



Generalized Spatio-Temporal RNN Beamformer for Target Speech Separation

Yong Xu¹, Zhuohuang Zhang^{2*}, Meng Yu¹, Shi-Xiong Zhang¹, Dong Yu¹

¹Tencent AI Lab, Bellevue, USA

²Indiana University, Bloomington, USA

{lucayongxu, raymondmyu, auszhang}@tencent.com, zhuozhan@iu.edu

Abstract

Although the conventional mask-based minimum variance distortionless response (MVDR) could reduce the non-linear distortion, the residual noise level of the MVDR separated speech is still high. In this paper, we propose a spatio-temporal recurrent neural network based beamformer (RNN-BF) for target speech separation. This new beamforming framework directly learns the beamforming weights from the estimated speech and noise spatial covariance matrices. Leveraging on the temporal modeling capability of RNNs, the RNN-BF could automatically accumulate the statistics of the speech and noise covariance matrices to learn the frame-level beamforming weights in a recursive way. An RNN-based generalized eigenvalue (RNN-GEV) beamformer and a more generalized RNN beamformer (GRNN-BF) are proposed. We further improve the RNN-GEV and the GRNN-BF by using layer normalization to replace the commonly used mask normalization on the covariance matrices. The proposed GRNN-BF obtains better performance against prior arts in terms of speech quality (PESQ), speech-to-noise ratio (SNR) and word error rate (WER).

Index Terms: MVDR, Spatio-temporal RNN beamformer, ADL-MVDR, GEV, GRNN-BF, speech separation

1. Introduction

The mask-based MVDR [1, 2, 3, 4, 5, 6] could obtain distortionless speech while existing purely “black box” neural network (NN) based speech separation methods [7, 8, 9, 10, 11] always suffer from the non-linear distortion problem [6]. However, the residual noise level of mask-based MVDR method is still high [12, 6]. Most of mask-based beamformers are optimized in the chunk-level [1, 3, 5, 6]. The calculated beamforming weights are hence chunk-level which is not optimal for each frame. Furthermore, the matrix inversion in the traditional beamformer (e.g., MVDR) comes with the numerical instability problem [13, 14, 15, 16], which is caused by the singularity in the matrix inversion [13]. Although this issue could be alleviated by using some techniques, e.g., diagonal loading [17, 14, 13], it is not fully solved. This problem could be worse in the end-to-end joint training system [13, 6]. Time-varying mask-based beamformers were investigated in [18, 19], however they also have the numerical instability problem [13, 14].

The recurrent neural network (RNN) was once proved to have the ability to solve the matrix inversion [20, 21] and the eigenvalue decomposition problems [22, 23], which are two main matrix operations in most of the beamformers’ solutions, e.g., MVDR [3, 5, 24] and Generalized Eigenvalue (GEV) beamformer [25, 26]. Although different types of beamformers [24, 3, 25, 27] have different optimization and constraint conditions, most of their solutions are derived from the estimated speech and noise covariance matrices. These prior studies inspired us to use RNNs to directly learn the beamforming

*work done during internship at Tencent

weights from the estimated speech and noise covariance matrices. Hence, we recently proposed an all-deep-learning MVDR (ADL-MVDR) method [28] which was superior to the traditional MVDR beamformer [3, 6]. In the ADL-MVDR [28], the matrix inversion and principal component analysis (PCA) operations of traditional MVDR are replaced by two RNNs with the estimated speech and noise covariance matrices as the input. In this work, we propose more advanced and generalized RNN-based beamformers (RNN-BFs). Note that there were also several other learning based beamforming methods [29, 30] which yielded worse performance than the traditional mask-based MVDR [3] approach due to the lack of explicitly using the speech and noise covariance matrices information [29].

In this work, three contributions are made to further improve the ADL-MVDR [28] beamformer. First, we propose a RNN-based GEV (RNN-GEV) beamformer, where it achieves slightly better performance than the ADL-MVDR [28]. It indicates that the RNNs could also be incorporated into other traditional beamforming algorithms. Second, a generalized RNN beamformer (GRNN-BF) is proposed and it is superior to the RNN-GEV and the ADL-MVDR [28]. The GRNN-BF directly learns the frame-level beamforming weights from covariance matrices without following conventional beamformers’ solutions. It suggests that the GRNN-BF could learn a better beamforming solution by automatically accumulating the covariance matrices across history frames. Finally, the layer normalization [31] is proposed to replace the commonly used mask normalization [32, 29, 6, 13] on the covariance matrices. The layer normalization is more flexible than the mask normalization and it can achieve better performance. These improvements make our proposed GRNN-BF perform the best in terms of PESQ, SNR and word error rate (WER) comparing to the traditional MVDR beamformers [6] and the ADL-MVDR beamformer [28].

