

CLOSING THE GAP BETWEEN TIME-DOMAIN MULTI-CHANNEL SPEECH ENHANCEMENT ON REAL AND SIMULATION CONDITIONS

Wangyou Zhang¹, Jing Shi², Chenda Li¹, Shinji Watanabe^{3†}, Yanmin Qian^{1†}

¹ MoE Key Lab of Artificial Intelligence, AI Institute
X-LANCE Lab, Department of Computer Science and Engineering
Shanghai Jiao Tong University, Shanghai, China

{wyz-97, lichenda1996, yanminqian}@sjtu.edu.cn

² Institute of Automation, Chinese Academy of Sciences shijing2014@ia.ac.cn

³ Carnegie Mellon University, USA shinjiw@ieee.org

ABSTRACT

The deep learning based time-domain models, e.g. Conv-TasNet, have shown great potential in both single-channel and multi-channel speech enhancement. However, many experiments on the time-domain speech enhancement model are done in simulated conditions, and it is not well studied whether the good performance can generalize to real-world scenarios. In this paper, we aim to provide an insightful investigation of applying multi-channel Conv-TasNet based speech enhancement to both simulation and real data. Our preliminary experiments show a large performance gap between the two conditions in terms of the ASR performance. Several approaches are applied to close this gap, including the integration of multi-channel Conv-TasNet into the beamforming model with various strategies, and the joint training of speech enhancement and speech recognition models. Our experiments on the CHiME-4 corpus show that our proposed approaches can greatly reduce the speech recognition performance discrepancy between simulation and real data, while preserving the strong speech enhancement capability in the frontend.

Index Terms— multi-channel speech enhancement, time domain, beamforming, automatic speech recognition

1. INTRODUCTION

With the development of deep learning, speech enhancement (SE), as well as speech separation, has witnessed remarkable advances in both single-channel and multi-channel scenarios [1–4]. Since surprisingly good performance has been achieved in the simulated conditions, more and more researches have drawn their interests in more realistic environments, such as noisy and reverberant speech recorded in various real-world scenarios.

When multiple microphones are available, the capacity of deep learning based speech enhancement models can be further boosted by leveraging the additional spatial information between different channels. A straightforward way is to apply single-channel speech enhancement techniques to the multi-channel speech by extracting the spatial feature as an auxiliary input [5–7]. However, such approaches inevitably introduce artifacts to the enhanced signal, which can be harmful to the downstream automatic speech recognition (ASR) task [8], even though the artifacts are imperceptible to human listeners. Another widely adopted method is known as the neural beamformer [9, 10]. It usually consists of a mask estimation network for predicting time-frequency masks and a conventional beamformer module such as the minimum variance distortionless

response (MVDR) [11] beamformer. The neural beamformer is favored for its good compatibility with the downstream ASR task, as it explicitly constrains the enhanced output to be distortionless and thus enjoys better generalizability in realistic scenarios.

More recently, the time-domain audio separation network (TasNet) [1, 12] was proposed for speech separation, and was later extended for denoising [13]. Different from conventional frequency-domain approaches, TasNet directly operates on the input waveform and performs speech enhancement on the learned representation space. It shows very promising results on several benchmarks in both single-channel speech enhancement and separation [1, 13–15].

While the aforementioned time-domain approaches bring significant performance improvement to speech enhancement, the performance gap between real and simulation conditions is still widely observed [13, 16, 17]. In this paper, we aim to reduce the gap between time-domain multi-channel speech enhancement on real and simulation conditions, which has not been well studied yet. One interesting direction is the combination of TasNet and neural beamforming. TasNet has strong modeling capability, and MVDR beamforming has the benefit of enhancement without distortion. But how these two methods can benefit from each other in the multi-channel speech enhancement task is not well studied. Previous work [8] proposed the Beam-TasNet to estimate the beamformer filter on the output of a multi-channel Conv-TasNet (MC-Conv-TasNet) [8, 18] for speech separation, which demonstrates superior performance over the vanilla MC-Conv-TasNet and oracle MVDR beamformer. However, the experiments were conducted only on simulated mixture data, without any background noise. Therefore, the performance and robustness of this approach on realistic data are still unknown.

