

# Enhanced Reverberation as Supervision for Unsupervised Speech Separation

Saijo, Kohei; Wichern, Gordon; Germain, François G; Pan, Zexu; Le Roux, Jonathan

TR2024-116 September 04, 2024

## Abstract

Reverberation as supervision (RAS) is a framework that allows for training monaural speech separation models from multi-channel mixtures in an unsupervised manner. In RAS, models are trained so that sources predicted from a mixture at an input channel can be mapped to reconstruct a mixture at a target channel. However, stable unsupervised training has so far only been achieved in over-determined source-channel conditions, leaving the key determined case unsolved. This work proposes enhanced RAS (ERAS) for solving this problem. Through qualitative analysis, we found that stable training can be achieved by leveraging the loss term to alleviate the frequency-permutation problem. Separation performance is also boosted by adding a novel loss term where separated signals mapped back to their own input mixture are used as pseudo-targets for the signals separated from other channels and mapped to the same channel. Experimental results demonstrate high stability and performance of ERAS.

*Interspeech 2024*



# Enhanced Reverberation as Supervision for Unsupervised Speech Separation

Kohei Saijo<sup>1,2\*</sup>, Gordon Wichern<sup>1</sup>, François G. Germain<sup>1</sup>, Zexu Pan<sup>1</sup>, Jonathan Le Roux<sup>1</sup>

<sup>1</sup>Mitsubishi Electric Research Laboratories (MERL), MA, USA <sup>2</sup>Waseda University, Tokyo, Japan  
 saiyo@pcl.cs.waseda.ac.jp, {wichern,germain,leroux}@merl.com

## Abstract

Reverberation as supervision (RAS) is a framework that allows for training monaural speech separation models from multi-channel mixtures in an unsupervised manner. In RAS, models are trained so that sources predicted from a mixture at an input channel can be mapped to reconstruct a mixture at a target channel. However, stable unsupervised training has so far only been achieved in over-determined source-channel conditions, leaving the key determined case unsolved. This work proposes enhanced RAS (ERAS) for solving this problem. Through qualitative analysis, we found that stable training can be achieved by leveraging the loss term to alleviate the frequency-permutation problem. Separation performance is also boosted by adding a novel loss term where separated signals mapped back to their own input mixture are used as pseudo-targets for the signals separated from other channels and mapped to the same channel. Experimental results demonstrate high stability and performance of ERAS.

**Index Terms:** unsupervised speech separation, reverberation as supervision, deep learning

## 1. Introduction

Speech separation has been intensively investigated for listening applications or as a front-end for applications such as automatic speech recognition [1,2] or speaker diarization [3,4]. Pioneered by deep clustering [5] and permutation invariant training [5,6], neural network (NN)-based approaches have become a major technique to achieve high-fidelity separation [7,8]. Most NN-based methods are trained in a supervised manner and rely on synthetic data as it is hard to collect pairs of mixtures and their individual sources in real environments. NNs are however known to be vulnerable to domain mismatch, and separation models trained on synthetic data often perform poorly in real environments [9]. Unsupervised speech separation techniques that can leverage recorded unlabeled mixtures can be the key to success in real-world applications [10–17]. We focus here on developing such a training technique for monaural separation.

Recently, MixIT [18] and remixing-based methods [19–21] have shown great success for unsupervised separation. They artificially create mixtures-of-mixtures or remix pseudo-mixtures to achieve unsupervised learning, but such artificial mixtures cause another kind of domain mismatch against the normal mixtures seen at inference. Another direction is exploiting *multi-channel* mixtures to train a monaural separation model. Prior work has utilized spatial cues as pseudo-targets [22,23], avoiding the domain mismatch issue altogether. However, performance is bounded by the pseudo-targets' quality.

\*This work was done while K. Saijo was an intern at MERL.