The rest of this paper is organized as follows. In section 2, traditional mask-based beamformers are described. Section 3 presents the proposed generalized RNN-based beamformers (GRNN-BFs). The experimental setups and results are provided in Section 4 and 5, respectively. Section 6 concludes this paper.

2. Traditional mask-based beamformers

This section describes the traditional mask-based beamformers. Given the M -channel speech mixture $\mathbf{y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_M]$, the corresponding M -channel target speaker’s speech and the noise (the sum of interfering speakers’ speech and background noise) waveforms are denoted as \mathbf{s} and \mathbf{n} , respectively. After applying short-time Fourier transform (STFT), we have $\mathbf{Y}, \mathbf{S}, \mathbf{N}$ in the time-frequency (T-F) domain,

$$\mathbf{Y}(t, f) = \mathbf{S}(t, f) + \mathbf{N}(t, f) \quad (1)$$

where (t, f) indicates the time and frequency indices of the T-F domain variables. In conventional mask-based beamforming [1, 3, 5, 6], a neural network is used to predict the real-valued




Figure 1: The system framework includes the dilated Conv-1D blocks based complex-valued ratio filter (cRF) estimator and the proposed spatio-temporal RNN beamformer (RNN-BF). The cRF estimator is actually a Conv-TasNet variant [9] with a fixed STFT encoder [33]). \odot indicates the complex-domain multiplication to estimate the multi-channel speech $\hat{\mathbf{S}}$ and noise $\hat{\mathbf{N}}$ through cRF (as shown in Eq. (5)). \otimes is the matrix multiplication of beamforming (see Eq. (13)). The speech covariance matrix Φ_S , noise covariance matrix Φ_N and beamforming weights \mathbf{w} are complex-valued variables, their real parts and imaginary parts are reshaped and concatenated together. The time-domain scale-invariant SNR (Si-SNR) loss [9] is applied for end-to-end training.

speech mask RM_S and the real-valued noise mask RM_N . Then the speech covariance matrix Φ_S is calculated with the predicted speech mask RM_S ,

$$\Phi_S(f) = \frac{\sum_{t=1}^T \text{RM}_S^2(t, f) \mathbf{Y}(t, f) \mathbf{Y}^H(t, f)}{\sum_{t=1}^T \text{RM}_S^2(t, f)} \quad (2)$$

Where T stands for the total number of frames in a chunk. \mathbf{H} is the Hermitian transpose. The noise covariance matrix Φ_N could be calculated in the same way with the noise mask RM_N . The MVDR solution [3] can be derived as,

$$\mathbf{w}_{\text{MVDR}}(f) = \frac{\Phi_N^{-1}(f) \mathbf{v}(f)}{\mathbf{v}^H(f) \Phi_N^{-1}(f) \mathbf{v}(f)}, \quad \mathbf{w}_{\text{MVDR}}(f) \in \mathbb{C}^M \quad (3)$$

where $\mathbf{v}(f)$ represents the steering vector at f -th frequency bin. $\mathbf{v}(f)$ could be derived by applying PCA on $\Phi_S(f)$, namely $\mathbf{v}(f) = \mathcal{P}\{\Phi_S(f)\}$. Another type of commonly used beamformer is generalized eigenvalue (GEV) [1], where its optimal solution is the generalized principle component [1, 26],

$$\mathbf{w}_{\text{GEV}}(f) = \mathcal{P}\{\Phi_N^{-1}(f) \Phi_S(f)\}, \quad \mathbf{w}_{\text{GEV}}(f) \in \mathbb{C}^M \quad (4)$$

However, the beamforming weights \mathbf{w} above are usually chunk-level [1, 3, 5, 6], which is not optimal for each frame. Furthermore, the matrix inversion involved in Eq. (3) and Eq. (4) has the numerical instability problem [13, 14, 15, 16]. Note that we have already applied the diagonal loading technique [17, 14, 13] to alleviate this problem in our MVDR baselines.

