A NEURAL BEAM FILTER FOR REAL-TIME MULTI-CHANNEL SPEECH ENHANCEMENT

Wenzhe Liu[*†], *Andong Li*[*†], *Chengshi Zheng*[*†], *Xiaodong Li*[*†]

[*]Key Laboratory of Noise and Vibration Research, Institute of Acoustics, Chinese Academy of
Sciences, Beijing, China
[†]University of Chinese Academy of Sciences, Beijing, China

ABSTRACT

Most deep learning-based multi-channel speech enhancement methods focus on designing a set of beamforming coefficients to directly filter the low signal-to-noise ratio signals received by microphones, which hinders the performance of these approaches. To handle these problems, this paper designs a causal neural beam filter that fully exploits the spatial-spectral information in the beam domain. Specifically, multiple beams are designed to steer towards all directions using a parameterized super-directive beamformer in the first stage. After that, the neural spatial filter is learned by simultaneously modeling the spatial and spectral discriminability of the speech and the interference, so as to extract the desired speech coarsely in the second stage. Finally, to further suppress the interference components especially at low frequencies, a residual estimation module is adopted to refine the output of the second stage. Experimental results demonstrate that the proposed approach outperforms many state-of-the-art multi-channel methods on the generated multi-channel speech dataset based on the DNS-Challenge dataset.

Index Terms— Multi-channel speech enhancement, neural beam filter, deep learning, causal

1. INTRODUCTION

Multi-channel speech enhancement (MCSE) can efficiently suppress directional interference and improve speech quality with beamforming and/or post-filtering [1]. It has been widely applied as a preprocessor in video conferencing systems, automatic speech recognition (ASR) systems, and smart TVs. The recent breakthrough in deep neural networks (DNNs) has facilitated the research in MCSE, which yields notable performance improvements over conventional statistical beamforming techniques [2–9].

Considering the success of DNNs in the single-channel speech enhancement (SCSE) area, a straightforward strategy is to extend the previous SCSE models to extract spatial features either heuristically or implicitly [2–7]. This paradigm is prone to cause non-linear speech distortion such as spectral blackhole in low signal-to-noise (SNR) scenarios since the advantage of the spatial filter with microphone-array beamforming is not fully exploited to null the directional interference and suppress the ambient noise [8, 9]. Another category follows the cascade-style regime. To be specific, in the first stage, an SC-based network was adopted to predict the mask of each acoustic channel in parallel, followed by spatial covariance matrices (SCMs) calculation *w.r.t.* speech and noise. In the second stage, traditional beamformers like minimum variance distortionless response (MVDR) or eigenvalue decomposition (GEV) were adopted for spatial filtering [8, 10–12]. These methods have shown their effectiveness in ASR when the latency is not a critical issue, *i.e.*, if it can be up to hundreds of milliseconds. When the latency should be much lower, *i.e.*, ≤ 20 ms, for many practical applications, such as speech communication, hearing aids, and transparency, these methods may degrade their performance significantly for these low-latency systems. Moreover, the performance heavily depends on the previous mask estimation, which can degrade a lot in complex acoustic scenarios.

As a solution, an intuitive tactic is to enforce the network to directly predict the beamforming weights, which can be done in either time domain [13, 14] or frequency domain [9, 15–17]. Nonetheless, according to the signal theory, the desired beam pattern is required to form its main beam towards the target direction and meanwhile form the null towards the interference direction, which tends to be difficult from the optimization perspective. Moreover, because the network directly determines the weights of the multi-channel input, such a paradigm usually inevitably learns the internal relation of the array geometry, index sequence and the predicted output. In other words, the network itself is not fully decoupled with the inherent information of the array geometry.