In this work, we show that both MC-Conv-TasNet and Beam-TasNet trained on simulated noisy data can suffer from severe performance degradation on the real data. To alleviate such degradation, we propose two training schemes to improve the performance and robustness of MC-Conv-TasNet and Beam-TasNet: (1) Exploring different integration approaches in the Beam-TasNet framework; (2) Joint training of MC-Conv-TasNet and ASR models. We evaluate different methods on the CHiME-4 [16] corpus, which consists of real and simulation data for both training and testing, allowing us to verify the performance gap in different conditions. For real data, since it is hard to measure the speech enhancement metrics directly due to the lack of reference signals, we instead evaluate the ASR performance. The experimental results show that the proposed methods can greatly improve the overall performance of both MC-Conv-TasNet and Beam-TasNet. More than 42% relative word error rate (WER) reduction is achieved on the evaluation set, while a comparable speech enhancement performance is preserved.

[†]Yanmin Qian and Shinji Watanabe are the corresponding authors.




Figure 1: Schematic diagram of the Beam-TasNet architecture for multi-channel speech enhancement. Gray lines denote different channel-rotated input signals, which are fed into the MC-Conv-TasNet* module separately to form the multi-channel enhanced signal $\hat{\mathbf{x}}^{(1)}$.

2. PROPOSED METHODS

2.1. Beam-TasNet

We first review the Beam-TasNet approach proposed in [8] and reformulate it in the context of speech enhancement. The Beam-TasNet system makes use of the MC-Conv-TasNet to estimate speech and noise covariance matrices based on its output signal, and then performs beamforming on the original multi-channel input.

To build this system, the MC-Conv-TasNet model is first trained on simulated multi-channel speech data. As shown in the left part of Figure 1, it consists of the multi-channel encoder¹, separator, and decoder. The multi-channel encoder aggregates multiple input channels into one hidden representation, which is then processed by the separator and decoder to generate a single-channel enhanced signal. In order to generate enhanced signals for all input channels, the MC-Conv-TasNet is trained in a channel-aware manner, which is hereafter referred to as MC-Conv-TasNet*. That is, the original C -channel training data is augmented by rotating the input channels anti-clockwise, so that each channel c can be placed as the first channel while preserving the original array geometry. Then, the MC-Conv-TasNet* model is trained to enhance each input signal with the first channel as the reference channel. In the inference phase, the multi-channel output can be obtained by rotating the input channels C times and feeding all channel-rotated signals into the MC-Conv-TasNet*. The above process can be formulated below:

$$\mathbf{y}^{(c)} = [y_c, y_{c+1}, \dots, y_C, y_1, y_2, \dots, y_{c-1}]^T, \quad (1)$$

$$\mathbf{w}_c = \text{MC-Conv-Encoder}(\mathbf{y}^{(c)}), \quad (2)$$

$$[\mathbf{m}_{c,X}, \mathbf{m}_{c,N}] = \text{Separator}(\mathbf{w}_c), \quad (3)$$

$$\hat{x}_c = \text{Conv-Decoder}(\mathbf{m}_{c,X} \otimes \mathbf{w}_c), \quad (4)$$

$$\hat{n}_c = \text{Conv-Decoder}(\mathbf{m}_{c,N} \otimes \mathbf{w}_c), \quad (5)$$