On the other hand, although the MVDR and GEV are two different beamformers, their solutions are both derived from the speech and noise covariance matrices, namely Φ_S and Φ_N . This is also our motivation to use RNNs to directly learn the beamforming weights from Φ_S and Φ_N .

3. Proposed generalized RNN beamformer

We aim to extract the target speaker's speech from the multi-channel multi-talker overlapped mixture. As shown in Fig. 1, the whole system consists of a complex-valued ratio filter (cRF) estimator and the proposed spatio-temporal RNN beamformer. The predicted cRFs are used to calculate the covariance matrices. Then the proposed GRNN-BF learns the beamforming weights from the covariance matrices. The details of the input features and the cRF estimator will be described in Sec. 4. Our proposed GRNN-BF will be first illustrated here.

Recently we proposed the ADL-MVDR method [28] which is superior to the traditional mask-based MVDR beamformers

[3, 5, 6]. The ADL-MVDR uses two RNNs to replace the matrix inversion and PCA in MVDR solution (defined in Eq. (3)). Here we explore to use RNNs to implement the GEV beamformer (defined in Eq. (4)) and another more generalized RNN beamformer. Layer normalization [31] is also proposed to replace the commonly used mask normalization [32, 29, 13], which is applied on the covariance matrices.

3.1. Layer normalization on covariance matrix

Before we use RNNs to learn the beamforming weights, the speech and noise covariance matrices should be first estimated. Real-valued masks [32], complex-valued ratio mask (cRM) [34, 6] or complex-valued ratio filter (cRF) [35, 28] could be used to estimate the speech and noise. In our previous ADL-MVDR work [28], the cRF [35] is demonstrated to be better than the cRM [34]. The cRF [35] is $K \times K$ size cRM [34] by using nearby $K \times K$ T-F bins around (t, f) . With the speech $\text{cRF}_S(t, f)$, the multi-channel target speech is estimated as,

$$\hat{\mathbf{S}}(t, f) = \sum_{\tau_1=-K}^{\tau_1=K} \sum_{\tau_2=-K}^{\tau_2=K} \text{cRF}_S(t, f, \tau_1, \tau_2) * \mathbf{Y}(t + \tau_1, f + \tau_2) \quad (5)$$

Then the frame-wise speech covariance matrix is calculated as,

$$\Phi_S(t, f) = \frac{\hat{\mathbf{S}}(t, f) \hat{\mathbf{S}}^H(t, f)}{\sum_{t=1}^T \text{cRM}_S^H(t, f) \text{cRM}_S(t, f)} \quad (6)$$

Where $\text{cRM}_S(t, f)$ is the center unit of the speech $\text{cRF}_S(t, f)$. Given the noise $\text{cRF}_N(t, f)$, the noise $\hat{\mathbf{N}}(t, f)$ and the frame-wise noise covariance matrix $\Phi_N(t, f)$ could be estimated in the same way. Different from Eq. (2) where the covariance matrix are averaged over a chunk of frames in the traditional mask-based MVDR, the covariance matrices here are frame-wise. This is because the covariance matrices are later fed into RNNs where the unidirectional RNN could automatically accumulate the statistics of covariance matrices across history frames in a recursive way. Note that the denominator in Eq. (6) is the commonly used mask normalization [32, 29, 6, 13] to normalize the covariance matrix and equalize the training loss and gradients of different chunks in a mini-batch. In this work, we propose to use the layer normalization [31] to normalize the covariance matrices to achieve better performance.

$$\Phi_S(t, f) = \text{LayerNorm}(\hat{\mathbf{S}}(t, f) \hat{\mathbf{S}}^H(t, f)) \quad (7)$$

Where the layer normalization [31] applies per-element scale and bias with learnable affine transform, which is more flexible than the mask normalization. Another layer normalization is also adopted for $\Phi_N(t, f)$.