In this paper, we design a neural beam filter for real-time multi-channel speech enhancement. In detail, the multi-channel signals are first processed by a set of pre-defined fixed beamformers, *i.e.*, we uniformly sample a beam set with various directions in the space. Then the network is utilized to learn the temporal-spectral-spatial discriminative features of target speech and noise, which aims to generate the bin-level filtering coefficients to automatically weight the beam set. Note that different from previous neural beamformer-based literature where the output weights are applied to multi-channel input signals directly, here the predicted coefficients are to filter the noise component of each pre-generated beam and fuse them. We dub it a neural beam filter to distinguish it from the existing neural beamformer literally. The rationale of such network design logic has three-fold. First, the target signal can be pre-extracted with the fixed beamformer, and the dominant part should exist within at least one or multiple directional beams, which serves as the SNR-improved target priori to guide the subsequent beam fusion process. Second, the interference-dominant beam can be obtained when the beam steers towards the interference direction, which can provide the interference priori for better distinguishment in a spatial-spectral sense. Besides, the target and interference components may co-exist within each beam while their distributions are dynamically changed due to their spectral difference. Therefore, the beam set can be viewed as a reasonable candidate to indicate the spectral and spatial characteristics. As the beam set is only discretely sampled in the space, the information loss tends to arise due to the limited spatial resolution at low frequencies, which causes speech distortion. To this end, a residual branch is designed to refine the fused beam. We have to emphasize that, although the multi-beam concept is used in both [18] and this study, they are very different as [18] is in essence a parallel single beam enhancement
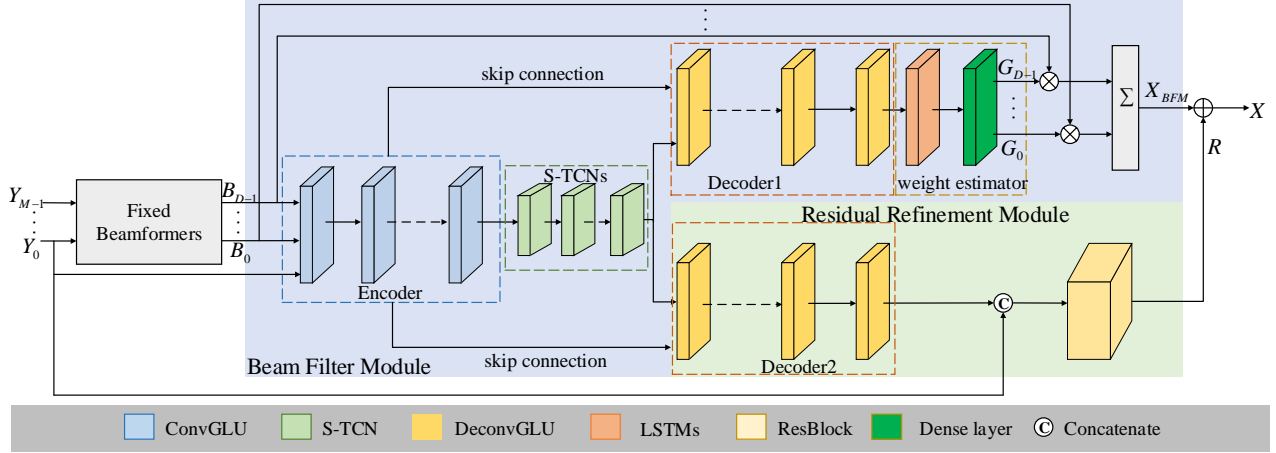
Fig. 1. Overview of the proposed framework. Different modules are highlighted with different colors.

process while the proposed system can be regarded as the filter and fusion process of multi-beam. Experiments conducted on the DNS-Challenge corpus [19] show that the proposed neural beam filter outperforms previous state-of-the-art (SOTA) baselines.

The remainder of the paper is organized as follows. We describe the proposed neural beam filter in Section 2. Experimental setting and results are given in Section 3 and Section 4, respectively. Finally, we draw some conclusions in Section 5.

2. PROPOSED SYSTEM

2.1. Forward Stream

Fig. 1 shows the overall diagram of the proposed architecture, which consists of three components, namely fixed beamforming module (FBM), beam filtering module (BFM), and residual refinement module (RRM).

