



Spatial Loss for Unsupervised Multi-channel Source Separation

Kohei Saijo^{1,2*}, Robin Scheibler²

¹Waseda University, Japan ²LINE Corporation, Japan

saijo@pcl.cs.waseda.ac.jp

Abstract

We propose a *spatial loss* for unsupervised multi-channel source separation. The proposed loss exploits the duality of direction of arrival (DOA) and beamforming: the steering and beamforming vectors should be aligned for the target source, but orthogonal for interfering ones. The spatial loss encourages consistency between the mixing and demixing systems from a classic DOA estimator and a neural separator, respectively. With the proposed loss, we train the neural separators based on minimum variance distortionless response (MVDR) beamforming and independent vector analysis (IVA). We also investigate the effectiveness of combining our spatial loss and a *signal loss*, which uses the outputs of blind source separation as the references. We evaluate our proposed method on synthetic and recorded (LibriCSS) mixtures. We find that the spatial loss is most effective to train IVA-based separators. For the neural MVDR beamformer, it performs best when combined with a signal loss. On synthetic mixtures, the proposed unsupervised loss leads to the same performance as a supervised loss in terms of word error rate. On LibriCSS, we obtain close to state-of-the-art performance without any labeled training data.

Index Terms: unsupervised learning, spatial loss, direction-of-arrival, beamforming, blind source separation

1. Introduction

Source separation has been used as front-ends of speech systems to recover sources corrupted by interference and background noise. On the one hand, blind source separation (BSS) such as independent component analysis (ICA) [1], independent vector analysis (IVA) [2, 3], and independent low-rank matrix analysis (ILRMA) [4] have been an area of intense research. On the other hand, supervised learning of deep neural networks (DNNs) for single-channel source separation has been eagerly investigated [5–7]. In the multi-channel setup, methods that exploit a single-channel separation network to estimate spatial cues such as DNN-based minimum variance distortionless response (MVDR) beamforming [8–10] or BSS with a neural source model [11–13] has led to stunning improvements in performance. Such linear filtering has also been shown empirically to be better front-ends for ASR than non-linear ones [14]. However, supervised learning requires access to the massive amount of mixtures and corresponding ground-truth signals. Because such a dataset of natural recordings cannot be obtained, many prior works rely on simulation instead [14, 15].

Recently, unsupervised source separation with a *signal loss*, which uses the outputs of BSS as pseudo-targets instead of ground-truth clean signals, has been proposed [16, 17]. In [16], time-frequency (TF) masks estimated by a blind spatial clustering technique were used to train a deep clustering model. [17] proposed a loss function that evaluates Kullback-Leibler Divergence (KLD) between the posterior probability density function

*This work was done during an internship at LINE Corporation.

of the separated signals of BSS and that of the DNN-based separator to avoid overfitting to the errors in the BSS outputs. These unsupervised losses enforce consistency between the output signals of BSS and those of DNN-based separators.

In contrast, we propose a *spatial loss* function for unsupervised multi-channel source separation that enforces consistency of the estimated spatial parameters. We exploit the duality of direction of arrival (DOA) and beamforming: since linear separation relies on the assumption that sources are mixed linearly, the mixing and the beamforming matrix should be inverse of each other. We train a neural separator so that the beamforming vector for a given source should have a large inner product with the corresponding steering vector, while being close to orthogonal to those of the other sources. The steering vectors are obtained by conventional DOA estimation, the multiple signal classification (MUSIC) algorithm [18]. We also investigate the effectiveness of combining the spatial loss and the signal loss [10, 17]. We train two types of neural separators with our proposed loss function, auxiliary function-based IVA (Aux-IVA) [19] with a neural source model [13], and neural MVDR beamformer [10]. We evaluate our proposed method with the recorded dataset LibriCSS [20]. Synthetic mixtures are also used for a more detailed analysis that requires the groundtruth.

The key contributions are summarized as follows. 1) We propose a spatial loss function for unsupervised multi-channel source separation. It gave the best performance for IVA, and improved the performance of MVDR when combined with the signal loss. 2) We conduct extensive experiments using both synthetic mixtures and real-world recordings, whereas most prior works only considered the former. The proposed loss leads to the same performance as a supervised loss in terms of word error rate on synthetic mixtures. On LibriCSS, the proposed method outperformed strong baselines [21] and obtains close to state-of-the-art performance [22] with no data besides that available in LibriCSS, nor any ground-truth.