3.2. Spatio-temporal RNN GEV beamformer

Similar to the ADL-MVDR [28], here the proposed spatio-temporal RNN GEV beamformer (RNN-GEV) also takes the estimated target speech covariance matrix $\Phi_S(t, f)$ and the noise covariance matrix $\Phi_N(t, f)$ as the input to predict the frame-wise beamforming weights. Following the solution of the traditional GEV beamformer defined in Eq. (4), we reformulate its form in the RNN-based beamforming framework as,

$$\hat{\Phi}_N^{-1}(t, f) = \text{RNN}(\Phi_N(t, f)) \quad (8)$$

$$\hat{\Phi}_S(t, f) = \text{RNN}(\Phi_S(t, f)) \quad (9)$$

$$\mathbf{w}_{\text{RNN-GEV}}(t, f) = \text{DNN}(\hat{\Phi}_N^{-1}(t, f)\hat{\Phi}_S(t, f)) \quad (10)$$

$$\hat{\mathbf{S}}(t, f) = (\mathbf{w}_{\text{RNN-GEV}}(t, f))^H \mathbf{Y}(t, f) \quad (11)$$

where $\mathbf{w}_{\text{RNN-GEV}}(t, f) \in \mathbb{C}^M$. $\hat{\Phi}_S(t, f)$ is the accumulated speech covariance matrix from the history frames by leveraging on the temporal modeling capability of RNNs. $\hat{\Phi}_N^{-1}(t, f)$ is assumed to be the matrix inversion of $\Phi_N(t, f)$. Instead of using the actual generalized PCA (as in Eq. (4)), a deep neural network (DNN) is utilized to calculate the beamforming weights for RNN-GEV. Hinton et al [36] shows that the DNN has the ability to conduct the non-linear generalized PCA.

3.3. Generalized spatio-temporal RNN beamformer

Finally, we propose a more generalized spatio-temporal RNN beamformer (GRNN-BF) without following any traditional beamformers' solutions. This is motivated by that, different beamformers (e.g., MVDR and GEV) have different solutions but almost all solutions are derived from the speech and noise covariance matrices. The neural networks could be able to learn a better solution from the speech and noise covariance matrices. The RNN-GEV and the ADL-MVDR [28] both have two RNNs to deal with the target speech covariance matrix $\Phi_S(t, f)$ and the noise covariance matrix $\Phi_N(t, f)$, respectively. But GRNN-BF here uses only one unified RNN-DNN model to predict the frame-level beamforming weights directly.

$$\mathbf{w}_{\text{GRNN-BF}}(t, f) = \text{RNN-DNN}([\Phi_N(t, f), \Phi_S(t, f)]) \quad (12)$$

$$\hat{\mathbf{S}}(t, f) = (\mathbf{w}_{\text{GRNN-BF}}(t, f))^H \mathbf{Y}(t, f) \quad (13)$$

where $\mathbf{w}_{\text{GRNN-BF}}(t, f) \in \mathbb{C}^M$. The input for the RNN-DNN is the concatenated tensor of $\Phi_N(t, f)$ and $\Phi_S(t, f)$. All of the covariance matrices and beamforming weights are complex-valued, and we concatenate the real and imaginary parts of any complex-valued tensors in the whole work.

4. Dataset and experimental setup

Dataset: The methods are evaluated on the mandarin audio-visual corpus [37, 33], which is collected from YouTube [38]. The dataset has 205500 clean speech segments (about 200 hours) over 1500 speakers. The audio sampling rate is 16 kHz. 512-point of STFT is used to extract features along 32ms Hann window with 50% overlap. There are one to three overlapped speaking speakers in the simulated 15-channel mixture signal. The signal-to-interference ratio (SIR) is ranging from -6 to 6 dB. Noise with 18-30 dB SNR is added to all the 15-channel mixtures [37]. We use a 15-element non-uniform linear array. Based on the image-source simulation method [39], the simulated dataset contains 190000, 15000 and 500 multi-channel mixtures for training, validation and testing. The virtual acoustic room size is ranging from 4m-4m-2.5m to 10m-8m-6m. The reverberation time T60 is sampled in a range of 0.05s to 0.7s.