Let us denote the N-point short-time Fourier transform (STFT) of M-channel input mixtures as $Y_m(t, f) \in \mathbb{C}$, with $m = 0, \cdots, M-1$, where $t \in \{0, \cdots, T-1\}$ and $f \in \{0, \cdots, F-1\}$ refer to the index of frames and that of frequency bins, respectively. Considering the symmetry of $Y_m(t, f)$ in frequency, $F = N/2 + 1$ is chosen throughout this paper. In FBM, the fixed beamformer is employed to sample the space uniformly and obtain multiple beams steering towards different directions, *i.e.*, $B_d \in \mathbb{C}^{T \times F}$, with $d = 0, \cdots, D-1$, where D denotes the number of resultant multi-beam. The process is thus given by:

$$\{B_d\}_{d=0,\cdots,D-1} = \mathcal{F}_{FBM}(\{Y_m\}_{m=0,\cdots,M-1}; \Phi_{FBM}), \quad (1)$$

where $\mathcal{F}_{FBF}$ is the function of FBM and Φ_{FBF} denotes the parameter set. From now on, we will omit the subscript (t, f) when no confusion arises.

We concatenate the beam set along the channel dimension, serving as the input of BFM, *i.e.*, $Cat(B_0, \cdots, B_{D-1}) \in \mathbb{R}^{2D \times T \times F}$. Here 2 means that both real and imaginary (RI) parts are considered. As muti-beams can represent both spectral and spatial characteristics, BFM is adopted to learn the temporal-spatial-spectral discriminative information between speech and interference and attempt to assign the filter weights $\widehat{G}_d \in \mathbb{C}^{T \times F}$ for each beam. It is worth noting that as the beam set is discretely sampled in the space, the information loss tends to arise due to the limited spatial resolution. To alleviate this problem, the complex spectrum of the reference channel is also incorporated into the input and meanwhile, similar to [27], the complex residual needs to be estimated with RRM, which aims to compensate for the inherent information loss of the filtered spec-

trum. This process can be presented as:

$$\widehat{G} = \mathcal{F}_{BFM}([\{B_d\}_{d=0,\cdots,D-1}, Y_0]; \Phi_{BFM}), \quad (2)$$

$$\widehat{R} = \mathcal{F}_{RRM}([\{B_d\}_{d=0,\cdots,D-1}, Y_0]; \Phi_{RRM}), \quad (3)$$

where $\widehat{G} \in \mathbb{C}^{D \times T \times F}$ and $\widehat{R} \in \mathbb{C}^{D \times T \times F}$ denote the complex filter weights and the complex residual estimated by BFM and RRM, respectively. By applying the estimated weights $\{\widehat{G}_d\}_{d=0,\cdots,D-1}$ to filter the beams $\{B_d\}_{d=0,\cdots,D-1}$ and then summing them along the channel axis, the fused beam $\widehat{X}_{BFM}$ can be obtained by:

$$\widehat{X}_{BFM} = \sum_d \widehat{G}_d * B_d, \quad (4)$$

where $*$ denotes the complex-valued multiplication operator. We then add the filtered beam and estimated complex residual together to obtain the final output $\widehat{X}$, *i.e*,

$$\widehat{X} = \widehat{X}_{BFM} + \widehat{R}. \quad (5)$$

2.2. Fixed Beamforming Module

In this module, the fixed beamformer is leveraged to transform input multi-channel mixtures into several beams, which steer towards different looking directions and uniformly sample the space. As the fixed beamformer is data-independent, it is robust in adverse environments and has low computational complexity. Moreover, filtering multi-channel mixtures with the fixed beamformer allows our system to be less sensitive to the array geometry and is suitable for arbitrary microphone arrays. In this paper, we choose the super-directivity (SD) beamformer as the default beamformer due to its promising performance in high directivity [20]. Note that other fixed beamformers can also be adopted, which is out of the scope of the paper. Assume the target directional angle is θ_d, the weights of the SD beamformer can be calculated as:

$$\mathbf{w}_d(f) = \frac{\Gamma_{nn}^{-1}(f)\mathbf{v}(\theta_d, f)}{\mathbf{v}^H(\theta_d, f)\Gamma_{nn}^{-1}(f)\mathbf{v}(\theta_d, f)}, \quad (6)$$