2. Background

Assuming N sources are captured by M microphones, the observed signal in the short-time Fourier transform (STFT) domain is represented with a mixing matrix $\mathbf{A}_f \in \mathbb{C}^{M \times N}$ as,

$$\mathbf{x}_{f,t} = \mathbf{A}_f \mathbf{s}_{f,t} + \mathbf{b}_{f,t}, \quad (1)$$

where $\mathbf{s}_{f,t}$ is the clean sources vector, $\mathbf{b}_{f,t}$ is the background noise, and $f = 1, \dots, F$ and $t = 1, \dots, T$ are the frequency bin and the time frame index. Linear source separation is the problem of estimating a demixing matrix $\mathbf{W}_f \in \mathbb{C}^{N \times M}$, that estimates the sources as,

$$\mathbf{s}_{f,t} \approx \mathbf{W}_f \mathbf{x}_{f,t}. \quad (2)$$

Thus, the optimal demixing matrix should satisfy $\mathbf{W}_f \mathbf{A}_f \approx \mathbf{I}$, where \mathbf{I} is the identity matrix. In the following, T and H denote the transpose and Hermitian transpose of vectors or matrices.




Figure 1: *Duality of beamforming and DOA. Beam is directed to target and null is directed to interference.*

2.1. DOA Estimation

Knowing the location of each microphone, $\mathbf{d}_m \in \mathbb{R}^3$, the steering vector of the n th source is,

$$\mathbf{a}_f(\mathbf{q}_n) = M^{-\frac{1}{2}} \left[e^{j\omega_f \mathbf{d}_1^\top \mathbf{q}_n}, \dots, e^{j\omega_f \mathbf{d}_M^\top \mathbf{q}_n} \right]^\top, \quad (3)$$

where $\omega_f = 2\pi \frac{f}{c}$ with sound speed c . $\mathbf{q}_n \in \mathbb{R}^3$ is a unit-length vector pointing towards the n th source, and is represented with elevation ϕ and azimuth θ as $\mathbf{q} = [\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi]^\top$. Here we consider the MUSIC method [18]. MUSIC assumes that $\mathbf{a}_f(\mathbf{q}_n)$ are orthogonal to the subspaces spanned by the $M - N$ least eigenvectors of the covariance matrix of (1), i.e., the noise subspace. Let \mathbf{E}_f be a matrix containing a basis for the noise subspace in its columns. Then, \mathbf{q}_n is estimated by finding local maxima of the following cost function,

$$\mathcal{G}(\mathbf{q}_n) = \frac{1}{F} \sum_f (|\mathbf{E}_f^H \mathbf{a}_f(\mathbf{q}_n)|^2)^{-1}. \quad (4)$$

Hereafter, we denote $\mathbf{a}_f(\mathbf{q}_n)$ as $\mathbf{a}_{n,f}$.

2.2. Source Separation

Here we mention two separation methods, AuxIVA and MVDR. **AuxIVA** estimates the demixing matrix by likelihood maximization assuming the prior distribution of sources $p(y_{n,f,t})$, where $y_{n,f,t}$ is the n th separated signal. The demixing matrix is estimated by minimizing the following negative log-likelihood,

$$\mathcal{J} = -2T \sum_f \log |\det \mathbf{W}_f| + \sum_{n,f,t} \frac{|\mathbf{w}_{n,f}^H \mathbf{x}_{f,t}|^2}{r_{n,f,t}}, \quad (5)$$

where $\mathbf{w}_{n,f}^H$ is the n th row vector of \mathbf{W}_f and $r_{n,f,t}$ the variance of the source prior $p(y_{n,f,t})$. \mathbf{W}_f can be updated with the techniques such as iterative projection [19] or iterative source steering (ISS) [23]. While conventional IVA uses a fixed source prior, replacing the source model with a trained DNN has recently demonstrated high performance [13].

MVDR computes the demixing vector that minimizes the noise variance under the distortionless constraint of the target source,

$$\mathbf{w}_{n,f} = \arg \min_{\mathbf{w}_{n,f}} \mathbf{w}_{n,f}^H \mathbf{R}_{n,f}^{(n)} \mathbf{w}_{n,f} \quad s.t. \quad \mathbf{w}_{n,f}^H \tilde{\mathbf{a}}_{n,f} = 1, \quad (6)$$

where $\mathbf{R}_{n,f}^{(n)}$ is the spatial covariance matrix of the noise. The steering vector $\tilde{\mathbf{a}}_{n,f}$ and $\mathbf{R}_{n,f}^{(n)}$ can be obtained using DNNs [8].

