

NEURAL NETWORK-BASED TIME-FREQUENCY-BIN-WISE LINEAR COMBINATION OF BEAMFORMERS FOR UNDERDETERMINED TARGET SOURCE EXTRACTION

Changda Chen¹, Yichen Yang^{1,2}, Wei Liu^{1,3}, and Shoji Makino¹

¹Waseda University, Japan

²Northwestern Polytechnical University, Xi'an, China

³School of Electronic Information, Wuhan University, Wuhan, China

ABSTRACT

Extracting a target source from underdetermined mixtures is challenging for beamforming approaches. Recently proposed time-frequency-bin-wise switching (TFS) and linear combination (TFLC) strategies mitigate this by combining multiple beamformers in each time-frequency (TF) bin and choosing combination weights that minimize the output power. However, making this decision independently for each TF bin can weaken temporal-spectral coherence, causing discontinuities and consequently degrading extraction performance. In this paper, we propose a novel neural network-based time-frequency-bin-wise linear combination (NN-TFLC) framework that constructs minimum power distortionless response (MPDR) beamformers without explicit noise covariance estimation. The network encodes the mixture and beamformer outputs, and predicts temporally and spectrally coherent linear combination weights via a cross-attention mechanism. On dual-microphone mixtures with multiple interferers, NN-TFLC-MPDR consistently outperforms TFS/TFLC-MPDR and achieves competitive performance with TFS/TFLC built on the minimum variance distortionless response (MVDR) beamformers that require noise priors.

Index Terms— Target source extraction, underdetermined situations, beamforming, linear combination, neural networks

1. INTRODUCTION

In acoustic signal processing, multi-microphone arrays provide spatial diversity for target source extraction (TSE) from observed signals [1–3]. Classical spatial filters such as the minimum power distortionless response (MPDR) and minimum variance distortionless response (MVDR) beamformers [4, 5], utilize array geometry to suppress interference while preserving the source signal from the target direction of arrival (DOA). Recently, neural network-based beamforming has advanced along several primary directions: estimating time-frequency (TF) masks to derive the noise covariance [6–8]; predicting the target relative transfer function (RTF) to strengthen MPDR or MVDR beamformers [9, 10]; and directly learning beamformer weights [11]. In parallel, time-domain methods such as FaS-Net achieve learned beamforming on the waveform for lower latency [12].

However, in underdetermined situations, where the number of active sources exceeds the number of microphones, the performance of beamforming methods severely degrades. A recently proposed strategy is to exploit TF sparsity by constructing multiple candidate beamformers that share the same target steering while placing nulls at different DOAs of the interferers. Time-frequency-bin-wise switching (TFS) [13, 14] makes a hard choice, selecting one beam-

former in each TF bin, whereas time-frequency-bin-wise linear combination (TFLC) [15] uses soft weights to combine multiple beamformers. Both TFS and TFLC choose combination weights that minimize the output power (min-selection) in each TF bin under a distortionless constraint and can be extended to MPDR or MVDR [14]. Here, MVDR typically requires prior information of the interferers to make the combination directly minimize their power, which is rarely attainable in practice. MPDR removes this requirement, but minimizing observed power risks suppressing the target. Moreover, independent per-bin selection weakens TF correlation, leading to spectral discontinuities and consequently poorer extraction performance. By contrast, neural networks can aggregate information over frequency bins within each frame, yielding more consistent weights and a better trade-off between target leakage and interference reduction. Existing neural beam combination approaches usually fix a set of pre-designed beamformers and learn a global selection at the frame level or spectrogram level [16, 17] to perform source separation, which cannot fully exploit TF sparsity and generally relies on the beamformers designed in overdetermined situations.

To this end, we propose a neural network-based time-frequency-bin-wise linear combination (NN-TFLC) framework that encodes the mixture and each beamformed signal with inplace convolutional gated linear unit (ICGLU) [18, 19] and frequency-independent Bi-LSTM to preserve the full TF resolution. A cross-attention module then consumes mixture and beam features to produce combination weights with temporal and spectral context in each TF bin. We use the prior RTF to form phase cues and to update several candidate MPDR beamformers without explicitly estimating noise covariances. The pipeline supports a variable number of input beamformers and requires only a single update step. In experiments, NN-TFLC-MPDR consistently outperforms TFS/TFLC-MPDR and achieves competitive performance with TFS/TFLC-MVDR in dual-microphone scenarios where multiple interferers coexist.