where $\mathbf{y}^{(c)}$ is the channel-rotated input signal with the c -th original channel placed at the first channel, $c = 1, 2, \dots, C$. MC-Conv-Encoder(\cdot) is the multi-channel encoder and \mathbf{w}_c is its output representation. $\mathbf{m}_{c,X}$ and $\mathbf{m}_{c,N}$ are the predicted speech and noise masks, respectively. \otimes denotes element-wise multiplication. \hat{x}_c and \hat{n}_c are the estimated speech and noise waveforms, respectively. For training the MC-Conv-TasNet and MC-Conv-TasNet*, we use the combination of time-domain signal-to-noise (SNR) losses on estimated speech and noise signals [13]:

$$\mathcal{L}_{\text{enh}} = -\text{SNR}(x_c, \hat{x}_c) - \text{SNR}(n_c, \hat{n}_c), \quad (6)$$

where \mathbf{x}_c and \mathbf{n}_c are the reference speech and noise signals, respectively. $\text{SNR}(x_c, \hat{x}_c) = 20 \log_{10} \frac{\|\mathbf{x}_c\|^2}{\|\mathbf{x}_c - \hat{x}_c\|^2}$, and $\|\cdot\|^2$ is the L_2 norm.

After training, the Beam-TasNet system is built upon the MC-Conv-TasNet* by calculating the speech and noise covariance matrices based on the enhanced multi-channel speech $\hat{\mathbf{x}}^{(1)} =$

$[\hat{x}_1, \hat{x}_2, \dots, \hat{x}_C]^T$ and then estimating the MVDR filter $\hat{\mathbf{h}}_f$:

$$\hat{\mathbf{h}}_f = \frac{(\Phi_{N,f})^{-1} \Phi_{X,f}}{\text{Trace}((\Phi_{N,f})^{-1} \Phi_{X,f})} \mathbf{u}, \quad (7)$$

where the subscript f is the frequency bin index. $\Phi_{X,f}$ and $\Phi_{N,f}$ denote the speech and noise covariance matrices, respectively, which will be discussed in detail in Section 2.2. $\text{Trace}(\cdot)$ is the matrix trace operator. \mathbf{u} is a one-hot vector denoting the reference channel q . Then the beamformed signal \hat{x}_q^{BF} can be derived as follows:

$$\hat{\mathbf{X}}_{f,q}^{\text{BF}} = \hat{\mathbf{h}}^H \mathbf{Y}_f^{(1)}, \quad (8)$$

$$\hat{x}_q^{\text{BF}} = \text{iSTFT}(\hat{\mathbf{X}}_q^{\text{BF}}), \quad (9)$$

where $\mathbf{Y}_f^{(1)}$ and $\hat{\mathbf{X}}_{f,q}^{\text{BF}}$ are the input noisy spectrum and beamformed spectrum, respectively. $(\cdot)^H$ denotes conjugate transpose. $\text{iSTFT}(\cdot)$ denotes the inverse short-time Fourier transform (STFT).

2.2. Integration approaches for Beam-TasNet

Following the introduction in Section 2.1 and [8], there are two main strategies to integrate the MC-Conv-TasNet into the Beam-TasNet architecture, i.e. sig-MVDR and mask-MVDR. The sig-MVDR uses the enhanced signals in Eq. (4) to calculate the covariance matrices in Eq. (7) directly:

$$\Phi_{X,f} = \frac{1}{T} \sum_t \hat{\mathbf{X}}_{t,f}^{(1)} (\hat{\mathbf{X}}_{t,f}^{(1)})^H, \quad (10)$$

$$\Phi_{N,f} = \frac{1}{T} \sum_t (\mathbf{Y}_{t,f}^{(1)} - \hat{\mathbf{X}}_{t,f}^{(1)}) (\mathbf{Y}_{t,f}^{(1)} - \hat{\mathbf{X}}_{t,f}^{(1)})^H, \quad (11)$$

where $\hat{\mathbf{X}}_{t,f}^{(1)}$ is the speech spectrum enhanced by MC-Conv-TasNet. The mask-MVDR estimates time-frequency masks from the enhanced signals for calculating the covariance matrices:

$$\Phi_{\alpha,f} = \frac{1}{T} \sum_t M_{\alpha,t,f} \mathbf{Y}_{t,f}^{(1)} (\mathbf{Y}_{t,f}^{(1)})^H, \quad (12)$$

where $\alpha \in \{X, N\}$. $M_{\alpha,t,f}$ is the estimated speech / noise mask.

