 1 and 2 are our novel explorations.

No.	Model	#Param	#MAC (G/s)		Test (CHiME-4 Simu)					Test (CHiME-4 Real)	
			1ch	2ch	PESQ-WB \uparrow	STOI ($\times 100$) \uparrow	SDR (dB) \uparrow	DNSMOS OVRL \uparrow	WER (%) \downarrow	DNSMOS OVRL \uparrow	WER (%) \downarrow
1	No processing	-	-	-	1.07	67.4	-0.7	1.30	28.3	1.46	6.7
2	USES (baseline) [12]	3.05 M	65.3	98.0	3.16	98.3	20.6	<u>3.22</u>	4.2	1.99 (2.94)	78.1 (11.0)
3	w/ decoupled proc.	3.05 M	60.8	98.0	2.62	97.0	18.2	3.08	5.2	2.32	27.8
4	No.3 + 2-stage training	3.05 M	60.8	98.0	2.23	94.9	15.6	2.99	6.8	2.41	23.8
5	No.3 + Att $\times 1$ + TAC $\times 2$	3.02 M	60.8	97.5	3.21	98.5	22.2	3.20	4.0	1.48	99.0
6	No.5 + 2-stage training	3.02 M	60.8	97.5	2.98	97.9	19.0	<u>3.22</u>	4.4	1.58	93.8
7	No.3 + TA $_n$ C $\times 3$	3.47 M	60.8	116.9	2.72	97.3	18.7	3.14	5.2	2.40	28.8
8	No.7 + 2-stage training	3.47 M	60.8	116.9	3.06	98.2	19.8	3.21	4.2	2.40	54.3
9	USES2-Swin (proposed)	2.92 M	37.7	75.5	2.94	98.1	20.6	3.21	4.2	2.84	22.1
10	No.9 + 2-stage training	2.92 M	37.7	75.5	2.98	98.2	<u>21.1</u>	<u>3.22</u>	<u>4.1</u>	2.80	24.9
11	USES2-Comp (proposed)	2.53 M	52.4	83.0	3.05	<u>98.3</u>	20.4	3.21	4.2	<u>2.89</u>	15.6
12	No.11 + 2-stage training	2.53 M	52.4	83.0	2.94	97.9	18.8	3.23	4.6	2.96	<u>12.1</u>




Fig. 2: Speech enhancement and recognition performance in diverse simulated conditions, averaged over the five corpora listed in Table 1. The models are trained only on 8 kHz data (via downsampling), and tested with the original sampling frequencies. “USES”, “excl”, and “U2-C” denote the SE model proposed in [12], corpus-exclusive SE, and the proposed USES2-Comp model, respectively. We follow the same training and evaluation configurations as in [12] to make the results comparable.

Table 3: Performance evaluation on various real recordings.

Test set	DNSMOS OVRL \uparrow			WER (%) \downarrow				
	noisy	USES	excl	U2-C	noisy	USES	excl	U2-C
CHiME-4 (Real, 5ch)	1.46	1.58	2.89	3.08	6.7	85.9	15.6	10.3
REVERB (Real, 8ch)	1.57	3.11	2.10	3.07	5.8	5.1	4.9	5.1

channel modeling in rows 7 and 8 gives a good balance of performance between simulated and real conditions. All metrics in simulated condition are comparable, while multi-channel performance in real condition improves significantly (WER 78.1% \rightarrow 28.8%). Rows 9 to 12 demonstrate USES2 performances with and without the proposed 2-stage training. Both architectures outperform the USES model, while the proposed USES2-Comp with 2-stage training demonstrates the best performance in the real condition in both DNSMOS and WER. Note that both proposed models reduce the amount of parameters and computational costs compared to USES. Since USES2-Comp shows overall better performance than USES2-Swin in Table 2, we adopt this architecture for the rest experiments.