2.3. Unsupervised Learning using Signal Loss

Recently, learning source separation using the separated signals from a BSS method as the pseudo-target signals has been proposed [16, 17]. Let $\tilde{\mathbf{y}}$ and $\hat{\mathbf{y}}$ be the separated signal of BSS and that of DNN-based separator. The training is simply done with the loss between the two signals, which we call *signal loss*, as

$$\mathcal{L}_{\text{sig}} = \mathcal{L}(\tilde{\mathbf{y}}, \hat{\mathbf{y}}), \quad (7)$$

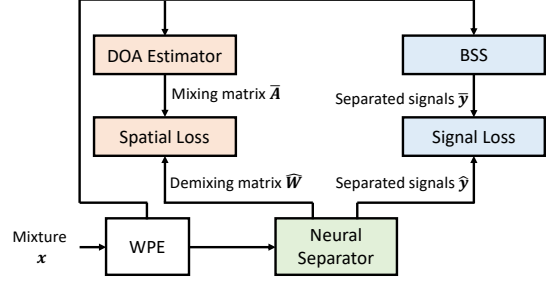


Figure 2: *Overview of our proposed method. We train a separator with blindly estimated speech signals and mixing matrices.*

where \mathcal{L} is the permutation invariant loss. Here we consider two loss functions, KLD [17] and CI-SDR [10].

KLD: Assuming that each TF bin of the source follows a time-varying complex Gaussian distribution, i.e., $y_{n,f,t} \sim \mathcal{N}(0, r_{n,f,t})$, KLD loss for a single TF bin is,

$$\mathcal{L}_{\text{kld}} = \frac{|\hat{y}_{\pi(n),f,t} - \bar{y}_{n,f,t}|^2}{\hat{r}_{\pi(n),f,t}} + \frac{\bar{r}_{n,f,t}}{\hat{r}_{\pi(n),f,t}} + \log \frac{\hat{r}_{\pi(n),f,t}}{\bar{r}_{n,f,t}} - 1, \quad (8)$$

where $\pi(n)$ denotes the optimal assignment among all the permutations to minimize the loss. The final loss is obtained by summing (8) over all TF bins and sources.

CI-SDR: Let $\tilde{\mathbf{s}}$ and $\hat{\mathbf{s}}$ be the *time-domain* separated signals of BSS and DNN. Further define a matrix containing K shifts of $\tilde{\mathbf{s}}$ in its columns, $\tilde{\mathbf{S}} = [\tilde{\mathbf{s}}_1, \dots, \tilde{\mathbf{s}}_I]^\top \in \mathbb{R}^{I \times K}$, with i th row $\tilde{\mathbf{s}}_i = [\tilde{s}_{i-1}, \dots, \tilde{s}_{i-K}]^\top \in \mathbb{R}^K$ where $i = 1, \dots, I$ is the time index. Then, let $\boldsymbol{\alpha}_n = (\tilde{\mathbf{S}}_n^\top \tilde{\mathbf{S}}_n)^{-1} \tilde{\mathbf{S}}_n^\top \hat{\mathbf{s}}_{\pi(n)}$, CI-SDR loss is,

$$\mathcal{L}_{\text{sdr}} = \sum_n -10 \log_{10} \left(\frac{\|\tilde{\mathbf{S}}_n \boldsymbol{\alpha}_n\|^2}{\|\tilde{\mathbf{S}}_n \boldsymbol{\alpha}_n - \hat{\mathbf{s}}_{\pi(n)}\|^2} \right). \quad (9)$$

3. Proposed Method

We propose a *spatial loss* function for unsupervised multi-channel source separation. We exploit the duality of DOA and beamforming. The beamforming vector for a given source should have a large inner product with the corresponding steering vector, while being close to orthogonal to those of competing sources. DOAs are estimated by a conventional MUSIC-based technique [24]. Different from the signal loss described in Sec.2.3, our proposed method explicitly imposes the spatial constraints on the demixing matrix.