In this section, we would like to put more emphasis on the mask-MVDR strategy, as the sig-MVDR based method may unexpectedly corrupt the appropriate spatial correlation in the multi-channel signal, while the mask-MVDR based method can mitigate such corruption, which will be shown later in our experiments. Since the speech mask is estimated from the enhanced signal, various types of masks can be investigated, including the well-known phase-sensitive mask (PSM) [19], and the voice activity detection (VAD) like 1-D mask. The PSM takes into account the phase information explicitly, which can be beneficial to the covariance matrix estimation. The VAD-like 1-D mask is shown to be robust against noise or interference signals [20–22], and it is calculated by averaging the power mask \mathbf{M}_X along the frequency dimension:

$$\mathbf{M}_{X,t}^{1-D} = \frac{1}{F} \sum_f M_{X,t,f}. \quad (13)$$

¹Also called parallel encoder in [18].




Figure 2: Joint training of SE and ASR models with truncated back-propagation through time. The blue and gray chunks denote the enhanced signal with and without the backward graph, respectively.

2.3. Joint training of MC-Conv-TasNet and ASR

Another direction is to jointly optimize the MC-Conv-TasNet frontend and the end-to-end ASR backend, which implicitly mitigates the mismatch between speech enhancement and ASR systems.

Since MC-Conv-TasNet directly operates on the raw waveform, the memory consumption can be very large when processing a full-length waveform, making it impractical for joint training. In order to jointly optimize MC-Conv-TasNet and end-to-end (E2E) ASR, we adopt the approximated truncated back-propagation through time (TBPTT) strategy used in [23]. That is, the backward graph is only retained for a randomly selected chunk instead of the full-length waveform, while the other part of the waveform is used only for the forward pass. Then, the full-length enhanced signal with the partially retained backward graph is fed into the ASR backend. This enables us to jointly train both frontend and backend with a flexibly adjustable memory cost, which is determined by the chunk size K .

When simulated and real data with transcripts are available for training, our proposed framework allows exploiting both data for optimizing the SE and ASR models. As illustrated in Figure 2, the final loss \mathcal{L} in the joint training is composed of two parts, i.e. the speech enhancement \mathcal{L}_{enh} and the ASR loss \mathcal{L}_{asr} . For real data, the final loss is equal to \mathcal{L}_{asr} , i.e. we train the SE and ASR models end-to-end without the need of signal-level references. For simulated data, the final loss is defined as $\mathcal{L} = \mathcal{L}_{\text{enh}} + \mathcal{L}_{\text{asr}}$.

Furthermore, the clean speech data can be additionally utilized to train only the ASR backend. The above multi-condition training strategy provides a flexible way for the MC-Conv-TasNet model to adapt to both simulated and real data, which is shown to greatly improve the ASR performance on real evaluation data in Section 3.2.

3. EXPERIMENTS

3.1. Experimental Setup

We conducted experiments on the 6-channel track of the CHiME-4 [16] corpus to evaluate our proposed methods. The CHiME-4 corpus consists of both real recordings and simulated speech data, so that we can easily evaluate the robustness and generalizability of our proposed approaches in unseen conditions. There are 42828 (9600), 1640 (1640), and 1320 (1320) simulated (real) samples for training, development and evaluation, respectively. The sample rate is 16 kHz for all speech data. We adopt the 5-th channel (CH5) as the reference channel for both training and evaluation. For evaluating the performance of frontend models on the real data, we adopt an E2E ASR model pretrained on the CHiME-4 dataset, which was also used in Section 4.1 in [24]. For the joint training of frontend and backend, we optionally include an additional dataset from the Wall Street Journal (WSJ) corpus [25] for training, which consists of 37416 clean speech samples. SpecAugment [26] is applied to the ASR input feature during the joint training. The chunk size K mentioned in Section 2.3 is set to 3 seconds. We use the Adam optimizer

for model training, and all models are trained until convergence.