cRF estimator: we use the complex-valued ratio filter (cRF) [28, 35] to estimate the covariance matrices. As shown in Fig. 1, the input to the cRF estimator includes a 15-channel mixture audio and a target direction of arrival (DOA) (θ). From the multi-channel audio, log-power spectra (LPS) and interaural phase difference (IPD) [37] features are extracted. For the simulated data, the ground-truth target DOA is known. For the real-world scenario, we have the hardware where the 180° wide-angle camera and the 15-linear microphone array are aligned [6]. Hence the target DOA (θ) could be roughly estimated from the camera view by locating the target speaker's face (see our actual hardware demo website: <https://yongxuustc.github.io/grnnbf>). Then the DOA guided directional feature (DF) [40], namely $d(\theta)$, is estimated by calculating the cosine similarity between the target steering vector \mathbf{v} and IPDs [40, 33]. Target DOA and $d(\theta)$ are speaker-dependent features which can be used to extract the target speech. LPS, IPD and $d(\theta)$ are merged and fed into a Conv-TasNet variant [9] with a fixed STFT encoder [37, 33]. A stack of eight successive dilated Conv-1D layers with 256 channels forms a network block and four blocks are piled together. The estimated cRF [35] size ($K \times K$) is empirically set to 3x3 [35, 28].

As for the RNN-BF module, the RNNs have 2-layer gated recurrent units (GRUs) with 500 hidden nodes. The non-linear DNN layer has 500 PReLU units. There are 30 linear units at the output DNN layer to predict the frame-wise beamforming weights. The model is trained in a chunk-wise mode with 4-second chunk size, using Adam optimizer. Initial learning rate is set to 1e-4. The objective is to maximize the time-domain scale-invariant source-to-noise ratio (Si-SNR) [9]. Pytorch 1.1.0 was used. Gradient norm is clipped with max norm 10. We evaluate the systems by using different metrics, including PESQ, Si-SNR (dB), signal-to-distortion ratio (SDR) (dB). A commercial general-purpose mandarin speech recognition Tencent API [41] is used to test the ASR performance in WER. Note this work only focuses on speech separation and denoising without dereverberation. Hence the reverberant clean (without dereverberation) is used as the reference signal.

5. Results and discussions

We evaluate the target speech separation performance in the overlapped multi-talker scenario. The spatial angle between the target speaker and others (interfering speakers) lies within 0-180°. The more overlapped speakers and the smaller spatial angle would lead to more challenging separation tasks. The detailed PESQ scores across different scenarios (i.e., angle between the target speaker and other speakers; number of overlapped speakers) are presented in Table 1. Other metrics are given with the average results.

GRNN-BF vs. traditional MVDRs: Two traditional MVDR systems, MVDR with mask normalization (ii) [6, 28] and multi-tap (i.e., $[t-1, t]$) MVDR with mask normalization (iii) [6, 28] are compared here. They also use the cRF estimator to calculate covariance matrices but replace the RNN-BF module (as shown in Fig. 1) with conventional MVDR or multi-tap MVDR [6] solutions. They both work reasonably well, e.g., the multi-tap MVDR (iii) achieves 13.52% WER. However, the proposed GRNN-BF with mask normalization (vii) could obtain significantly better performance. The proposed GRNN-BF (vii) increases the average PESQ to 3.52 from 3.08 of multi-tap MVDR (iii) and 2.92 of MVDR (ii). The WER of the proposed GRNN-BF (vii) is better than the multi-tap MVDR (iii), i.e., 11.86 vs. 13.52. The corresponding Si-SNR and SDR are in-

Table 1: PESQ, Si-SNR(dB) [9], SDR(dB) and WER(%) results among Conv-TasNet with STFT [33], several MVDRs and proposed GRNN-BF systems. "MN" and "LN" denote mask normalization and layer normalization on the covariance matrices, respectively.

systems/metrics	PESQ [-0.5, 4.5]								Si-SNR	SDR	WER
	Angle between target & others				# of overlapped speakers						
	0-15	15-45	45-90	90-180	1SPK	2SPK	3SPK	Avg.			
Reverberant clean reference	4.50	4.50	4.50	4.50	4.50	4.50	4.50	4.50	∞	∞	8.26
Mixture (overlapped speech+noise)	1.88	1.88	1.98	2.03	3.55	2.02	1.77	2.16	3.39	3.50	55.14
Conv-TasNet with STFT (i) [33]	2.75	2.95	3.12	3.09	3.98	3.06	2.76	3.10	12.50	13.01	22.07
MVDR w/ MN (ii) [6]	2.55	2.77	2.96	2.89	3.82	2.90	2.55	2.92	11.31	12.58	15.91
Multi-tap MVDR w/ MN (iii) [6]	2.67	2.95	3.15	3.10	3.92	3.06	2.72	3.08	12.66	14.04	13.52
ADL-MVDR w/ MN (iv) [28]	3.04	3.30	3.48	3.48	4.17	3.41	3.07	3.42	14.80	15.45	12.73
Prop. RNN-GEV w/ MN (v)	3.11	3.36	3.55	3.54	4.19	3.48	3.14	3.48	15.34	15.88	12.07
Prop. RNN-GEV w/ LN (vi)	3.15	3.39	3.57	3.56	4.19	3.51	3.17	3.51	15.55	16.07	11.75
Prop. GRNN-BF w/ MN (vii)	3.17	3.40	3.58	3.59	4.21	3.53	3.19	3.52	15.48	16.03	11.86
Prop. GRNN-BF w/ LN (viii)	3.23	3.45	3.62	3.60	4.23	3.57	3.24	3.56	15.84	16.38	11.36