2. TFS AND TFLC OF BEAMFORMERS

Consider an M -element microphone array observing one target source, $N - 1$ interferers, and additive background noise. The observation in the short-time Fourier transform (STFT) domain is

$$\mathbf{x}_{f,t} = \mathbf{a}_f S_{\text{ref},f,t} + \sum_{n=1}^{N-1} \mathbf{h}_{n,f} I_{n,f,t} + \mathbf{n}_{f,t}, \quad (1)$$

where $f = 1, \dots, F$ and $t = 1, \dots, T$ are the indices of the frequency bins and time frames, respectively, $\mathbf{x}_{f,t} \in \mathbb{C}^M$ contains the observed signals, $\mathbf{a}_f \in \mathbb{C}^M$ is the RTF of the target source signal, and $S_{\text{ref},f,t} \in \mathbb{C}$ is the target source image at the reference microphone, $I_{n,f,t} \in \mathbb{C}$ and $\mathbf{h}_{n,f} \in \mathbb{C}^M$ denote the signal of the n -th

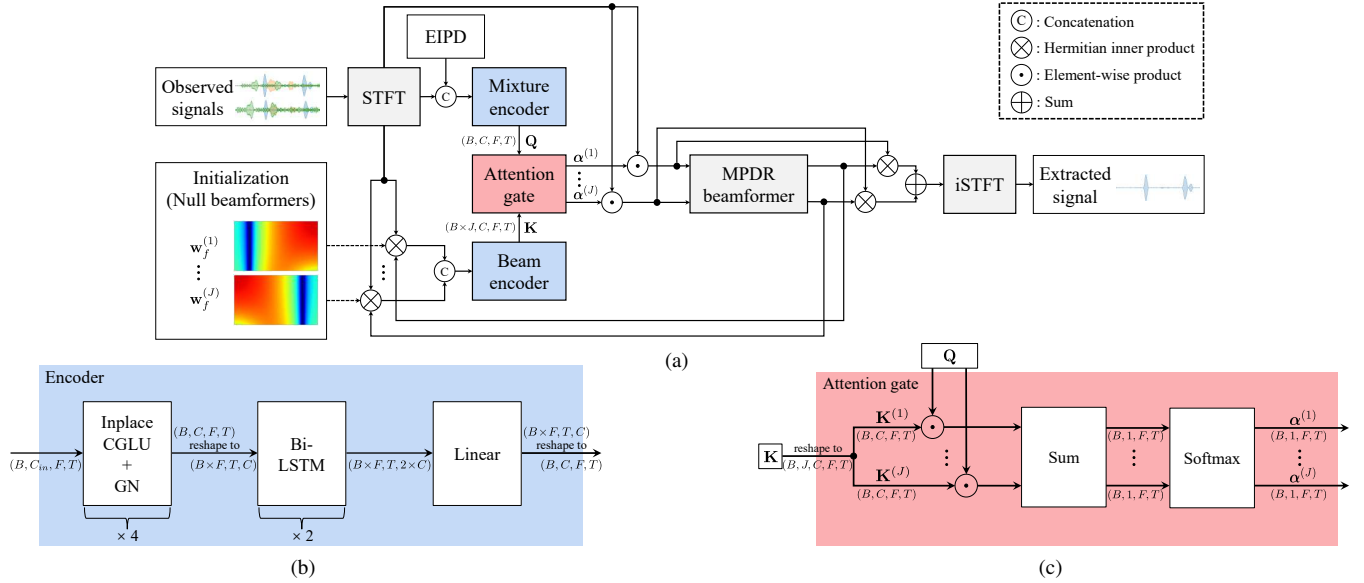


Fig. 1. Overview and detailed components of the proposed framework. (a) Overall structure of NN-TFLC-MPDR; (b) Structure of mixture encoder and beam encoder; (c) Structure of attention gate.

interferer and its corresponding transfer function, respectively, and $\mathbf{n}_{f,t} \in \mathbb{C}^M$ is the additive background noise. Our goal is to estimate $S_{\text{ref},f,t}$.