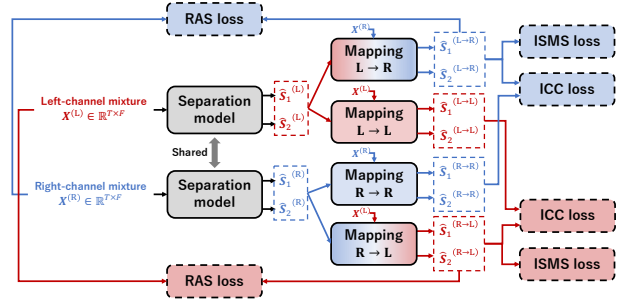


Figure 1: Overview of ERAS training. Separated signals at the left (L) or right (R) channel are mapped to the opposite channel by relative RIR estimation, and the model is trained to reconstruct mixtures as the sum of the mapped sources (RAS loss). ERAS improves training stability by strongly penalize undesirable solution by ISMS loss and boosts performance by introducing an inter-channel consistency (ICC) loss aiming to make sources mapped to the same channel closer.

Reverberation as supervision (RAS) was proposed for effectively leveraging multi-channel mixtures to train monaural separation models [24]. In RAS, separated signals from an input channel are mapped to another target channel by relative room impulse response (RIR) estimation (e.g., Wiener filtering), and the model is trained to reconstruct the mixture at the target channel. The idea is that recovering the mixture by mapping from well-separated signals is much easier than from the mixture itself, so the model will learn to separate the sources to reconstruct the mixture well. Since both input and target are mixtures, RAS ideally overcomes both domain-mismatch and pseudo-target-quality problems. In the original RAS [24], however, unsupervised learning of a two-speaker separation model using two-speaker mixtures failed, which implies that there are undesirable solutions where the model outputs signals that are not well separated but from which it is easy to recover the mixture. Unsupervised neural speech separation leveraging over-determined mixtures (UNSSOR) has shown that we can avoid such undesirable solutions when we have more channels than sources [25]. Intuitively, more constraints are imposed on the model outputs by using more microphones, because the model has to estimate signals from which the mixtures at all the microphones can be reconstructed. Still, fully-unsupervised RAS training in the determined condition remains unsolved.

In this paper, we tackle unsupervised RAS training in the determined setup, particularly training a monaural two-speaker separation model using two-channel mixtures. We have discovered that sources which are not well separated but are frequency-permuted result in a undesirable solution with respect to RAS loss, but such a solution can be avoided by

leveraging the loss to alleviate the frequency-permutation problem [25]. To boost performance, we introduce an inter-channel consistency loss in which separated signals mapped back to their input mixture are used as pseudo-targets for those separated from other channels and mapped to the same channel. Our observation is that the former signals have higher quality and thus improve the quality of the latter. By further introducing an effective two-stage training strategy, the proposed method, called enhanced RAS (ERAS), achieves both stable training and high separation performance.

## 2. Reverberation as Supervision

Assuming  $N$  sources are recorded with an  $M$ -channel microphone array, the mixture signal  $x$  observed at the  $m$ -th microphone,  $m = 1, \dots, M$ , can be written as

$$x_i^{(m)} = \sum_{n=1}^N (h_n^{(m)} * \bar{s}_n)_l, \quad (1)$$

where  $\bar{s}_n$  denotes the dry source,  $n = 1, \dots, N$  the speaker index,  $l = 1, \dots, L$  the time sample index, and  $h_n^{(m)}$  the room impulse response (RIR) from the  $n$ -th source to the  $m$ -th microphone, with  $h_n^{(m)} * \bar{s}_n$  being referred to as the source-image signal. We denote by  $s_n^{(m)}$  the direct-path signal. In the short-time Fourier transform (STFT) domain, we respectively denote as  $X_{t,f}^{(m)}$  and  $S_{n,t,f}^{(m)}$  the mixture and the  $n$ -th source at time frame  $t = 1, \dots, T$  and frequency index  $f = 1, \dots, F$ .

### 2.1. RAS and UNSSOR

RAS has been proposed to train monaural speech separation in an unsupervised manner by using multi-channel mixtures [24]. Specifically, RAS has been used to train a two-speaker separation model using two-channel mixtures ( $M = N = 2$ ).