creased to 15.48 dB and 16.03 dB, respectively. Fig. 2 also shows that the proposed GRNN-BF could estimate the spectra with less residual noise than traditional MVDR. The traditional MVDR has limited noise reduction capability [12, 6]. Finally, the proposed GRNN-BF with layer normalization (viii) achieves the best performance among all systems across all metrics. Moreover, the PESQ scores of our proposed GRNN-BF (viii) are above 3.2 at all scenarios, especially the two most difficult cases, namely small angle ($<15^\circ$) and three overlapped speakers (3SPK).

GRNN-BF vs. RNN-MVDR/GEV: Our proposed RNN-GEV uses RNNs to implement the GEV beamformer following Eq. (4) while the ADL-MVDR [28] following Eq. (3). With a more flexible structure (as shown in Sec. 3.2), the proposed RNN-GEV (v) is slightly better than the ADL-MVDR (iv), e.g., PESQ: 3.48 vs. 3.42; WER: 12.07 vs. 12.73. However, the proposed GRNN-BF (vii) is better than both of them. Compared to the ADL-MVDR (iv), the proposed GRNN-BF (vii) further improves the average PESQ from 3.42 to 3.52 and the average Si-SNR from 14.80 dB to 15.48 dB. Fig. 2 also shows that the proposed GRNN-BF can enhance the spectrogram with less residual noise than the ADL-MVDR. These results suggest that there is no need to follow any beamformers' solutions. The RNNs could learn a better solution from the speech and noise covariance matrices directly. The layer normalization is better than the mask normalization to normalize covariance matrices for both of the proposed RNN-GEV (vi) and GRNN-BF (viii).

GRNN-BF vs. Conv-TasNet: Conv-TasNet with a fixed STFT encoder [33] is our cRF estimator (as shown in Fig. 1), which is a variant of the original Conv-TasNet [9]. It is a purely "black-box" neural network system with the same multi-channel input. It predicts the target speech as Eq. (5) defined. The proposed GRNN-BF with layer normalization (viii) beats the Conv-TasNet with STFT (i) by a large margin, i.e., PESQ: 3.56 vs. 3.10; Si-SNR: 15.84 vs. 12.50; WER: 11.36 vs. 22.07. The Conv-TasNet results in the worst WER 22.07% among all systems due to the non-linear distortion which is quite common in most of purely neural network based speech separation systems [7, 42, 6]. This non-linear distortion can also be found in the separated spectrogram of Conv-TasNet in Fig. 2.

Layer normalization vs. mask normalization: As the denominator defined in Eq. (6), the mask normalization [32, 29, 6, 13] on the covariance matrix are always applied to stabilize the training. However, the proposed layer normalization on the covariance matrix (as defined in Eq. (7)) is more flexible than

the mask normalization. The proposed GRNN-BF with layer normalization (viii) obtains better performance than the GRNN-BF with mask normalization (vii), e.g., WER: 11.36 vs. 11.86; PESQ: 3.52 vs. 3.56.

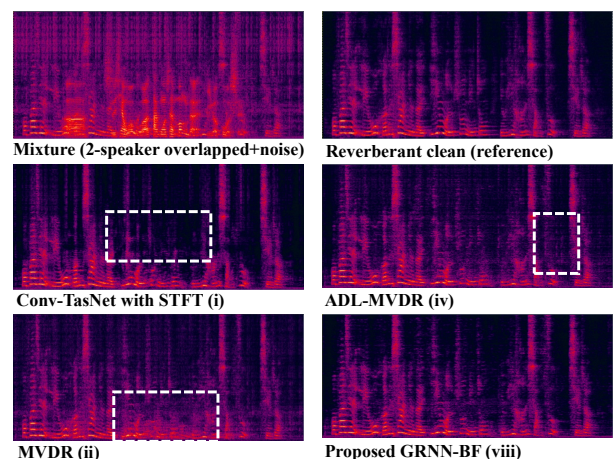


Figure 2: Sample separated spectrograms of different target speech separation systems. More testing demos (including real-world recording demos to verify the generalization capability) could be found at: <https://yongxuustc.github.io/grnnbf>.