Theoretically, an M -dimensional time-invariant beamformer can impose at most M independent linear constraints (one unit response plus $M - 1$ nulls), so when more than $M - 1$ interferers are simultaneously active, conventional beamforming cannot fully suppress them. To address this limitation, we consider combining different beamformers to process the observed signals.

Under the M -disjoint orthogonality assumption mentioned in [15], at most $M - 1$ sources are active in each TF bin. In the dual-microphone scenarios, this assumption reduces to W -disjoint orthogonality [20, 21]. Therefore, we construct J ($\geq N - 1$) candidate beamformers, each suppressing one interferer, and estimate the target via a combination in each TF bin, expressed as

$$\hat{S}_{\text{ref},f,t} = \sum_{j=1}^J \alpha_{f,t}^{(j)} (\mathbf{w}_f^{(j)})^H \mathbf{x}_{f,t}, \quad (2)$$

where $\alpha_{f,t}^{(j)}$ is the corresponding linear weight of the j -th beamformer $\mathbf{w}_f^{(j)}$, satisfying $0 \leq \alpha_{f,t}^{(j)} \leq 1$ and $\sum_{j=1}^J \alpha_{f,t}^{(j)} = 1$, and $(\cdot)^H$ denotes the Hermitian transpose. Estimating the linear weights in each TF bin reduces to output-power minimization: the hard case $\alpha_{f,t}^{(j)} \in \{0, 1\}$ yields TFS [13], whereas the soft case $\alpha_{f,t}^{(j)} \in [0, 1]$ yields TFLC [15].

In practice, it is difficult to directly obtain predesigned beamformers whose nulls are precisely aligned with the DOAs of the interferers. Therefore, prior work [14, 15] adopts an iterative update strategy, where beamformers initialized with random values or with nulls placed in random DOAs are used as the starting points. Then the linear weights $\alpha_{f,t}^{(j)}$ are updated, and the updated weights are applied as masks to derive new beamformers, as

$$\mathbf{w}_f^{(j)} = \frac{(\Phi_f^{(j)})^{-1} \mathbf{a}_f}{\mathbf{a}_f^H (\Phi_f^{(j)})^{-1} \mathbf{a}_f} \quad \text{for } 1 \leq j \leq J, \quad (3)$$

where $\Phi_f^{(j)} = \mathbb{E}[(\alpha_{f,t}^{(j)} \mathbf{x}_{f,t})(\alpha_{f,t}^{(j)} \mathbf{x}_{f,t})^H]$ is the covariance matrix for the j -th beamformer. Here, when $\mathbf{x}_{f,t}$ is the full mixture including

the target, this update corresponds to MPDR, and when $\mathbf{x}_{f,t}$ contains only the noise components, this update corresponds to MVDR.

3. PROPOSED METHOD

3.1. An overview of the proposed framework

The overall structure of the proposed NN-TFLC-MPDR is illustrated in Fig. 1. Given dual-microphone observed signals, the waveforms are transformed to the STFT domain. The cosine and sine of the expected inter-channel phase difference (EIPD) [22] derived from the RTF serve as the target phase cues, which are concatenated with the real and imaginary parts of the mixture along the channel dimension and fed into a mixture encoder. We prepare J null beamformers $\{\mathbf{w}_f^{(j)}\}_{j=1}^J$ with randomly selected null DOAs $\{\theta^{(j)}\}_{j=1}^J$ to provide initial directional information. Likewise, the real and imaginary parts of their beamformed signals are concatenated along the channel dimension and encoded by a shared beam encoder.

An attention gate takes the outputs of the mixture encoder and beam encoder as its inputs, and yields TF-bin-wise linear weights. We adopt a single iteration, in which the weights are used once to refine the MPDR beamformers. The updated candidate beamformed signals are re-encoded, and the same attention gate then yields the final combination weights, which are used to synthesize the estimated target source signal.

3.2. Mixture encoder and Beam encoder

Each initial null beamformer is constructed with two linear constraints: a unit response toward the target DOA, and a null toward a DOA sampled from the angular range of the interferers. In the dual-microphone scenarios, this construction is determined.