Referring to the channel used as input as the reference channel and denoting its index as  $m_r$  ( $1 \leq m_r \leq M$ ), the separation model  $\mathcal{F}$  takes  $x^{(m_r)}$  and separates it into  $N$  sources:

$$(\hat{s}_1^{(m_r)}, \dots, \hat{s}_N^{(m_r)}) = \mathcal{F}(x^{(m_r)}), \quad (2)$$

with  $\hat{s}_n^{(m_r)}$  the  $n$ -th separated signal. RAS aims to reconstruct the mixture at another microphone  $m$  ( $1 \leq m \leq M, m \neq m_r$ ) using the separated signals. However, the model cannot directly estimate the signals at the  $m$ -th channel since it does not know the position of each source and thus cannot learn to compensate for its RIR to the  $m$ -th channel. RAS thus relies on an external relative RIR estimation (e.g., the Wiener filter) to map each separated signal at the  $m_r$ -th channel to the  $m$ -th channel:

$$\hat{s}_n^{(m_r \rightarrow m)} = \mathcal{M}(\hat{s}_n^{(m_r)}, x^{(m)}), \quad (3)$$

where  $\mathcal{M}$  denotes the mapping via relative RIR estimation, and  $\hat{s}_n^{(m_r \rightarrow m)}$  is the signal obtained by mapping  $\hat{s}_n^{(m_r)}$  to the  $m$ -th channel. Now an estimate of the mixture  $x^{(m)}$  at the  $m$ -th microphone can be obtained by summing up the estimated sources  $\hat{s}_n^{(m_r \rightarrow m)}$  for all  $n$ . The RAS loss encourages the model to output signals from which the mixture can be reconstructed well:

$$\mathcal{L}_{\text{RAS}}^{(m_r \rightarrow m)} = \mathcal{L}\left(x^{(m)}, \sum_n \hat{s}_n^{(m_r \rightarrow m)}\right), \quad (4)$$

where  $\mathcal{L}$  is the signal-level loss function. Table 1 shows how well mixtures can be reconstructed from various signals at another microphone by the two filtering methods considered so far in RAS-related works, Wiener filtering [24] and forward convolutive prediction (FCP) [26] (detailed in Section 2.2). The re-

Table 1: *SI-SNR [dB] when predicting a mixture at a channel from signals at another channel with Wiener filter or FCP.*

Filtering input	NF-WHAMR!		SMS-WSJ	
	Wiener	FCP	Wiener	FCP
Mixture ↓	6.2	<b>4.3</b>	6.8	<b>5.1</b>
Source-image signals ↑	12.1	<b>13.4</b>	11.2	<b>11.4</b>
Direct-path + early reflections ↑	12.7	<b>14.4</b>	12.2	<b>13.1</b>
Dry sources ↑	14.2	<b>17.0</b>	15.2	<b>17.4</b>

sults show that recovering a mixture at some channel using well-separated sources is much easier than using the mixture at another channel. Minimizing Eq. (4) thus pushes the model to separate sources. However, the original RAS, where  $M = N = 2$  is assumed, failed in a fully unsupervised setup, which implies that there are undesirable solutions where the model's outputs are not well separated but can easily reconstruct mixtures.

In [25], it has been shown that such undesirable solutions in RAS can be avoided by leveraging over-determined mixtures ( $M > N$ ). The method, called UNSSOR, aims to reconstruct mixtures at all the  $M$  channels available:

$$\mathcal{L}_{\text{UNSSOR}}^{(m_r)} = \sum_{m=1}^M \alpha^{(m)} \mathcal{L}_{\text{RAS}}^{(m_r \rightarrow m)}, \quad (5)$$

where  $\alpha^{(m)}$  ( $> 0$ ) is the loss weight of each microphone. Unlike RAS, UNSSOR tries to recover the mixture even on the reference microphone (i.e.,  $m$  can be equal to  $m_r$ ). Since the loss  $\mathcal{L}_{\text{RAS}}^{(m_r \rightarrow m_r)}$  has a trivial solution,  $\alpha^{(m_r)}$  is set to less than 1, while  $\alpha^{(m \neq m_r)}$  is set to 1. Intuitively, the more microphones we have, the more constraints are imposed on the model outputs because the model has to estimate the signals from which the mixtures at all microphones can be reconstructed. Thus, having more RAS loss terms (Eq. (4)) can improve training stability, and UNSSOR can achieve unsupervised learning in over-determined conditions. However, unsupervised RAS training in the determined condition ( $M = N$ ) remains unsolved.