where $\mathbf{v}(\theta_d, f)$ is the steering vector, $(\cdot)^H$ is the complex transpose operator and $\Gamma_{nn}(f)$ denotes the covariance matrix of a diffuse noise field with the diagonal loading to control the white noise gain, where we denote it as the parameterized SD beamformer. Note that the diagonal loading level often needs to be chosen carefully to make a good balance between the white noise gain and the array gain [21]. In this paper, the diagonal loading level is fixed to 1e-5 and its impact on performance will be studied in the near future. The (i, j)-th element of $\Gamma_{nn}(f)$ represents the coherence between the signals received by two microphones with indices i and j in an

isotropic diffuse field, which can be formulated as:

$$\mathbf{\Gamma}_{nn}^{(i,j)}(f) = \mathrm{sinc}\left(\frac{2\pi f_s f l_{ij}/N}{c}\right),\qquad(7)$$

where $\mathrm{sinc}(x) = \sin(x)/x$ and l_{ij} is the distance between the i-th and j-th microphones, c is the speed of sound and f_s is the sampling rate.

Define $\mathbf{Y}(t,f) = \{Y_0(t,f),\cdots,Y_{M-1}(t,f)\}$, the output of the d-th SD beamformer can be expressed as:

$$B_d(t,f) = \mathbf{w}_d^H(f)\mathbf{Y}(t,f).\qquad(8)$$

2.3. Beam Filter Module

As shown in Fig. 1, the beam filter module (BFM) consists of a causal convolutional encoder-decoder (CED) architecture and a weight estimator. For the encoder, it comprises six two-dimensional (2-D) convolutional gated linear units (ConvGLUs) [23], which consecutively halve the feature size and extract high-level features. Each GLU is followed by instance normalization (IN) and Parameter ReLU (PReLU). The decoder is the mirror version of the encoder except that all the convolution operations are replaced by the deconvolutional version (dubbed DeconvGLU). Similar to [22], a stack of squeezed temporal convolutional networks (S-TCNs) is inserted as the bottleneck of CED to model the temporal correlations among adjacent frames. After that, in the weight estimator, we simulate the filter generation process, where T-F bin-level filter coefficients are assigned for each beam. To be specific, the output embedding tensor of the decoder is first normalized by layer normalization (LN) and then the LSTM is employed to update the feature frame by frame with ReLU serving as the intermediate nonlinear activation function. The weights $\widehat{\mathbf{G}}$ are obtained after the output linear layer. Then these weights are applied to each beam to obtain the target beam.

2.4. Residual Refinement Module

Because the SD beamformer tends to amplify the white noise to ensure the array gain at low frequencies, the weighted beam output with BFM often contains lots of residual noise components, which need to be further suppressed to improve speech quality. Meanwhile, speech distortion is often introduced due to the mismatch between the main beam steering toward the predefined direction and the true direction of the target speech, which is because the number of the fixed beamformers is limited. To refine the target beam, a residual refinement module (RRM) is proposed, which comprises a decoder module similar to that of BFM and a ResBlock containing three residual blocks, as shown in Fig. 1. The output of S-TCNs serves as the input feature of the RRM. After decoding from multiple DeGLUs, the output tensor is concatenated with the original complex spectrum of the reference microphone Y_0 and is fed to a series of residual blocks. Finally, the complex residual spectrum $\widehat{\mathbf{R}}$ is derived by the output 1×1-conv to reduce the dimensions to 2, and applied to refine the filtered beam output $\widehat{X}_{BFM}$.

3. EXPEREIMENTAL SETUP

3.1. Datasets

DNS-Challenge corpus is selected to convolve with multi-channel RIRs to generate multi-channel pairs for experiments. To be specific, the clean clips are randomly sampled from *neutral clean speech* set, which includes about 562 hours speaking by 11,350 speakers. We split it into two parts without overlap, namely for training and testing. For training set, around 20,000 types of noise are selected, with the duration time about 55 hours [22, 27]. For testing set, four types of unseen noise are chosen, namely babble, factory1, white noises

from NOISEX92 [25], and cafe noise from CHiME3 [26]. We generate the RIRs with the image method [24] based on a uniform linear array with 9 microphones, whose spacing distance is around 4 cm between two neighboring microphones. The room size is sampled from $3\times3\times2.5$ m^3 to $10\times10\times3$ m^3 and the reverberation time RT60 ranges from 0.05s to 0.7s. The source is randomly located in angle from $0°$ to $180°$, and the distance between the source and the array center ranges from 0.5 to 3.0 m. The signal-to-interference ratio (SIR) ranges from -6 to 6 dB. Totally, about 80,000 and 4000 multi-channel reverberant noisy mixtures are generated for training and validation. For testing set, five SIRs are set, *i.e.*, $\{-5,-2,0,2,5\}$ dB and 150 pairs are generated for each case.