Both the mixture encoder and beam encoder share the same architecture in Fig. 1b. The encoder begins with four ICGLU blocks [18], whose stride is set to 1 so that the TF resolution does not change. This preserves an alignment with the later prediction of TF-bin-wise linear weights. In the beam encoder, we stack J beamformed signals along the batch dimension to share weights, enabling

the model to naturally support a variable number of beamformer inputs. In addition, each ICGLU block is followed by a batch-size-agnostic group normalization (GN) [23], which is more robust than batch normalization due to the distribution shifts induced by different beam patterns. After GN, the exponential linear unit (ELU) acts as the activation function.

Temporal context is then modeled by a two-layer frequency-independent Bi-LSTM shared across frequency bins to consistently capture the time-delay characteristics over frequencies. A final linear layer halves the channel dimension, and the features are reshaped back to the original TF structure before entering the attention gate.

3.3. Attention gate

In each TF bin, a cross-attention mechanism is employed. The linear weights of every candidate beamformer are obtained by the attention gate using softmax on a scaled dot product, expressed as

$$\alpha_{f,t}^{(j)} = \frac{\exp(\mathbf{Q}_{f,t}^T \mathbf{K}_{f,t}^{(j)} / \sqrt{C})}{\sum_{j'=1}^J \exp(\mathbf{Q}_{f,t}^T \mathbf{K}_{f,t}^{(j')} / \sqrt{C})} \text{ for } 1 \leq j \leq J, \quad (4)$$

where $\mathbf{Q}_{f,t} \in \mathbb{R}^C$ is the query from the mixture encoder, $\mathbf{K}_{f,t}^{(j)} \in \mathbb{R}^C$ is the j -th key from the beam encoder, $(\cdot)^T$ denotes the transpose, and C is the encoded channel dimension.

The linear weights $\alpha_{f,t}^{(j)}$ first serve as masks to form masked covariances and update J MPDR beamformers via (3), where $\mathbf{x}_{f,t}$ includes the target. Then, the refined beamformers are fed back through the shared beam encoder, re-encoded, and passed through the attention gate to obtain the final combination weights, which yield the target estimate in (2), followed by inverse STFT (iSTFT).

3.4. Training strategy

We use two loss functions to train the neural network: a scale-invariant signal-to-distortion ratio (SI-SDR) [24], and an entropy regularization term.

To help the network learn to recover the target source signals from the mixture, the negative SI-SDR loss is employed, expressed as

$$\mathcal{L}_{\text{SI-SDR}} = -10 \log_{10} \left[\frac{\mathbb{E}(|\beta \mathbf{s}|^2)}{\mathbb{E}(|\beta \mathbf{s} - \hat{\mathbf{s}}|^2)} \right], \quad (5)$$

where \mathbf{s} and $\hat{\mathbf{s}}$ denote the time-domain reference and enhanced signals, respectively, via iSTFT from $S_{\text{ref},f,t}$ and $\hat{S}_{\text{ref},f,t}$, and $\beta = \hat{\mathbf{s}}^T \mathbf{s} / \|\mathbf{s}\|_2^2$ is the scaling factor. In the early stage of training, we observed that the final estimated linear weights are nearly uniform, e.g., $\alpha_{f,t}^{(j)} \approx 1/J$, which severely stalls the training process. Therefore, an entropy regularization term is proposed, expressed as

$$\mathcal{L}_{\text{Ent}} = -\frac{1}{FTJ} \sum_{f=1}^F \sum_{t=1}^T \sum_{j=1}^J \alpha_{f,t}^{(j)} \ln(\alpha_{f,t}^{(j)} + \epsilon), \quad (6)$$

where $\epsilon > 0$ prevents taking $\ln 0$. Minimizing this term not only helps the linear weights deviate from the uniform distribution, which accelerates training and convergence, but also encourages more decisive selection, allowing the difference in covariance statistics during the update of MPDR beamformers to be amplified, thereby enhancing the complementarity among the candidate beamformers.

The final loss function is summarized as

$$\mathcal{L} = \mathcal{L}_{\text{SI-SDR}} + \lambda \mathcal{L}_{\text{Ent}}, \quad (7)$$

where λ is the weight of the entropy regularization term.