Evaluation of USES2 in diverse conditions. Table 3 addresses extensive experimental results on various real datasets. We train a single USES2-Comp model (denoted as “U2-C”) with two-stage training on a combined dataset [12] as introduced in Section 3.1 to evaluate its universality in diverse conditions. We also train a corpus-exclusive SE model with the USES2-Comp architecture on each corpus alone (denoted as “excl”) to validate the efficacy of the “U2-C” model. All models are trained on downsampled 8 kHz data for the sake of efficiency, and then evaluated with the original sampling frequencies. It can be seen that *the multi-channel SE performance is significantly improved on CHiME-4 real data*, with $\sim 90\%$ relative improvement in both DNSMOS OVRL and WER. This is attributed

to the proposed TA $_n$ C module in Section 2.3 that overcomes the limitation of the vanilla TAC design, which can well handle microphone channels with unevenly distributed energy as in CHiME-4. Meanwhile, the SE and ASR performances on other real data are comparable, showing the robustness of our proposed USES2-Comp model in unseen scenarios.

Finally, we validate whether our proposed USES2-Comp model can preserve the strong performance on simulated conditions in addition to the generalization in real conditions. Fig. 2 shows the average performance over five different datasets as listed in Table 1. It can be seen that both USES and U2-C models show comparable performances that outperform the corpus-exclusive models (“excl”). Combining with the observation in Table 3, we believe that the proposed USES2-Comp model can be a drop-in replacement for the conventional USES model with reduced parameters and computational costs.

4. CONCLUSION

In this paper, we proposed two novel architectures, USES2-Comp and USES2-Swin, for improving the input condition invariant SE in terms of performance and generalization. Both USES2 models significantly improved the multi-channel processing capability in real conditions, while preserving the strong performance in simulation conditions. We further proposed a two-stage training strategy to improve the training efficiency, which also circumvents the data balancing issue between single- and multi-channel data. Extensive experiments on various public corpora demonstrated the effectiveness of the proposed methods. In future work, we plan to develop a universal SE model that covers diverse distortions in addition.