All our models are built based on the ESPnet toolkit [24, 27]. The MC-Conv-TasNet model uses a Conv1D layer with 5 input channels and 256 output channels for the multi-channel encoder, with a kernel size of 20 and stride of 10. The separator consists of 4 stacked dilated convolutional blocks, each composed of 8 Conv1D blocks with 256 bottleneck channels and 512 hidden channels. The decoder is a transposed Conv1D layer with 256 input channels and 1 output channel, and the kernel size and stride are the same as the multi-channel encoder. The E2E ASR model is a joint connectionist temporal classification (CTC)/attention-based encoder-decoder [28] model, which consists of 12 and 6 transformer blocks with 2048 hidden units and four 64-dimensional attention heads for the encoder and decoder, respectively.

For speech enhancement performance, we adopt the short-time objective intelligibility (STOI) [29], perceptual evaluation of speech quality (PESQ) [30], and signal-to-distortion ratio (SDR) for evaluation. For ASR performance, the WER is evaluated.

3.2. Performance evaluation on simulation and real data

The performance of the baseline methods (No. 1~No. 5) and proposed methods (No. 6~No. 13) are listed in Table 1. Here we take the official result [16] as the ASR baseline, which uses a DNN-HMM acoustic model with language model rescoring. The speech enhancement baselines include the BeamformIt [32], neural beamformer (denoted as “BLSTM MVDR”) from [24], and the time-domain filter-and-sum network (FaSNet)² [31]. The speech recognition performance is evaluated using the same pretrained E2E ASR model on CHiME-4 for models from No. 2 to No. 10.

For the speech enhancement performance on simulated data, it can be observed that all MC-Conv-TasNet based models outperform the baselines, and the best performance is achieved by the MC-Conv-TasNet* model, which is trained using the channel-rotated data as described in Section 2.1. Compared to No. 2, the WERs of the MC-Conv-TasNet models on simulated data are also greatly reduced. However, the WERs on the real data become worse than the ASR baseline, which indicates the over-training of the MC-Conv-TasNet models in simulation conditions. In contrast, the frequency-domain BLSTM MVDR model shows better generalizability.³


After applying the trained MC-Conv-TasNet* model to the Beam-TasNet framework (No. 8~No. 10), we can observe significant WER reduction on both development and evaluation sets, especially on the real data. This is attributed to the distortionless constraint that is explicitly enforced in the design of MVDR beamforming. On the other hand, the speech enhancement performance of Beam-TasNet is worse than the MC-Conv-TasNet*, which could result from the fact that MVDR beamforming does not fully eliminate the noise and the residual noise level may be higher than that in MC-Conv-TasNet*. In addition, we can observe that the mask-MVDR based Beam-TasNet achieves much better speech recognition performance on real data than the sig-MVDR based one, while sacrificing some speech enhancement performance. This illustrates that the masking based integration approach can effectively mitigate the artifacts introduced in the TasNet output, making it more practical in realistic scenarios. And the VAD-like 1-D mask shows better ASR performance than PSM, which can be attributed to the averaging operation in Eq. (13) that further eliminates some inaccurate

²We use the open-source implementation at <https://github.com/y1uo42/TAC/blob/master/FaSNet.py#L176>.

³Another way to define the performance gap is the relative WER ratio between Real and Simu conditions, and we observed similar conclusions with this metric.

Table 1: Performance (PESQ/STOI/SDR [dB]/WER [%]) on the CHiME-4 6-channel track. The same ASR model pretrained on CHiME-4 is used for evaluating WER on models No. 2~No. 10. For PESQ, STOI, and SDR, larger is better. For WER (gray cells), smaller is better.