3.2. Baselines

In this paper, two state-of-the-art (SOTA) single-channel speech enhancement systems namely CTSNet [22] and GaGNet [27] and three multi-channel speech enhancement systems namely MC-Conv-TasNet [28], FaSNet-TAC [13], and MIMO-UNet [17] are chosen as the comparative systems. CTSNet is a two-stage framework that ranked 1st in the 2nd DNS Challenge [19]. As the parallel version of CTSNet, GaGNet is proposed and achieves higher performance. In this paper, we select these two systems to represent the SOTA performance in the monaural speech enhancement area. MC-Conv-TasNet is the multiple-input version of Conv-TasNet which is one of the most effective time-domain speech enhancement and separation models. FaSNet-TAC is an end-to-end filter-and-sum style time-domain multi-channel speech enhancement system, which can achieve better performance than mask-based beamformers. MIMO-UNet is a frequency-domain neural beamformer which is the winner of the INTERSPEECH Far-field Multi-Channel Speech Enhancement Challenge for Video Conferencing [29]. Note that all the models are set with causal configuration, that is, no future frames are involved into the calculation and inference of current frame.

3.3. Experiment Setup

3.3.1. Network Detail

In the encoder and decoder part, the kernel size and stride of the 2-D convolution layers are $(2,3)$ and $(1,2)$, respectively, and the number of channels remains 64 by default. 3 S-TCNs are adopted, each of which consists of 6 temporal convolutional modules (TCMs) with kernel size and dilation rate being 5 and $(1,2,4,8,16,32)$. In the weight estimator module, 2 uni-directional LSTM layers are employed with 64 hidden nodes. In the ResBlock, the kernel size and stride are $(2,3)$ and $(1,1)$, respectively.

3.3.2. Training Detail

All the utterances are sampled at 16 kHz. 32 ms Hann window is utilized, with 50% overlap between adjacent frames. Accordingly, 512-point FFT is utilized, leading to 257-D spectral features. Adam optimizer is applied with the initial learning rate set to 5e-4. If validation loss does not decrease for consecutive two epochs, the learning rate will be halved. All models are trained for 60 epochs.

4. RESULTS ANALYSIS

We choose perceptual evaluation speech quality (PESQ) [30] and extended short-time objective intelligibility (ESTOI) [31] as objective metrics to compare the performance of different models. [1]

---

[1] Audio samples are available online: https://wenzhe-liu.github.io/NBF/

Table 1. Objective result comparisons among different causal MCSE models in terms of PESQ and ESTOI for unseen noises. **BOLD** indicates the best score in each case.