6. Conclusions

In summary, we proposed a generalized RNN beamformer (GRNN-BF) that learns the frame-level beamforming weights directly from the estimated speech and noise covariance matrices. The layer normalization achieves better performance than the mask normalization for normalizing the speech and noise covariance matrices. The proposed GRNN-BF with layer normalization achieves the best objective scores (PESQ, Si-SNR, SDR) and the lowest WER among all evaluated systems. It achieves relative 10.8% and 16% WER reduction against the prior art methods, ADL-MVDR [28] and the conventional multi-tap MVDR [6], respectively. Although we only tested the proposed GRNN-BF on the DOA-guided target speech separation task, it could also be used for general scenarios without DOA information, e.g., the multi-channel speech enhancement task or the permutation invariant training (PIT) [10] based multi-channel speech separation task.

7. References

- [1] J. Heymann, L. Drude, and et al., “Neural network based spectral mask estimation for acoustic beamforming,” in *ICASSP*, 2016.
- [2] Z.-Q. Wang and D. Wang, “Mask weighted stft ratios for relative transfer function estimation and its application to robust asr,” in *ICASSP*, 2018, pp. 5619–5623.
- [3] H. Erdogan, J. R. Hershey, and et al., “Improved MVDR beamforming using single-channel mask prediction networks,” in *Interspeech*, 2016.
- [4] J. Heymann, L. Drude, and et al., “Beamnet: End-to-end training of a beamformer-supported multi-channel ASR system,” in *ICASSP*, 2017.
- [5] X. Xiao, S. Zhao, and et al., “On time-frequency mask estimation for MVDR beamforming with application in robust speech recognition,” in *ICASSP*, 2017.
- [6] Y. Xu, M. Yu, and et al., “Neural spatio-temporal beamformer for target speech separation,” *Interspeech*, 2020.
- [7] J. Du, Q. Wang, and et al., “Robust speech recognition with speech enhanced deep neural networks,” in *Interspeech*, 2014.
- [8] Y. Wang, A. Narayanan, and D. Wang, “On training targets for supervised speech separation,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 22, no. 12, pp. 1849–1858, 2014.
- [9] Y. Luo and N. Mesgarani, “Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [10] D. Yu, M. Kolbæk, and et al., “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *ICASSP*, 2017.
- [11] Y. Xu, J. Du, and et al., “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.
- [12] E. A. Habets and J. Benesty, “A two-stage beamforming approach for noise reduction and dereverberation,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 5, pp. 945–958, 2013.
- [13] W. Zhang, C. Boeddeker, S. Watanabe, and et al., “End-to-end dereverberation, beamforming, and speech recognition with improved numerical stability and advanced frontend,” *arXiv preprint arXiv:2102.11525*, 2021.
- [14] S. Chakrabarty and E. A. Habets, “On the numerical instability of an LCMV beamformer for a uniform linear array,” *IEEE Signal Processing Letters*, vol. 23, no. 2, pp. 272–276, 2015.
- [15] C. Y. Lim, C.-H. Chen, and W.-Y. Wu, “Numerical instability of calculating inverse of spatial covariance matrices,” *Statistics & Probability Letters*, vol. 129, pp. 182–188, 2017.
- [16] S. Zhao and D. L. Jones, “A fast-converging adaptive frequency-domain MVDR beamformer for speech enhancement,” in *Interspeech*, 2012.
- [17] X. Mestre and M. A. Lagunas, “On diagonal loading for minimum variance beamformers,” in *Proceedings of the 3rd IEEE International Symposium on Signal Processing and Information Technology*, 2003, pp. 459–462.
- [18] Z.-Q. Wang, H. Erdogan, and et al., “Sequential multi-frame neural beamforming for speech separation and enhancement,” *arXiv preprint arXiv:1911.07953*, 2019.