No.	Model	Dev (Simu)				Test (Simu)				Dev (Real)		Test (Real)	
		PESQ	STOI	SDR	WER	PESQ	STOI	SDR	WER	WER	WER		
1	Official baseline in [16]	-	-	-	6.8	-	-	-	10.9	5.8	11.5		
2	Noisy Input (CH5)	2.17	0.86	5.78	12.6	2.18	0.87	7.54	19.9	10.9	19.5		
3	BeamformIt in [24]	2.31	0.88	5.51	8.4	2.20	0.86	6.25	13.9	7.3	13.2		
4	BLSTM MVDR in [24]	2.68	0.95	13.40	5.3	2.68	0.95	14.10	8.0	5.9	9.8		
5	FaSNet [31]	2.64	0.93	10.56	8.3	2.43	0.89	9.73	18.4	10.3	22.5		
6	MC-Conv-TasNet	3.08	0.96	18.32	6.2	2.92	0.95	17.52	10.4	22.1	33.2		
7	MC-Conv-TasNet*	3.08	0.97	18.62	6.7	2.90	0.95	17.99	11.4	18.5	41.5		
8	→ Beam-TasNet (sig-MVDR)	2.57	0.95	15.31	5.3	2.58	0.95	16.17	7.3	11.8	25.7		
9	→ Beam-TasNet (mask-MVDR, PSM)	2.61	0.95	14.78	5.4	2.65	0.95	15.25	7.4	7.2	13.5		
10	→ Beam-TasNet (mask-MVDR, 1-D)	2.62	0.95	14.11	5.7	2.66	0.95	15.94	7.7	6.3	10.7		
11	Joint MC-Conv-TasNet + ASR	3.06	0.97	18.27	9.3	2.96	0.96	17.52	13.7	19.0	42.5		
12	+ real training data	3.07	0.96	18.19	8.3	2.93	0.95	17.39	12.1	9.1	17.0		
13	++ clean training data	3.05	0.96	18.10	6.5	2.93	0.95	17.24	10.0	7.3	13.5		



(a) Noisy speech (CH5) (b) Close-talk “reference” (CH0) (c) MC-Conv-TasNet (d) Beam-TasNet (sig-MVDR) (e) Beam-TasNet (mask-MVDR, 1-D) (f) Jointly trained MC-Conv-TasNet + ASR

Figure 3: Spectrograms of a real speech recording randomly selected from the CHiME-4 evaluation set.

estimations. Compared to the baseline ASR model, more than 42% relative WER reduction is achieved on all subsets.

Finally, the joint training⁴ in the last three lines leads to a comparable speech enhancement performance to the MC-Conv-TasNet models. Although the ASR performance on the simulated data is slightly worse than the MC-Conv-TasNet result (No. 6), this can be regarded as the effect of regularization from the real data during training. When only the simulated data is used for training (No. 11), the speech recognition on real data is severely degraded compared to the ASR baseline, which indicates the over-training in the simulation condition. After introducing the real data (No. 12) and additionally the clean data (No. 13) for training, we can observe pronounced WER reduction on the real data. This illustrates the advantage of joint training that various training data from different conditions can be well utilized to improve the overall performance of the entire system. Although the ASR performance does not outperform the BLSTM MVDR model, the performance gap is largely reduced, and much better enhancement performance is achieved.