3.1. Spatial Loss Function

We propose a spatial loss function that exploits DOAs estimated by MUSIC as the pseudo-targets. As shown in Fig.1, there is a duality between beamforming and DOA, i.e., the beam should be directed to the target source direction and the null to the interference direction. Let $\bar{\mathbf{A}}_f$ be the mixing matrix computed with the estimated DOAs and $\hat{\mathbf{W}}_f$ be the demixing matrix estimated by the neural separator. Our proposed loss function is,

$$\mathcal{L}_{\text{doa}} = \min_{\boldsymbol{\Pi}} \sum_f \mathbf{1}^\top (|\boldsymbol{\Pi} - |\hat{\mathbf{W}}_f \bar{\mathbf{A}}_f|) \mathbf{1}, \quad (10)$$

where $|\hat{\mathbf{W}}_f \bar{\mathbf{A}}_f|$ is normalized so that all the elements range from 0 to 1. $\mathbf{1}$ is a vector whose all the elements are one and $\mathbf{1}^\top \mathbf{X} \mathbf{1}$ computes the sum of all the elements of a matrix \mathbf{X} . $\boldsymbol{\Pi}$ denotes a permutation matrix, where only one element in each row and each column is 1 and the others are 0. For example, if $N = 2$, $\boldsymbol{\Pi}$ is either $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ or $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Here we consider two options of normalizing $|\hat{\mathbf{W}}_f \bar{\mathbf{A}}_f|$. One is to normalize rows and

Table 1: Average SDR in decibels, and STOI, PESQ, and WER of separated signals of IVA from the test set of WSJ-mix dataset. Training is done with 3 channels, and evaluation is done with 2, 3 and 6 channels. The proposed method is indicated by a star (*).

Algo.	Loss	2 channels				3 channels				6 channels			
		SDR [↑]	STOI [↑]	PESQ [↑]	WER [↓]	SDR [↑]	STOI [↑]	PESQ [↑]	WER [↓]	SDR [↑]	STOI [↑]	PESQ [↑]	WER [↓]
Unprocess.	-	-0.4	0.738	1.213	114.9%	-0.4	0.738	1.213	114.9%	-0.4	0.738	1.213	114.9%
Gauss-IVA	-	8.8	0.865	1.630	40.6%	12.5	0.907	1.920	26.3%	16.8	0.937	2.340	21.2%
DNN-IVA	KLD	9.8	0.877	1.666	36.0%	13.4	0.917	1.967	22.5%	16.9	0.941	2.292	19.0%
DNN-IVA	SDR	10.5	0.886	1.701	33.6%	14.2	0.923	2.016	21.0%	18.0	0.950	2.408	18.4%
DNN-IVA*	DOA ₁	10.4	0.888	1.720	31.5%	14.2	0.927	2.056	19.0%	19.4	0.959	2.596	12.7%
DNN-IVA*	DOA ₂	11.2	0.899	1.737	28.7%	14.4	0.931	2.050	17.0%	18.0	0.957	2.471	13.6%
DNN-IVA*	KLD+DOA ₂	11.0	0.894	1.742	27.9%	14.8	0.930	2.091	17.3%	18.7	0.954	2.522	14.2%
DNN-IVA*	SDR+DOA ₂	10.7	0.890	1.723	30.0%	14.5	0.927	2.056	17.5%	18.6	0.954	2.494	13.5%
DNN-IVA	Sup. SDR	11.4	0.898	1.765	29.3%	15.6	0.934	2.133	17.2%	20.4	0.960	2.621	13.1%
DNN-IVA	Sup. DOA ₂	11.2	0.899	1.742	28.6%	14.4	0.931	2.049	17.6%	18.0	0.957	2.475	13.2%

columns of $|\hat{W}_f|$ and $|\bar{A}_f|$, respectively, and the other is to normalize the rows of $|\hat{W}_f \bar{A}_f|$. We denote the former as DOA₁ loss and the latter as DOA₂ loss, respectively.

3.2. DOA Estimator

We estimate DOAs using the MUSIC-based method [24]. When all the sources are highly active in a mixture, we simply obtain q_n from the entire signal x with length T . However, when some sources are not very active, we found out that such way of estimation often missed the DOAs of the less-active ones. In such case, we obtain DOAs at regular intervals with sliding window and cluster them. The detailed procedure is as follows.

1. Obtain N DOAs for each window of length L and shift S .
2. Group them into $K \geq N$ clusters C_1, \dots, C_K with E_1, \dots, E_K elements by two-dimensional k-means on elevation ϕ and azimuth θ . The centroid of C_k , (ϕ_k, θ_k) , represents the DOA of the k th source.
3. Remove C_k if E_k is less than a threshold E_{thres} .
4. Remove the cluster with the lower number of elements that satisfies $|\theta_k - \theta_{k'}| < \theta_{thres}$ ($1 \leq k < k' \leq K$).
5. Select top- N DOAs with the largest E_k .