### 2.2. Relative RIR estimation methods for channel mapping

In the RAS training, we have several choices of relative RIR estimation methods for the mapping  $\mathcal{M}$ . The original RAS utilized the Wiener filter [24], while UNSSOR used FCP [26]. The Wiener filtering is done in the time domain, and a source at the  $m_r$ -th channel is mapped to the  $m$ -th channel as:

$$\hat{s}_n^{(m_r \rightarrow m)} = \mathcal{M}(\hat{s}_n^{(m_r)}, x^{(m)}) = \hat{w}_n^{(m)} * \hat{s}_n^{(m_r)}, \\ \hat{w}_n^{(m)} = \arg \min_{w_n^{(m)}} \sum_l |x_l^{(m)} - (w_n^{(m)} * \hat{s}_n^{(m_r)})_l|^2, \quad (6)$$

where  $\hat{w}_n^{(m)}$  is the Wiener filter. In contrast, FCP operates in the STFT domain. The mapping is done with a time-invariant  $K$ -tap filter  $\mathbf{g}_{n,f}^{(m)}$  (with  $K_{\text{past}}$  past frames, 1 current frame, and  $K_{\text{future}}$  future frames) as:

$$\hat{S}_{n,t,f}^{(m_r \rightarrow m)} = (\mathcal{M}(\hat{S}_n^{(m_r)}, X^{(m)}))_{t,f} = (\mathbf{g}_{n,f}^{(m)})^H \tilde{\hat{S}}_{n,t,f}^{(m_r)}, \\ \hat{\mathbf{g}}_{n,f}^{(m)} = \arg \min_{\mathbf{g}_{n,f}^{(m)}} \sum_t \frac{1}{\lambda_{t,f}^{(m)}} \left| X_{t,f}^{(m)} - (\mathbf{g}_{n,f}^{(m)})^H \tilde{\hat{S}}_{n,t,f}^{(m_r)} \right|^2, \quad (7)$$

where  $\tilde{\hat{S}}_{n,t,f} \in \mathbb{C}^K$  stacks TF bins from  $K_{\text{past}}$  past frames and  $K_{\text{future}}$  future frames of  $\hat{S}_{n,\cdot,f}$  with the bin from the current frame  $\hat{S}_{n,t,f}$ . Scaling parameter  $\lambda$  is defined in [25] as  $\lambda_{t,f}^{(m)} = (\frac{1}{M} \sum_m |X_{t,f}^{(m)}|^2) + 10^{-4} \times \max(\frac{1}{M} \sum_m |X_{t,f}^{(m)}|^2)$ .

In [25], it is shown that FCP leads to better performance.

This conforms with the results in Table 1, which imply that FCP leads to higher performance and more stable RAS training. We thus use FCP in this work<sup>1</sup>. Because FCP is applied independently for each frequency, it is clearly susceptible to the frequency-permutation problem. In [25], a loss term called intra-source magnitude scattering (ISMS) loss was introduced to alleviate this problem:

$$\mathcal{L}_{\text{ISMS}}^{(m_r \rightarrow m)} = \frac{\sum_t \frac{1}{N} \sum_n \text{var}(\log(|\hat{\mathbf{S}}_{n,t}^{(m_r \rightarrow m)}|))}{\sum_t \text{var}(\log(|\mathbf{X}_t^{(m)}|))}, \quad (8)$$

where  $|\cdot|$  is magnitude and  $\text{var}(\cdot)$  is the variance over frequencies. The ISMS loss favors magnitude spectrograms with smaller variance and thus mitigates the frequency-permutation problem. The RAS loss when using FCP is replaced by:

$$\mathcal{L}_{\text{RAS+ISMS}}^{(m_r \rightarrow m)} = \mathcal{L}_{\text{RAS}}^{(m_r \rightarrow m)} + \beta \mathcal{L}_{\text{ISMS}}^{(m_r \rightarrow m)}, \quad (9)$$

where  $\beta$  is a weight for the ISMS loss.