To better illustrate the discrepancy between different speech enhancement methods, we further visualize the original spectrogram of a real sample and its enhanced versions from different models in Figure 3. Subfigures (a) and (b) are the noisy speech recorded by the distant microphone and close-talk microphone, respectively. Subfigures (c)–(f) show the corresponding enhanced signals by different models discussed above. We can observe that the MC-Conv-TasNet model severely corrupts the speech pattern in the spectrum, while the Beam-TasNet and jointly trained MC-Conv-TasNet models can restore the speech pattern and suppress the noise to some

⁴We jointly finetuned the pretrained MC-Conv-TasNet and E2E ASR models instead of training from scratch.

extent. This observation also coincides with the results in Table 1.⁵

4. CONCLUSION

In this paper, we investigate the performance of multi-channel Conv-TasNet based models for time-domain speech enhancement. A large performance gap is observed between simulation and real conditions. And several approaches are proposed to reduce this gap and improve the robustness of MC-Conv-TasNet based models, including the integration into the Beam-TasNet framework, and the joint training of MC-Conv-TasNet and ASR models. Experimental results on the CHiME-4 data show the difficulty of achieving good performance on real data, and that well-trained speech enhancement models on the simulated data do not necessarily remain advantageous when evaluated on real data. Our proposed approaches are shown to effectively mitigate the ASR performance gap, while still preserving a comparable speech enhancement capability. In the future work, we would like to incorporate more approaches to mitigating the mismatch between real and simulation conditions, such as better simulation strategies.

5. ACKNOWLEDGMENT

The authors would like to thank Dr. Tsubasa Ochiai for his helpful comments about the gap between Beam-TasNet enhancement results on real and simulation conditions. This work was supported by the China NSFC projects (No. 62071288 and No. U1736202) and Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102). Experiments have been carried out on the PI super-computer at Shanghai Jiao Tong University.

⁵We also conducted subjective evaluation and the result is available at https://x-lance.sjtu.edu.cn/members/wangyou-zhang/waspaa21_subjective_evaluation.