4. EXPERIMENTS

4.1. Mixing conditions

We synthesize mixtures from clean utterances of LibriSpeech [25]. For each mixture, a simulated room is generated with length sampled from [6, 10] m, width from [5, 8] m, height from [2.5, 3.5] m, and reverberation time T_{60} from [0.2, 0.5] s. A dual-microphone linear array with 2 cm spacing is randomly placed within the room at a height of 1.5 m, and its center is at least 2.5 m from any wall. The source-to-array distance is chosen from [1.5, 2.0] m, and the source height is between 1.4 m and 1.6 m. The target DOA is selected from $[80^\circ, 100^\circ]$, whereas DOAs of the interferers are limited to $[0^\circ, 65^\circ]$ and $[115^\circ, 180^\circ]$, with at most two interferers in each range. Clean signals are convolved with simulated RIRs using the image method [26]. Each interferer is then scaled to achieve an input SIR of [0, 5] dB. Additive noise comprises simulated diffuse noise [27] and white Gaussian noise. The diffuse-to-white power ratio is in [15, 25] dB, and the overall additive noise is adjusted to yield an SNR in [10, 25] dB with respect to the reverberant target. All the above random choices follow a uniform distribution.

In total, we create 25,000 mixtures for training, 3,000 for validation, and 3,000 for testing, each 6-second-long. The training set contains 15,000 mixtures with two interferers (2I), 5,000 with three (3I), and 5,000 with four (4I). The validation and test sets have a similar composition with 2,000 (2I), 500 (3I), and 500 (4I).

4.2. Model configurations

In each ICGLU block, the kernel size is set as (5×1) , and the number of encoded channels is set to $C = 32$. All audio is sampled at 16 kHz, with a 1024-sample Hanning window and a 256-sample shift. The weight of the entropy regularization term is $\lambda = 0.05$. The Adam optimizer [28] is employed, together with a StepLR scheduler that reduces the learning rate by a factor of 0.8 every 10 epochs. The batch size is set to $B = 4$.

In our proposed method, we evaluate two model variants. **NN-TFLC-MPDR (w/o Full)** is trained only on the 2I subset with an initial learning rate of 6×10^{-4} , and its inputs are two beamformers whose null DOAs are randomly drawn from $[10^\circ, 55^\circ]$ and $[125^\circ, 170^\circ]$ respectively, to augment the directional diversity of the training data. The best model is chosen by the average SI-SDR on the 2I validation split over 100 epochs. **NN-TFLC-MPDR (w/ Full)** is initialized from the 2I model and then continues training on the 3I and 4I subsets, with an initial learning rate of 2×10^{-4} , using four input beamformers with nulls randomly drawn from $[10^\circ, 30^\circ]$, $[35^\circ, 55^\circ]$, $[125^\circ, 145^\circ]$, and $[150^\circ, 170^\circ]$. Its best model is selected by the average SI-SDR on the entire validation set over another 100 epochs.

For validation and testing, we fix the null angles. In the 2I case, the nulls of two initial beamformers are set to 32.5° and 147.5° . In the 3I and 4I cases, we use 16.25° , 48.75° , 131.25° , and 163.75° .

4.3. Results and discussions

For comparison, MVDR, TFS-MVDR [14], and TFLC-MVDR [15], TFS-MPDR [14], and TFLC-MPDR [15] are used as the baselines. All methods, including ours, share the same RTF, estimated as the principal eigenvector of the covariance matrix of the reverberant target source signals. For MVDR-based methods, the reverberant clean signals of interferers act as additional priors, and the min-selection rule is applied directly to minimize the interference and build the noise covariance for each beamformer, which is considered to have

Table 1. Average SI-SDR (dB), SI-SIR (dB), and PESQ scores for processed signals with 2/3/4 interferers (mean±standard deviation).