- [19] Y. Kubo, T. Nakatani, and et al., “Mask-based MVDR beamformer for noisy multisource environments: introduction of time-varying spatial covariance model,” in *ICASSP*, 2019.
- [20] J. Wang, “A recurrent neural network for real-time matrix inversion,” *Applied Mathematics and Computation*, vol. 55, no. 1, pp. 89–100, 1993.
- [21] Y. Zhang and S. S. Ge, “Design and analysis of a general recurrent neural network model for time-varying matrix inversion,” *IEEE Transactions on Neural Networks*, vol. 16, no. 6, pp. 1477–1490, 2005.
- [22] L. Liu, H. Shao, and et al., “Recurrent neural network model for computing largest and smallest generalized eigenvalue,” *Neurocomputing*, vol. 71, no. 16-18, pp. 3589–3594, 2008.
- [23] X. Wang, M. Che, and et al., “Recurrent neural network for computation of generalized eigenvalue problem with real diagonalizable matrix pair and its applications,” *Neurocomputing*, vol. 216, pp. 230–241, 2016.
- [24] J. Benesty, J. Chen, and Y. Huang, *Microphone array signal processing*. Springer Science & Business Media, 2008, vol. 1.
- [25] J. Heymann, L. Drude, A. Chinaev, and R. Haeb-Umbach, “BLSTM supported GEV beamformer front-end for the 3rd CHiME challenge,” in *ASRU*, 2015, pp. 444–451.
- [26] F. Grondin, J.-S. Lauzon, and et al., “GEV beamforming supported by DOA-based masks generated on pairs of microphones,” *arXiv preprint arXiv:2005.09587*, 2020.
- [27] T. Van den Bogaert, S. Doclo, and et al., “Speech enhancement with multichannel wiener filter techniques in multimicrophone binaural hearing aids,” *The Journal of the Acoustical Society of America*, vol. 125, no. 1, pp. 360–371, 2009.
- [28] Z. Zhang, Y. Xu, and et al., “ADL-MVDR: All deep learning MVDR beamformer for target speech separation,” *ICASSP*, 2021.
- [29] X. Xiao, C. Xu, and et al., “A study of learning based beamforming methods for speech recognition,” in *CHiME 2016 workshop*, 2016.
- [30] Z. Meng, S. Watanabe, and et al., “Deep long short-term memory adaptive beamforming networks for multichannel robust speech recognition,” in *ICASSP*, 2017.
- [31] J. L. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” *arXiv preprint arXiv:1607.06450*, 2016.
- [32] C. Boeddeker, H. Erdogan, T. Yoshioka, and R. Haeb-Umbach, “Exploring practical aspects of neural mask-based beamforming for far-field speech recognition,” in *ICASSP*, 2018, pp. 6697–6701.
- [33] R. Gu, S.-X. Zhang, and et al., “Multi-modal multi-channel target speech separation,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 3, pp. 530–541, 2020.
- [34] D. S. Williamson, Y. Wang, and et al., “Complex ratio masking for monaural speech separation,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 24, no. 3, pp. 483–492, 2015.
- [35] W. Mack and E. A. Habets, “Deep filtering: Signal extraction and reconstruction using complex time-frequency filters,” *IEEE Signal Processing Letters*, vol. 27, pp. 61–65, 2019.
- [36] G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [37] K. Tan, Y. Xu, and et al., “Audio-visual speech separation and dereverberation with a two-stage multimodal network,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 3, pp. 542–553, 2020.
- [38] S.-X. Zhang, Y. Xu, and et al., “M³: Multi-Modal Multi-channel dataset for cocktail party problems,” in *preparation*, 2020.
- [39] E. A. Habets, “Room impulse response generator,” *Technische Universiteit Eindhoven, Tech. Rep.*, vol. 2, no. 2.4, 2006.
- [40] Z. Chen, X. Xiao, and et al., “Multi-channel overlapped speech recognition with location guided speech extraction network,” in *SLT*, 2018.
- [41] “Tencent ASR,” <https://ai.qq.com/product/aaiasr.shtml>.
- [42] Y. Luo, C. Han, and N. Mesgarani, “Distortion-controlled training for end-to-end reverberant speech separation with auxiliary autoencoding loss,” *arXiv preprint arXiv:2011.07338*, 2020.