Metrics	PESQ						ESTOI (in %)					
SNR (dB)	-5	-2	0	2	5	Avg.	-5	-2	0	2	5	Avg.
Unprocessed	1.36	1.44	1.58	1.70	1.90	1.60	29.99	38.90	44.72	50.66	59.22	44.70
CTSNet	1.85	2.14	2.28	2.45	2.63	2.27	44.55	57.12	63.10	68.67	75.61	61.81
GaGNet	1.86	2.18	2.33	2.54	2.73	2.33	46.54	59.11	64.68	70.22	77.08	63.53
MC-Conv-TasNet	2.47	2.59	2.77	2.82	2.97	2.72	72.16	74.82	79.64	80.28	84.66	78.31
MIMO-UNet	2.43	2.63	2.80	2.88	2.95	2.74	66.73	72.91	78.18	79.51	82.23	75.91
FaSNet-TAC	2.61	2.74	2.92	2.97	3.18	2.88	76.71	79.84	85.50	85.30	89.75	83.42
Proposed-19beams (w/o RRM)	2.95	3.11	3.30	3.36	3.57	3.26	78.53	82.34	86.89	86.87	91.10	85.15
Proposed-10beams	3.00	3.16	3.36	3.39	3.60	3.30	79.27	83.00	87.35	87.25	91.52	85.68
Proposed-19beams	3.10	3.24	3.45	3.47	3.67	3.39	81.54	84.56	88.93	88.47	92.42	87.18
Proposed-37beams	**3.14**	**3.28**	**3.49**	**3.51**	**3.71**	**3.43**	**82.18**	**85.21**	**89.39**	**88.91**	**92.77**	**87.69**
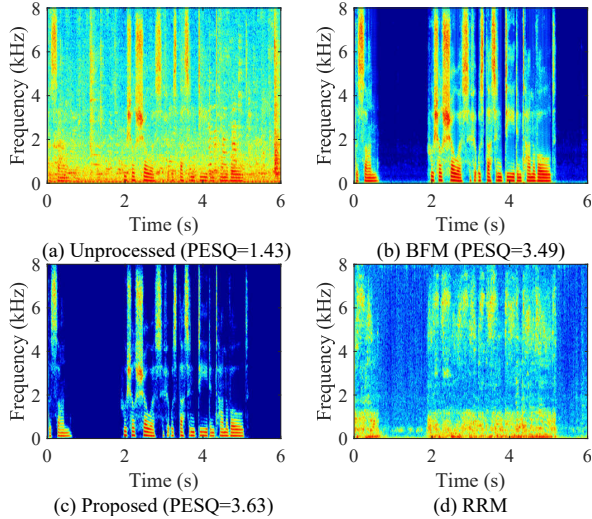
Fig. 2. Visualization results of (a) unprocessed (SIR=0dB), (b) BFM, (c) proposed-19beams, (d) RRM.

Fig. 3. Visualization results of (a) the spectrogram of input signal, and (b) the estimated filter weights.

4.1. Results Comparison in Objective Metrics

Table 1 shows the objective results of different SE systems. For comparison, the number of beams D is set to 19, *i.e.*, the sampling resolution is $10°$. We evaluate these systems in terms of PESQ and ESTOI. From the table, several observations can be made. First, compared with single-channel speech enhancement approaches such as CT-SNet and GaGNet, all the multi-channel speech enhancement systems yields notable performance improvements consistently thanks to the utilization of spatial information, which indicates that multi-channel information is beneficial for distinguishing speech and interference components. Second, the proposed system outperforms neural beamforming-based approaches by a large margin. For example, compare with FaSNet-TAC, our system achieves 0.51, and 3.76% improvements in terms of PESQ and ESTOI, respectively. Moreover, our model outperforms MIMO-UNet by 0.65 and 11.27% in PESQ and ESTOI, respectively. This demonstrates the superiority of filtering the beams over the best neural spatial filters based on frequency and time domains.

4.2. Ablation Analysis

We also validate the role of FBM, BFM and RRM. To analyze the effectiveness of FBM, we set another two candidates of D, *i.e.*, 10 $(20°)$, 37 $(5°)$, where 10 $(20°)$ means $D = 10$ and each main beam width is about $20°$; 37 $(5°)$ analogously. It can be seen from Table 1 that the performance of the beam neural filter gradually improves with the increase of D, which reveals the importance of the FBM. However, the relative performance improvement decreases as
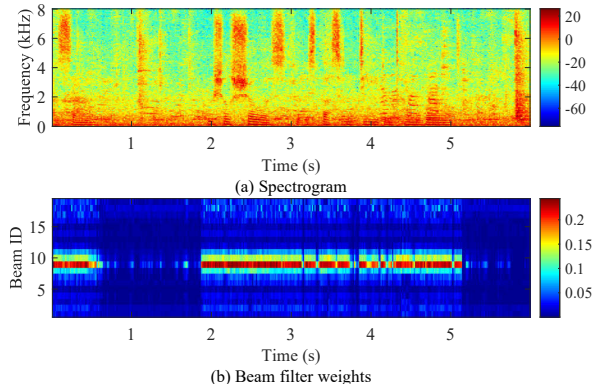
the spatial sampling interval becomes progressively smaller although there is still a mismatch between beam pointing and source direction, which indicates that the proposed model is robust to direction mismatch whereas the spatial filter is more sensitive to direction estimation error. To show the effectiveness of BFM, we visualize the norm of estimated complex weights in Fig. 3. The input signals are mixed by a speech radiating from $85°$ and a Factory1 noise source from $45°$. We can find that greater weights are assigned to beams steering toward the surroundings of the target direction, and beams steering to other directions, including those steering toward the interference direction, are given little weights during speaking, while all weights are small in non-speech segments. Besides, the proposed system with RRM achieves PESQ improvements of 0.13 and ES-TOI improvements of 2.03%. Comparing the visualization results of the model with and without RRM, one can see that the residual noise components are further suppressed at low frequencies and some missing speech components are recovered, which confirms the effectiveness of RRM in the proposed system.