Because sudden estimation errors occur at some time frames and they have large impact on clustering, we set K more than N in step 2, and remove such outliers in step 3. In step 4, because the clusters with close centroids are considered to belong to the same source, we remove it. Finally, we obtain $N' \leq N$ DOAs from the centroids of remaining cluster. Since we define the spatial loss assuming the existence of N sources, mixtures with $N' < N$ are not used for training.

The spatial loss requires knowledge of the microphone locations to obtain DOAs. It is thus applicable to datasets where most samples are recorded with a few known devices. We believe this to be reasonable. For most applications of interest, the geometry of the microphone array is known, e.g., smart speakers and conferencing systems. If it is unknown, but sufficient recordings from the same device are available, it can be estimated using blind calibration techniques [25, 26].

3.3. Neural Separator

As in Fig. 2, we first apply weighted prediction error (WPE) [27], and then the dereverberated mixture is separated by the neural separator. We consider two neural separators, AuxIVA with a neural source model (DNN-IVA) [13] and DNN-based MVDR beamforming (DNN-MVDR) [8–10]. In DNN-IVA, the DNN estimates $1/r_{n,f,t}$ in (5) in the form of a TF mask, and spatial model update is done with ISS [23]. The network is composed of three gate linear units [28] and a transpose

Table 2: Average SDR in decibels, and STOI, PESQ, and WER of separated signals of MVDR from the test set of WSJ-mix dataset. The proposed method is indicated by a star (*).

Ch.Algo.	Loss	SDR [↑]	STOI [↑]	PESQ [↑]	WER [↓]
6 DNN-MVDR	SDR	17.1	0.944	2.333	17.1%
DNN-MVDR*	DOA ₂	12.8	0.906	2.033	25.9%
DNN-MVDR*	SDR+DOA ₂	17.1	0.945	2.337	16.6%
DNN-MVDR	Sup.SDR	19.0	0.952	2.512	15.4%

convolution layer as in [13]. From [13], the size of the intermediate feature is changed to 256, and group normalization [29] with four groups is used instead of batch normalization. In DNN-MVDR, the DNN estimates TF masks of the target source and the noise for each source, which are used to compute the demixing matrix. The network consists of three bi-directional long short-term memory layers and two feed forward layers as in [10], where we changed the number of units to 512.

4. Experiments

4.1. Datasets and Experimental Setup

WSJ-mix: We used six channel synthetic mixtures to evaluate the separation performance with speech metrics using the groundtruth. It consisted of speech from the WSJ0 [30] and WSJ1 [31], and noise from the CHIME3 dataset [32] sampled in 16 kHz. The reverberation times were chosen randomly from 200 ms to 600 ms. The number of sources was two, with relative power from -5 dB to 5 dB. The noise was scaled to attain an SNR between 10 dB to 30 dB. Training, validation, and test sets contained 37416, 503, and 333 mixtures, approximately 98.5, 1.33 and 0.85 hours of mixtures, respectively. To evaluate WER, we trained an ASR system with clean anechoic signals using the `wsj/asr1` recipe from the ESPNet framework [33]. WER for the clean, anechoic test set was 9.25 %.

During training, the batch size was 16 and the input signal length was 7 seconds. Network parameters were optimized using the Adam optimizer [34] with learning rate 10^{-5} . When combining signal and spatial loss, we took weighted sum of them, i.e., $\mathcal{L}_{doa} + \alpha \mathcal{L}_{sig}$, where α was set to 0.2 for DNN-IVA and 1.0 for DNN-MVDR. When training DNN-IVA, three of six channels were used and the number of iterations was 15. DNN-MVDR was trained using all six channels. MUSIC and AuxIVA also used six channels. We assumed time-varying complex Gaussian prior for AuxIVA (Gauss-IVA). When using WSJ-mix, we obtain DOAs from the entire signal, not using clustering described in 3.2. We used different STFT parameters for WPE and separation. The window/shift size were set to 512/128 and 4096/1024, respectively. The number of iterations, the delay and the tap length of WPE was set to 3, 3 and 10.

Table 3: WER of LibriCSS dataset. Sessions 1 to 9 were used for training. The proposed method is indicated by a star (*). Methods with a dagger (†) are initialized with weights pre-trained on WSJ-mix. The others are trained from scratch.