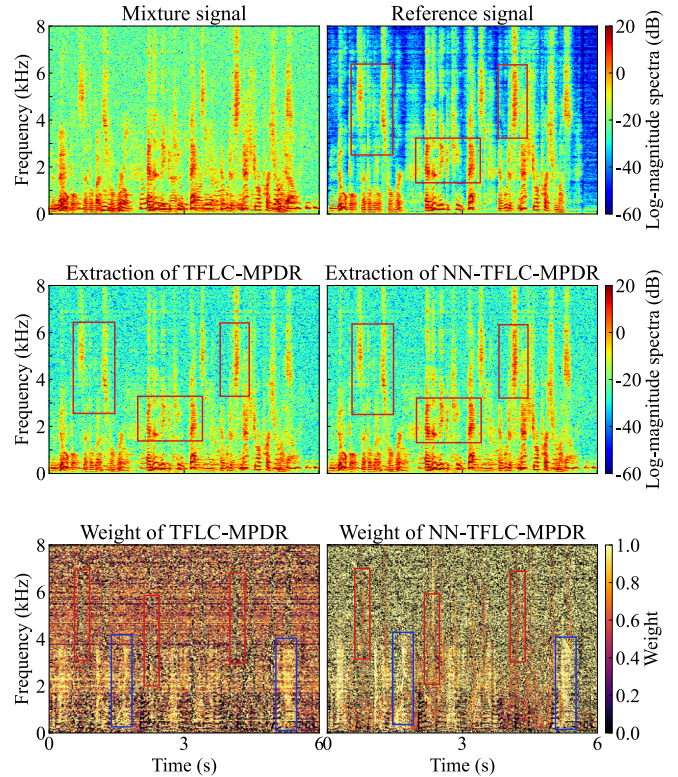
Method	2I (2 beamformers)			3I (4 beamformers)			4I (4 beamformers)		
	SI-SDR	SI-SIR	PESQ	SI-SDR	SI-SIR	PESQ	SI-SDR	SI-SIR	PESQ
Unprocessed	-0.81±1.00	-0.69±1.04	1.13±0.07	-2.48±0.82	-2.40±0.83	1.09±0.07	-3.88±0.77	-3.81±0.78	1.09±0.09
MVDR	0.93±1.07	2.61±1.40	1.16±0.08	-0.95±0.90	0.33±1.09	1.11±0.08	-2.46±0.81	-1.35±0.95	1.09±0.06
TFS-MVDR [14]	4.16±1.38	8.35±2.16	1.24±0.12	3.98±1.29	8.62±1.81	1.22±0.13	2.84±1.09	6.88±1.59	1.16±0.08
TFLC-MVDR [15]	4.52±1.43	8.04±2.02	1.25±0.13	4.54±1.32	7.87±1.61	1.23±0.14	3.37±1.13	6.16±1.40	1.17±0.08
TFS-MPDR [14]	2.45±1.51	6.06±2.36	1.20±0.10	0.03±1.55	4.04±2.21	1.13±0.08	-0.51±1.42	3.09±1.98	1.10±0.05
TFLC-MPDR [15]	2.86±1.55	5.56±2.12	1.21±0.10	1.31±1.58	3.82±1.97	1.14±0.09	0.32±1.39	2.53±1.72	1.11±0.05
NN-TFLC-MPDR (w/o Full)	4.80±1.55	7.70±2.19	1.28±0.12	3.19±1.44	5.85±1.93	1.20±0.11	1.27±1.31	3.67±1.73	1.14±0.06
NN-TFLC-MPDR (w/ Full)	4.51±1.52	7.00±2.06	1.26±0.12	4.71±1.54	9.82±2.55	1.26±0.13	2.65±1.52	7.08±2.42	1.17±0.07

oracle performance in the original framework. All TFS/TFLC baselines run for five iterations. After processing, we calculate their SI-SDR, scale-invariant signal-to-interference ratio (SI-SIR) [29], and perceptual evaluation of speech quality (PESQ) scores [30] as the objective metrics. The first microphone of the array is set as the reference microphone, and the reverberant clean source image of the target is set as the reference signal for all the above metrics.

First, we compare the baselines with the proposed NN-TFLC-MPDR in the 2I, 3I, and 4I cases of the test set, respectively. As shown in Table 1, the performance of a single MVDR beamformer degrades sharply in underdetermined scenes. In contrast, our NN-TFLC-MPDR consistently outperforms TFS/TFLC MPDR across 2I/3I/4I. Although NN-TFLC-MPDR (w/o Full) is trained only on 2I with two initial beamformers as inputs, it can still yield around 1 dB higher SI-SDR than TFS/TFLC-MPDR when more interferers exist in the environment, which demonstrates the extensibility of our model to adapt to varying numbers of input beamformers. NN-TFLC-MPDR (w/ Full) undergoes the full training process, and can then have substantial improvements in both SI-SDR and SI-SIR in 3I/4I, with a minor performance loss in 2I. Importantly, our method can match or surpass TFS/TFLC-MVDR in 2I/3I and yields a SI-SDR gap of less than 1 dB in 4I without requiring the prior information of the interferers, while achieving comparable PESQ scores.