5. CONCLUSIONS

In this paper, we propose a causal beam neural filter for real-time multi-channel speech enhancement. It comprises three components, namely FBM, BFM and RRM. Firstly, FBM is adopted to separate the sources from different directions. Then BFM maps filter weights by jointly learning the temporal-spectral-spatial discriminability of speech and interference. Finally, RRM is adopted to refine the weighting beam output. Experimental results show that the proposed system achieves better speech quality and intelligibility over previous spatial neural filters.

6. REFERENCES

[1] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Trans. Audio Speech Lang. Proc.*, vol. 26, no. 10, pp. 1702–1726, 2018.

[2] N. Tawara, T. Kobayashi, and T. Ogawa, "Multi-channel speech enhancement using time-domain convolutional denoising autoencoder," in *Proc. Interspeech*, pp. 86–90, 2019.

[3] C. Liu, S. Fu, Y. Li, J. Huang, H. Wang, and Y. Tsao, "Multi-channel speech enhancement by raw waveform-mapping using fully convolutional networks," *IEEE/ACM Trans. Audio Speech Lang. Proc.*, vol. 28, pp. 1888–1900, 2020.

[4] K. Tan, Y. Xu, S. Zhang, M. Yu, and D. Yu, "Audio-visual speech separation and dereverberation with a two-stage multimodal network," *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 3, pp. 542–553, 2020.

[5] J. Wu, Z. Chen, J. Li, T. Yoshioka, Z. Tan, E. Lin, Y. Luo, and L. Xie, "An end-to-end architecture of online multi-channel speech separation," in *Proc. Interspeech*, pp. 81–85, 2020.

[6] R. Gu, L. Chen, S. 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. Interspeech*, pp. 4290–4294, 2019.

[7] Y. Fu, J. Wu, Y. Hu, M. Xing, and L. Xie, "Desnet: A multi-channel network for simultaneous speech dereverberation, enhancement and separation," in *Proc. SLT*, pp. 857–864, 2021.

[8] Y. Xu, M. Yu, S. Zhang, L. Chen, C. Weng, J. Liu, and D. Yu, "Neural Spatio-Temporal Beamformer for Target Speech Separation," in *Proc. Interspeech*, pp. 56–60, 2020.

[9] Z. Zhang, Y. Xu, M. Yu, S. Zhang, L. Chen, and D. Yu, "Adl-mvdr: All deep learning mvdr beamformer for target speech separation," in *Proc. ICASSP*, pp. 6089–6093, 2021.

[10] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *Proc. ICASSP*, pp. 196–200, 2016.

[11] X. Zhang, Z. Wang, and D. Wang, "A speech enhancement algorithm by iterating single- and multi-microphone processing and its application to robust asr," in *Proc. ICASSP*, pp. 276–280, 2017.

[12] R. Gu, S. Zhang, Y. Zou, and D. Yu, "Complex neural spatial filter: Enhancing multi-channel target speech separation in complex domain," *IEEE Signal Proc. Let.*, 2021.

[13] Y. Luo, C. Han, N. Mesgarani, E. Ceolini, and S. Liu, "Fasnet: Low-latency adaptive beamforming for multi-microphone audio processing," in *Proc. ASRU*, pp. 260–267, 2019.

[14] Y. Luo, Z. Chen, N. Mesgarani, and T. Yoshioka, "End-to-end microphone permutation and number invariant multi-channel speech separation," in *Proc. ICASSP*, pp. 6394–6398, 2020.