Algo.	Loss	Ch.	WER for session 0-9 (all) [%]							WER for session 0 (test) [%]						
			OS	OL	10.0	20.0	30.0	40.0	Avg	OS	OL	10.0	20.0	30.0	40.0	Avg
Unprocess.	-	-	11.8	11.7	18.8	27.2	35.6	43.3	26.4	11.3	7.3	19.3	25.9	34.5	38.2	24.4
	Chen <i>et al.</i> [21]	7	7.2	7.5	9.6	11.3	13.7	15.1	11.2	-	-	-	-	-	-	-
	Wang <i>et al.</i> [22]	7	5.8	5.8	5.9	6.5	7.7	8.3	6.8	-	-	-	-	-	-	-
DNN-IVA	Sup. (WSJ-mix)	3	7.1	7.3	9.0	11.3	13.5	14.4	10.9	6.6	6.5	9.7	11.7	10.7	12.0	9.8
Gauss-IVA	-	3	7.5	7.4	9.2	12.1	15.4	18.8	12.3	8.6	5.4	10.5	11.7	13.0	13.0	10.8
DNN-IVA	KLD	3	7.7	7.8	9.9	12.8	15.5	17.0	12.3	8.4	5.9	9.6	11.9	12.9	10.7	10.2
DNN-IVA	SDR	3	8.2	8.0	9.5	12.1	15.2	17.5	12.3	8.6	5.7	10.0	12.1	12.5	12.1	10.5
DNN-IVA*	DOA ₁	3	7.1	7.1	9.4	12.2	15.0	17.3	11.9	6.4	5.2	9.6	10.8	14.0	11.5	10.0
DNN-IVA*	DOA ₂	3	7.1	7.3	8.8	10.3	13.4	15.1	10.7	6.6	5.5	9.4	10.0	11.9	10.7	9.3
DNN-IVA*	KLD+DOA ₂	3	7.2	7.3	8.9	10.4	13.8	14.8	10.8	6.3	5.4	9.2	9.7	11.8	10.1	9.0
DNN-IVA*†	DOA ₂	3	6.5	6.9	9.1	10.8	13.0	14.1	10.5	5.9	5.6	9.0	9.7	9.9	9.1	8.4
DNN-IVA*†	KLD+DOA ₂	3	6.8	7.2	9.8	12.7	14.2	14.8	11.4	6.1	5.3	10.9	9.9	12.7	11.3	9.7

In test of IVA, we evaluated the performance using two, three and six channels with 30, 25 and 15 iterations. When using more than two channels, two separated signals with the highest power were evaluated. For testing, we used the averaged model parameters of the five epochs with lowest unsupervised loss on the validation set.

LibriCSS: We used LibriCSS [20] to evaluate the performance on real-recorded speeches. Utterances taken from LibriSpeech test clean set were played back from loudspeakers and recorded by a seven-channel circular microphones with radius of 4.25 cm. LibriCSS had 10 hours of audio recordings and they were separated into 10 sessions. Each session contained six mini-sessions with different overlap ratios ranging from 0 to 40%, OS, OL, 10, 20, 30 and 40, where OS/OL are overlap-free recordings with short/long inter-utterance silence. We conducted *utterance-wise* evaluation by evaluating WERs using the ASR systems given in [20]. We used sessions 1 to 9 as training set, and session 0 as test set.

Experimental setup was basically the same as that on WSJ-mix, but we used only three channels for Gauss-IVA and DNN-IVA to avoid degradation of WER due to over-separation. Note that MUSIC used all seven channels to estimate DOAs, where we used clustering described in 3.2, setting the window size L to 15, the shift S to 1, the number of clusters K to 3, the azimuth threshold θ_{thres} to 10 and the threshold to remove the outliers E_{thres} to $0.1 \sum_k E_k$. When training with KLD or SDR loss only, the learning rate was 10^{-5} ; otherwise 5×10^{-5} .

4.2. Results on WSJ0-mix dataset

Table 1 shows the evaluation results of IVA algorithms on the WSJ-mix test set. The evaluation metrics are SDR [35,36], the short-time objective intelligibility (STOI) [37], the perceptual evaluation of speech quality (PESQ) [38] and WER. The performance of supervised learning was also evaluated as an upper bound for unsupervised learning. At all channel numbers, DNN-IVA trained with our proposed DOA loss outperformed conventional Gauss-IVA and DNN-IVA trained with KLD or SDR loss. In terms of WER, our proposed method achieved comparable or even better performance than supervised learning with SDR or DOA₂ loss. We conjecture this significant WER reduction to be due to the fact that our spatial loss gives the distortionless constraint to the target source direction as in MVDR. In addition, it is less susceptible to inter-frequency permutation problem, because the DOA q_n is the same regardless of frequency. Furthermore, the spatial loss focuses on the direct signal, which would make the model more robust against rever-

beration. When we used both the spatial and the signal loss, no significant improvement was confirmed. This implies that the spatial information is sufficient to train DNN-IVA.