Next, we visualize the spectrograms of estimated signals and the beamformer combination weights in the TF plane processed by TFLC-MPDR and NN-TFLC-MPDR, respectively, on a sample from the 2I case. In Fig. 2, the bright bands of TFLC-MPDR appear dimmer and more fragmented than the target's, as highlighted by the red boxes. This indicates that MPDR with per-bin min-selection does not fully preserve the target under the distortionless constraint, leading to potential suppression. While the bright bands of NN-TFLC-MPDR nearly coincide with those of the reference, which demonstrates the NN-based selection reduces target distortion and achieves a better trade-off between target preservation and suppression of interference.

We further plot the combination weights for one of the two candidate beamformers in Fig. 2. It is obvious that per-bin min-selection of TFLC yields fragmented weight maps with weak cross-bin correlation. The discontinuities make the weights noisy and irregular across TF bins, potentially introducing phase inconsistencies and artifacts. By contrast, the NN-derived weights are coherent over time and frequency, and transition smoothly, especially in low-frequency regions and along several frequency bands. This behavior arises because attention in each TF bin operates on features that already encode TF context. As highlighted by the red boxes, in these target-dominant regions, the network tends to blend two beamformers rather than producing the patchwork mixing seen in TFLC-MPDR. This reduces the risk of directly selecting a beamformer whose null lies near the target, and thus helps preserve the target component. Also, in interference-dominant regions, shown by the blue boxes,

**Fig. 2.** Log-magnitude spectrogram and combination weights of a beamformer in the TF plane for a sample from the 2I case.

the network makes more decisive selections, including at high frequencies where background noise is prominent, which strengthens beam selectivity and improves suppression of different interferers.

5. CONCLUSIONS

In this paper, we present an NN-TFLC framework that constructs MPDR beamformers. It uses a neural network to estimate TF-bin-wise linear weights for combining multiple beamformers to extract the target source from underdetermined mixtures, without explicit noise covariance estimation. Across mixtures with multiple interferers, our method consistently outperforms conventional TFS/TFLC-MPDR and achieves competitive performance with TFS/TFLC-MVDR, which leverages prior information of the interferers to directly minimize their power. Moreover, experiments show that the proposed network adapts seamlessly to different numbers of input beamformers, demonstrating strong scalability and practicality.