[15] X. Xiao, S. Watanabe, H. Erdogan, L. Lu, J. Hershey, M. L Seltzer, G. Chen, Y. Zhang, M. Mandel, and D. Yu, "Deep beamforming networks for multi-channel speech recognition," in *Proc. ICASSP*, pp. 5745–5749, 2016.

[16] Y. Xu, Z. Zhang, M. Yu, S. Zhang, and D. Yu, "Generalized spatio-temporal rnn beamformer for target speech separation," *arXiv preprint arXiv:2101.01280*, 2021.

[17] X. Ren, X. Zhang, L. Chen, X. Zheng, C. Zhang, L. Guo, and B. Yu, "A causal u-net based neural beamforming network for real-time multi-channel speech enhancement," in *Proc. Interspeech*, pp. 1832–1836, 2021.

[18] Z. Chen, J. Li, X. Xiao, T. Yoshioka, H. Wang, Z. Wang, and Y. Gong, "Cracking the cocktail party problem by multi-beam deep attractor network," in *Proc. ASRU*, pp. 437–444, 2017.

[19] C. Reddy, H. Dubey, V. Gopal, R. Cutler, S. Braun, H. Gamper, R. Aichner, and S. Srinivasan, "Icassp 2021 deep noise suppression challenge," in *Proc. ICASSP*, pp. 6623–6627, 2021.

[20] A. T. Parsons, "Maximum directivity proof for three-dimensional arrays," *J. Acoust. Soc. Amer.*, vol. 82, no. 1, pp. 179–182, 1987.

[21] C. Pan, J. Chen and J. Benesty, "Reduced-Order Robust Superdirective Beamforming With Uniform Linear Microphone Arrays," *IEEE/ACM Trans. Audio Speech Lang. Proc.*, vol. 24, no. 9, pp. 1548-1559, 2016.

[22] A. Li, W. Liu, C. Zheng, C. Fan, and X. Li, "Two heads are better than one: A two-stage complex spectral mapping approach for monaural speech enhancement," *IEEE/ACM Trans. Audio Speech Lang. Proc.*, vol. 29, pp. 1829–1843, 2021.

[23] K. Tan and D. Wang, "Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement," *IEEE/ACM Trans. Audio Speech Lang. Proc.*, vol. 28, pp. 380–390, 2019.

[24] J. Allen and D. Berkley, "Image method for efficiently simulating small-room acoustics," *J. Acoust. Soc. Am.*, vol. 65, no. 4, pp. 943–950, 1979.

[25] A. Varga and H. Steeneken, "Assessment for automatic speech recognition: Ii. noisex-92: A database and an experiment to study the effect of additive noise on speech recognition systems," *Speech Commun.*, vol. 12, no. 3, pp. 247–251, 1993.

[26] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third 'chime'speech separation and recognition challenge: Dataset, task and baselines," in *Proc. ASRU*, pp. 504–511, 2015.

[27] A. Li, C. Zheng, L. Zhang, and X. Li, "Glance and gaze: A collaborative learning framework for single-channel speech enhancement," *arXiv preprint arXiv:2106.11789*, 2021.

[28] J. Zhang, C. Zorilă, R. Doddipatla and J. Barker, "On End-to-end Multi-channel Time Domain Speech Separation in Reverberant Environments," in *Proc. ICASSP*, pp. 6389-6393, 2020.

[29] W. Rao, Y. Fu, Y. Hu, X. Xu, Y. Jv, J. Han, Z. Jiang, L. Xie, Y. Wang, S. Watanabe, et al., "Interspeech 2021 conferencingspeech challenge: Towards far-field multi-channel speech enhancement for video conferencing," *arXiv preprint arXiv:2104.00960*, 2021.

[30] A. Rix, J. Beerends, M. Hollier, and A. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *Proc. ICASSP*, vol. 2, pp. 749–752, 2001.

[31] J. Jensen and C. Taal, "An algorithm for predicting the intelligibility of speech masked by modulated noise maskers," *IEEE/ACM Trans. Audio Speech Lang. Proc.*, vol. 24, no. 11, pp. 2009–2022, 2016.