6. REFERENCES

- [1] 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.
- [2] Y. Liu and D. Wang, "Divide and conquer: A deep CASA approach to talker-independent monaural speaker separation," *IEEE/ACM Trans. ASLP*, vol. 27, no. 12, pp. 2092–2102, 2019.
- [3] T. Nakatani, R. Takahashi, T. Ochiai, K. Kinoshita, R. Ikeshita, M. Delcroix, and S. Araki, "DNN-supported mask-based convolutional beamforming for simultaneous denoising, dereverberation, and source separation," in *Proc. IEEE ICASSP*, 2020, pp. 6399–6403.
- [4] Z.-Q. Wang, H. Erdogan, S. Wisdom, K. Wilson, D. Raj, S. Watanabe, Z. Chen, and J. R. Hershey, "Sequential multi-frame neural beamforming for speech separation and enhancement," in *Proc. IEEE SLT*, 2021, pp. 905–911.
- [5] Z. Chen, X. Xiao, T. Yoshioka, H. Erdogan, J. Li, and Y. Gong, "Multi-channel overlapped speech recognition with location guided speech extraction network," in *Proc. IEEE SLT*, 2018, pp. 558–565.
- [6] R. Gu, L. Chen, S.-X. Zhang, J. Zheng, Y. Xu, M. Yu, D. Su, Y. Zou, and D. Yu, "Neural spatial filter: Target speaker speech separation assisted with directional information," in *Proc. ISCA Interspeech*, 2019, pp. 4290–4294.
- [7] R. Gu, S.-X. Zhang, L. Chen, Y. Xu, M. Yu, D. Su, Y. Zou, and D. Yu, "Enhancing end-to-end multi-channel speech separation via spatial feature learning," in *Proc. IEEE ICASSP*, 2020, pp. 7319–7323.
- [8] T. Ochiai, M. Delcroix, R. Ikeshita, K. Kinoshita, T. Nakatani, and S. Araki, "Beam-TasNet: Time-domain audio separation network meets frequency-domain beamformer," in *Proc. IEEE ICASSP*, 2020, pp. 6384–6388.
- [9] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *Proc. IEEE ICASSP*, 3 2016, pp. 196–200.
- [10] H. Erdogan, J. R. Hershey, S. Watanabe, M. I. Mandel, and J. Le Roux, "Improved MVDR beamforming using single-channel mask prediction networks," in *Proc. ISCA Interspeech*, 2016, pp. 1981–1985.
- [11] B. D. Van Veen and K. M. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE ASSP Magazine*, vol. 5, no. 2, pp. 4–24, 1988.
- [12] Y. Luo and N. Mesgarani, "TasNet: time-domain audio separation network for real-time, single-channel speech separation," in *Proc. IEEE ICASSP*, 2018, pp. 696–700.
- [13] K. Kinoshita, T. Ochiai, M. Delcroix, and T. Nakatani, "Improving noise robust automatic speech recognition with single-channel time-domain enhancement network," in *Proc. IEEE ICASSP*, 2020, pp. 7009–7013.
- [14] G. Wichern *et al.*, "WHAM!: extending speech separation to noisy environments," in *Proc. ISCA Interspeech*, 2019, pp. 1368–1372.
- [15] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "LibriMix: An open-source dataset for generalizable speech separation," *arXiv preprint:2005.11262*, 2020.
- [16] 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.
- [17] X. Liu and J. Pons, "On permutation invariant training for speech source separation," in *Proc. IEEE ICASSP*, 2021, pp. 6–10.
- [18] R. Gu, J. Wu, S.-X. Zhang, L. Chen, Y. Xu, M. Yu, D. Su, Y. Zou, and D. Yu, "End-to-end multi-channel speech separation," *arXiv preprint:1905.06286*, 2019.
- [19] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. IEEE ICASSP*, 2015, pp. 708–712.
- [20] Y.-H. Tu, J. Du, L. Sun, F. Ma, H.-K. Wang, J.-D. Chen, and C.-H. Lee, "An iterative mask estimation approach to deep learning based multi-channel speech recognition," *Speech Communication*, vol. 106, pp. 31–43, 2019.
- [21] A. S. Subramanian *et al.*, "An investigation of end-to-end multichannel speech recognition for reverberant and mismatch conditions," *arXiv preprint:1904.09049*, 2019.
- [22] W. Zhang, C. Boeddeker, S. Watanabe, T. Nakatani, M. Delcroix, K. Kinoshita, T. Ochiai, N. Kamo, R. Haeb-Umbach, and Y. Qian, "End-to-end dereverberation, beamforming, and speech recognition with improved numerical stability and advanced frontend," in *Proc. IEEE ICASSP*, 2021, pp. 6898–6902.
- [23] T. von Neumann, K. Kinoshita, L. Drude, C. Boeddeker, M. Delcroix, T. Nakatani, and R. Haeb-Umbach, "End-to-end training of time domain audio separation and recognition," in *Proc. IEEE ICASSP*, 2020, pp. 7004–7008.
- [24] 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.
- [25] LDC, *LDC Catalog: CSR-I (WSJ0) Complete*, University of Pennsylvania, 1993.
- [26] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Proc. ISCA Interspeech*, 2019, pp. 2613–2617.
- [27] S. Watanabe *et al.*, "ESPnet: End-to-end speech processing toolkit," in *Proc. ISCA Interspeech*, 2018, pp. 2207–2211.
- [28] S. Kim, T. Hori, and S. Watanabe, "Joint CTC-attention based end-to-end speech recognition using multi-task learning," in *Proc. IEEE ICASSP*, 2017, pp. 4835–4839.
- [29] 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.
- [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*, 2001, pp. 749–752.
- [31] Y. Luo, C. Han, N. Mesgarani, E. Ceolini, and S.-C. Liu, "FaSNet: Low-latency adaptive beamforming for multi-microphone audio processing," in *Proc. IEEE ASRU*, 2019, pp. 260–267.
- [32] X. Anguera, C. Wooters, J. Hern, and o, "Acoustic beamforming for speaker diarization of meetings," *IEEE Trans. ASLP*, vol. 15, no. 7, pp. 2011–2021, 2007.