6. REFERENCES

- [1] J. Benesty, S. Makino, and J. Chen, *Speech enhancement*, Springer, 2006.
- [2] J. Benesty, I. Cohen, and J. Chen, *Fundamentals of signal enhancement and array signal processing*, John Wiley & Sons, 2017.
- [3] S. Makino, *Audio Source Separation*, Springer, Switzerland, 2018.
- [4] J. Li, P. Stoica, and Z. Wang, “On robust capon beamforming and diagonal loading,” *IEEE Trans. Signal Process.*, vol. 51, no. 7, pp. 1702–1715, Jul. 2003.
- [5] J. Benesty, J. Chen, and Y. Huang, *Microphone array signal processing*, Springer, 2008.
- [6] J. Heymann, L. Drude, and R. Haeb-Umbach, “Neural network based spectral mask estimation for acoustic beamforming,” in *Proc. IEEE ICASSP*, 2016, pp. 196–200.
- [7] 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. Interspeech*, 2016, pp. 1981–1985.
- [8] X. Xiao, S. Zhao, D. L. Jones, E. S. Chng, and H. Li, “On time-frequency mask estimation for MVDR beamforming with application in robust speech recognition,” in *Proc. IEEE ICASSP*, 2017, pp. 3246–3250.
- [9] A. Aroudi and S. Braun, “DBNet: DOA-driven beamforming network for end-to-end reverberant sound source separation,” in *Proc. IEEE ICASSP*, 2021, pp. 211–215.
- [10] Y. Yang, C. Pan, Q. Gao, J. Benesty, S. Makino, and J. Chen, “Neural network-assisted joint DOA estimation and beamforming with first-order reflection modeling,” in *Proc. APSIPA ASC*, 2025, pp. 19–23.
- [11] Z. Zhang, Y. Xu, M. Yu, S.-X. Zhang, L. Chen, and D. Yu, “ADL-MVDR: All deep learning MVDR beamformer for target speech separation,” in *Proc. IEEE ICASSP*, 2021, pp. 6089–6093.
- [12] 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.
- [13] K. Yamaoka, A. Brendel, N. Ono, S. Makino, M. Buerger, T. Yamada, and W. Kellermann, “Time-frequency-bin-wise beamformer selection and masking for speech enhancement in underdetermined noisy scenarios,” in *Proc. EUSIPCO*, 2018, pp. 1582–1586.
- [14] K. Yamaoka, N. Ono, S. Makino, and T. Yamada, “Time-frequency-bin-wise switching of minimum variance distortionless response beamformer for underdetermined situations,” in *Proc. IEEE ICASSP*, 2019, pp. 7908–7912.
- [15] K. Yamaoka, N. Ono, and S. Makino, “Time-frequency-bin-wise linear combination of beamformers for distortionless signal enhancement,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3461–3475, Nov. 2021.
- [16] Z. Chen, T. Yoshioka, X. Xiao, L. Li, M. L. Seltzer, and Y. Gong, “Efficient integration of fixed beamformers and speech separation networks for multi-channel far-field speech separation,” in *Proc. IEEE ICASSP*, 2018, pp. 5384–5388.
- [17] R. Liu, Y. Zhou, H. Liu, X. Xu, J. Jia, and B. Chen, “A new neural beamformer for multi-channel speech separation,” *J. Signal Process. Syst.*, vol. 94, no. 10, pp. 977–987, May 2022.
- [18] J. Liu and X. Zhang, “Inplace gated convolutional recurrent neural network for dual-channel speech enhancement,” in *Proc. Interspeech*, 2021, pp. 1852–1856.
- [19] J. Bai, H. Li, X. Zhang, and F. Chen, “Attention-based beamformer for multi-channel speech enhancement,” in *Proc. IEEE ICASSP*, 2025, pp. 1–5.
- [20] A. Jourjine, S. Rickard, and O. Yilmaz, “Blind separation of disjoint orthogonal signals: Demixing N sources from 2 mixtures,” in *Proc. IEEE ICASSP*, 2000, vol. 5, pp. 2985–2988.
- [21] O. Yilmaz and S. Rickard, “Blind separation of speech mixtures via time-frequency masking,” *IEEE Trans. Signal Process.*, vol. 52, no. 7, pp. 1830–1847, Jul. 2004.
- [22] Z. Wang and D. Wang, “On spatial features for supervised speech separation and its application to beamforming and robust ASR,” in *Proc. IEEE ICASSP*, 2018, pp. 5709–5713.
- [23] Y. Wu and K. He, “Group normalization,” in *Proc. ECCV*, 2018, pp. 3–19.
- [24] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR—half-baked or well done?,” in *Proc. IEEE ICASSP*, 2019, pp. 626–630.
- [25] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *Proc. IEEE ICASSP*, 2015, pp. 5206–5210.
- [26] J. B. Allen and D. A. Berkley, “Image method for efficiently simulating small-room acoustics,” *J. Acoust. Soc. Am.*, vol. 65, no. 4, pp. 943–950, Apr. 1979.
- [27] E. A. P. Habets, I. Cohen, and S. Gannot, “Generating non-stationary multisensor signals under a spatial coherence constraint,” *J. Acoust. Soc. Am.*, vol. 124, no. 5, pp. 2911–2917, Nov. 2008.
- [28] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [29] R. Scheibler, “SDR—medium rare with fast computations,” in *Proc. IEEE ICASSP*, 2022, pp. 701–705.
- [30] J. G. Beerends, A. P. Hekstra, A. W. Rix, M. P. Hollier, et al., “Perceptual evaluation of speech quality (PESQ): The new ITU standard for end-to-end speech quality assessment part ii: psychoacoustic model,” *J. Audio Eng. Soc.*, vol. 50, no. 10, pp. 765–778, Oct. 2002.