5. REFERENCES

- [1] P. C. Loizou, *Speech enhancement: theory and practice*. CRC press, 2013. 1
- [2] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, “A regression approach to speech enhancement based on deep neural networks,” *IEEE/ACM Trans. ASLP*, vol. 23, no. 1, pp. 7–19, 2014. 1
- [3] Y. Wang, A. Narayanan, and D. Wang, “On training targets for supervised speech separation,” *IEEE/ACM Trans. ASLP*, vol. 22, no. 12, pp. 1849–1858, 2014. 1
- [4] D. S. Williamson *et al.*, “Time-frequency masking in the complex domain for speech dereverberation and denoising,” *IEEE/ACM Trans. ASLP*, vol. 25, no. 7, pp. 1492–1501, 2017. 1
- [5] Z.-Q. Wang, P. Wang, and D. Wang, “Complex spectral mapping for single-and multi-channel speech enhancement and robust ASR,” *IEEE/ACM Trans. ASLP*, vol. 28, pp. 1778–1787, 2020. 1
- [6] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation,” *IEEE/ACM Trans. ASLP*, vol. 27, no. 8, pp. 1256–1266, 2019. 1
- [7] A. Pandey and D. Wang, “A new framework for CNN-based speech enhancement in the time domain,” *IEEE/ACM Trans. ASLP*, vol. 27, no. 7, pp. 1179–1188, 2019. 1
- [8] Y. Luo, Z. Chen, and T. Yoshioka, “Dual-path RNN: efficient long sequence modeling for time-domain single-channel speech separation,” in *Proc. IEEE ICASSP*, 2020, pp. 46–50. 1
- [9] S. Zhao, B. Ma, K. N. Watcharasupat, and W.-S. Gan, “FR-CRN: Boosting feature representation using frequency recurrence for monaural speech enhancement,” in *Proc. IEEE ICASSP*, 2022, pp. 9281–9285. 1
- [10] Z.-Q. Wang, S. Cornell, S. Choi, Y. Lee, B.-Y. Kim, and S. Watanabe, “TF-GridNet: Making time-frequency domain models great again for monaural speaker separation,” in *Proc. IEEE ICASSP*, 2023, pp. 1–5. 1
- [11] L. Liu, H. Guan, J. Ma, W. Dai, G. Wang, and S. Ding, “A mask free neural network for monaural speech enhancement,” in *Proc. Interspeech*, 2023, pp. 2468–2472. 1
- [12] W. Zhang, K. Saijo, Z.-Q. Wang, S. Watanabe, and Y. Qian, “Toward universal speech enhancement for diverse input conditions,” in *Proc. IEEE ASRU*, 2023. 1, 2, 3, 4
- [13] L. Yang, W. Liu, and W. Wang, “TFPSNet: Time-frequency domain path scanning network for speech separation,” in *Proc. IEEE ICASSP*, 2022, pp. 6842–6846. 1, 2
- [14] Y. Luo, Z. Chen, N. Mesgarani, and T. Yoshioka, “End-to-end microphone permutation and number invariant multi-channel speech separation,” in *Proc. IEEE ICASSP*, 2020, pp. 6394–6398. 1
- [15] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, “Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks,” in *Proc. Interspeech*, 2016, pp. 352–356. 1, 3
- [16] C. K. Reddy *et al.*, “The INTERSPEECH 2020 deep noise suppression challenge: Datasets, subjective testing framework, and challenge results,” in *Proc. Interspeech*, 2020, pp. 2492–2496. 1, 3
- [17] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, “An analysis of environment, microphone and data simulation mismatches in robust speech recognition,” *Computer Speech & Language*, vol. 46, pp. 535–557, 2017. 1, 3
- [18] K. Kinoshita, M. Delcroix, T. Yoshioka, T. Nakatani, A. Sehr, W. Kellermann, R. Maas, S. Gannot, and B. Raj, “The RE-VERB challenge: A common evaluation framework for dereverberation and recognition of reverberant speech,” in *Proc. IEEE WASPAA*, 2013, pp. 1–4. 1, 3
- [19] M. Maciejewski, G. Wichern, E. McQuinn, and J. Le Roux, “WHAMR!: Noisy and reverberant single-channel speech separation,” in *Proc. IEEE ICASSP*, 2019, pp. 696–700. 1, 3
- [20] R. Gu and Y. Zou, “Temporal-spatial neural filter: Direction informed end-to-end multi-channel target speech separation,” *arXiv preprint arXiv:2001.00391*, 2020. 2
- [21] J.-w. Jung, H.-S. Heo, H. Tak, H.-j. Shim, J. S. Chung, B.-J. Lee, H.-J. Yu, and N. Evans, “AASIST: Audio anti-spoofing using integrated spectro-temporal graph attention networks,” in *Proc. IEEE ICASSP*, 2022, pp. 6367–6371. 2
- [22] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *Proc. ICCV*, 2021, pp. 10012–10022. 2
- [23] H. Zhao, C. Zhang, B. Zhu, Z. Ma, and K. Zhang, “S3T: Self-supervised pre-training with Swin transformer for music classification,” in *Proc. IEEE ICASSP*, 2022, pp. 606–610. 2
- [24] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, “SwinIR: Image restoration using Swin transformer,” in *Proc. ICCV*, 2021, pp. 1833–1844. 2
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017. 2
- [26] J. Chen, Q. Mao, and D. Liu, “Dual-path transformer network: Direct context-aware modeling for end-to-end monaural speech separation,” in *Proc. Interspeech*, 2020, pp. 2642–2646. 2
- [27] A. Pandey, B. Xu, A. Kumar, J. Donley, P. Calamia, and D. Wang, “TPARN: Triple-path attentive recurrent network for time-domain multichannel speech enhancement,” in *Proc. IEEE ICASSP*, 2022, pp. 6497–6501. 3
- [28] C. Li *et al.*, “ESPnet-SE: End-to-end speech enhancement and separation toolkit designed for ASR integration,” in *Proc. IEEE SLT*, 2021, pp. 785–792. 3
- [29] Y.-J. Lu, S. Cornell, X. Chang, W. Zhang, C. Li, Z. Ni, Z.-Q. Wang, and S. Watanabe, “Towards low-distortion multi-channel speech enhancement: The ESPnet-SE submission to the L3DAS22 challenge,” in *Proc. IEEE ICASSP*, 2022, pp. 9201–9205. 3
- [30] 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*, vol. 2, 2001, pp. 749–752. 3
- [31] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, “An algorithm for intelligibility prediction of time–frequency weighted noisy speech,” *IEEE Trans. ASLP*, vol. 19, no. 7, pp. 2125–2136, 2011. 3
- [32] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE Trans. ASLP*, vol. 14, no. 4, pp. 1462–1469, 2006. 3
- [33] C. K. Reddy, V. Gopal, and R. Cutler, “DNSMOS P.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors,” in *Proc. IEEE ICASSP*, 2022, pp. 886–890. 3