Table 2 shows the evaluation results of DNN-MVDR. Although the DOA loss worked well for IVA, it gave inferior performance to the SDR loss for MVDR. This would be due to the fact that IVA estimates the demixing matrix so that each separated signal is independent, whereas MVDR independently estimates the demixing matrix for each source. Thus, only the spatial information was not enough to train the network. However, we found that combining DOA and SDR loss in learning was better than either individually.

4.3. Results on LibriCSS dataset

Table 3 shows the WERs of separated signals from LibriCSS dataset. WERs per overlap ratio and average of them are listed. Compared to the signals losses, our proposed spatial loss led to higher performance. Both SDR and KLD loss gave only the comparable performance to pseudo-teacher, i.e., Gauss-IVA. The poor performance on overlap-free data implies that the signal loss leads to over-separation. Compared to DNN-IVA trained with WSJ-mix, which contained ten times more data than LibriCSS, our proposed method achieved higher performance. In addition, our proposed method also outperformed Conformer-based MVDR beamforming [21] with less than one twentieth the amount of data. Furthermore, the performance was further improved by fine-tuning the supervised trained model on WSJ-mix to LibriCSS using the proposed unsupervised learning. Although our proposed method did not exceed state-of-the-art performance [22], we showed high performance with the small amount of data, which demonstrates the effectiveness of unsupervised learning with in-domain data.

5. Conclusions

We proposed a spatial loss function that utilizes the DOAs estimated by classic techniques. It trains a neural separator so that it directs beam to the target source direction and null to the interference direction. We evaluated two neural separators, DNN-IVA and DNN-MVDR, using synthetic mixtures and real-recorded LibriCSS. In experiments with synthetic mixtures, we showed that the spatial loss worked well especially for DNN-IVA, which led to the same performance as the supervised loss. It also performed well for DNN-MVDR when combined with the signal loss. The spatial loss also gave high performance on LibriCSS. It outperformed the strong baselines with the small amount of training data without any labels.

6. References

- [1] P. Comon, “Independent component analysis, a new concept?” *Signal Processing*, vol. 36, no. 3, pp. 287–314, 1994.
- [2] T. Kim, T. Eltoft, and T.-W. Lee, “Independent vector analysis: An extension of ica to multivariate components,” in *International conference on independent component analysis and signal separation*. Springer, 2006, pp. 165–172.
- [3] A. Hiroe, “Solution of permutation problem in frequency domain ica, using multivariate probability density functions,” in *International Conference on Independent Component Analysis and Signal Separation*. Springer, 2006, pp. 601–608.
- [4] D. Kitamura, N. Ono, H. Sawada, H. Kameoka, and H. Saruwatari, “Determined blind source separation unifying independent vector analysis and nonnegative matrix factorization,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 24, no. 9, pp. 1626–1641, 2016.
- [5] D. Yu, M. Kolbæk, Z. H. Tan, and J. Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *Proc. IEEE ICASSP*, 2017, pp. 241–245.
- [6] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proc. IEEE ICASSP*, 2016, pp. 31–35.
- [7] 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.
- [8] 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.
- [9] T. Ochiai, S. Watanabe, T. Hori, and J. R. Hershey, “Multi-channel end-to-end speech recognition,” in *Proc. ICML*, 2017, p. 2632–2641.
- [10] C. Boeddeker et al., “Convolutional transfer function invariant SDR training criteria for multi-channel reverberant speech separation,” in *Proc. IEEE ICASSP*, 2021, pp. 8428–8432.
- [11] A. A. Nugraha, A. Liutkus, and E. Vincent, “Multichannel audio source separation with deep neural networks,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 24, no. 9, pp. 1652–1664, 2016.
- [12] N. Makishima et al., “Independent deeply learned matrix analysis for determined audio source separation,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 27, no. 10, pp. 1601–1615, 2019.
- [13] R. Scheibler and M. Togami, “Surrogate source model learning for determined source separation,” in *Proc. IEEE ICASSP*, 2021, pp. 176–180.
- [14] R. Haeb-Umbach, J. Heymann, L. Drude, S. Watanabe, M. Delcroix, and T. Nakatani, “Far-field automatic speech recognition,” *Proceedings of the IEEE*, vol. 109, no. 2, pp. 124–148, 2020.
- [15] R. Scheibler, E. Bezzam, and I. Dokmanić, “Pyroomacoustics: A python package for audio room simulation and array processing algorithms,” in *Proc. IEEE ICASSP*, 2018, pp. 351–355.
- [16] L. Drude, D. Hasenklever, and R. Haeb-Umbach, “Unsupervised training of a deep clustering model for multichannel blind source separation,” in *Proc. IEEE ICASSP*, 2019, pp. 695–699.
- [17] M. Togami, Y. Masuyama, T. Komatsu, and Y. Nakagome, “Unsupervised training for deep speech source separation with kullback-leibler divergence based probabilistic loss function,” in *Proc. IEEE ICASSP*, 2020, pp. 56–60.
- [18] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [19] N. Ono, “Stable and fast update rules for independent vector analysis based on auxiliary function technique,” in *Proc. IEEE WAS-PAA*, 2011, pp. 189–192.
- [20] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, X. Xiao, and J. Li, “Continuous speech separation: Dataset and analysis,” in *Proc. IEEE ICASSP*, 2020, pp. 7284–7288.
- [21] S. Chen, Y. Wu, Z. Chen, J. Wu, J. Li, T. Yoshioka, C. Wang, S. Liu, and M. Zhou, “Continuous speech separation with conformer,” in *Proc. IEEE ICASSP*, 2021, pp. 5749–5753.
- [22] Z.-Q. Wang, P. Wang, and D. Wang, “Multi-microphone complex spectral mapping for utterance-wise and continuous speech separation,” *IEEE/ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 2001–2014, 2021.
- [23] R. Scheibler and N. Ono, “Fast and stable blind source separation with rank-1 updates,” in *Proc. IEEE ICASSP*, 2020, pp. 236–240.
- [24] R. Scheibler and M. Togami, “Refinement of direction of arrival estimators by majorization-minimization optimization on the array manifold,” in *Proc. IEEE ICASSP*, 2021, pp. 436–440.
- [25] S. Thrun, “Affine structure from sound,” in *Advances in Neural Information Processing Systems*, Y. Weiss, B. Schölkopf, and J. Platt, Eds., vol. 18. MIT Press, 2005.
- [26] M. Kreković, G. Baechler, I. Dokmanić, and M. Vetterli, “Structure from sound with incomplete data,” in *Proc. IEEE ICASSP*, 2018, pp. 3539–3543.
- [27] T. Nakatani, T. Yoshioka, K. Kinoshita, M. Miyoshi, and B.-H. Juang, “Speech dereverberation based on variance-normalized delayed linear prediction,” *IEEE Trans. Audio Speech Lang. Process.*, vol. 18, no. 7, pp. 1717–1731, 2010.
- [28] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, “Language modeling with gated convolutional networks,” in *Proc. ICML*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, Aug 2017, pp. 933–941.
- [29] Y. Wu and K. He, “Group normalization,” in *Proc. ECCV*, 2018, pp. 3–19.
- [30] J. S. Garofolo et al., *CSR-I (WSJ0) Complete LDC93S6A*, Linguistic Data Consortium, Philadelphia, 1993, web Download.
- [31] Linguistic Data Consortium, and NIST Multimodal Information Group, *CSR-II (WSJ1) Complete LDC94S13A*, Linguistic Data Consortium, Philadelphia, 1994, web Download.
- [32] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, “The third ‘CHiME’ speech separation and recognition challenge: Dataset, task and baselines,” in *Proc. IEEE ASRU*, 2015, pp. 504–511.
- [33] S. Watanabe et al., “ESPnet: End-to-end speech processing toolkit,” in *Proc. ISCA INTERSPEECH*, 2018, pp. 2207–2211.
- [34] P. K. Diederik and B. Jimmy, “Adam: A method for stochastic optimization,” in *ICLR*, 2015.
- [35] E. Vincent, R. Gribonval, and C. Fevotte, “Performance measurement in blind audio source separation,” *IEEE Trans. Audio Speech Lang. Process.*, vol. 14, no. 4, pp. 1462–1469, Jun. 2006.
- [36] R. Scheibler, “SDR — Medium rare with fast computations,” in *Proc. IEEE ICASSP*, SG, May 2022.
- [37] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, “An algorithm for intelligibility prediction of time–frequency weighted noisy speech,” *IEEE Trans. Audio Speech Lang. Process.*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [38] 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. IEEE ICASSP*, vol. 2, 2001, pp. 749–752.