### 3. Enhanced Reverberation as Supervision

This work tackles the unsolved problem of unsupervised learning of a monaural separation system on determined mixtures, particularly focusing on training a monaural two-speaker separation model using two-channel mixtures ( $N = M = 2$ )<sup>2</sup>, as in the original RAS work.

#### 3.1. Stabilizing the training

Interestingly, in preliminary experiments, we found that just replacing the BLSTM-based mask estimation network in the original RAS [24] with TF-GridNet [8] made training lead to successful separation even in the determined condition for some initial seed values (with  $\beta = 0.1$ ). For the other seeds, it led to irrelevant solutions with very low separation quality, despite converging in terms of RAS loss. In the failure cases, we found that the sources in the output were often frequency-permuted.

In UNSSOR, where training stability was not mentioned as an issue thanks to the over-determined setting, the ISMS loss was introduced solely to make the method frequency-permutation-solver-free, and its weight  $\beta$  was set to a low value ( $\beta = 0.02$  for  $M = 6$  and  $\beta = 0.06$  for  $M = 3$ )<sup>3</sup>. Since the ISMS loss was shown there to be effective for avoiding frequency-permuted outputs, we consider here leveraging it as a regularizer to stabilize training. Indeed, we find the ISMS loss value to be always higher when the training fails than when the training goes well. However, with weight values  $\beta$  similar to those in UNSSOR or even  $\beta = 0.1$  as mentioned above, the training could still fail for some seed values. We thus propose to strongly penalize such undesirable solutions by setting the ISMS loss weight to a high value such as  $\beta = 0.3$ . We empirically demonstrate in Section 4.2 that the ISMS loss then greatly benefits the training stability.

#### 3.2. Inter-channel consistency loss

While the modification in Section 3.1 is for stabilizing the training, we propose an inter-channel consistency (ICC) loss to improve the separation performance. Let us denote the two channels as L (left) and R (right), and define  $m \in [L, R]$ , for simplicity. Suppose we have separated signals from both L ( $\hat{s}_n^{(L)}$ ) and

<sup>1</sup>We also tried the Wiener filter [24]: it worked but was less stable.

<sup>2</sup>Our proposed approach algorithmically works regardless of  $N$  or  $M$  but we leave the experiments on other configurations as future work.

<sup>3</sup>Personal communication.

Table 2: Oracle ISMS loss value (Eq. (8)) on validation set. In SMS-WSJ\*, we only consider fully-overlapped segments.

	NF-WHAMR!	SMS-WSJ	SMS-WSJ*
Mixture, mixture	1.00	1.00	1.00
Mixture, zero signal	0.50	0.50	0.50
Zero signal, zero signal	0.00	0.00	0.00
Source-image signals	1.05	1.53	1.61
Freq. permuted clean signals	1.45	10.83	2.22

R ( $\hat{s}_n^{(R)}$ )<sup>4</sup>, we can consider four signals from mappings L→R, L→L, R→R, and R→L (see Fig. 1). Then, for instance, both  $\hat{s}_n^{(L \rightarrow L)}$  and  $\hat{s}_n^{(R \rightarrow L)}$  are mapped to channel L, but we also find that  $\hat{s}_n^{(L \rightarrow L)}$  has better quality than  $\hat{s}_n^{(R \rightarrow L)}$ , likely since mapping to the reference channel itself (L→L) is much easier than mapping to another channel (R→L). This motivates us to use  $\hat{s}_n^{(L \rightarrow L)}$  as pseudo-label of  $\hat{s}_n^{(R \rightarrow L)}$  while stopping the gradient of  $\hat{s}_n^{(L \rightarrow L)}$ , i.e., the loss term:

$$\mathcal{L}_{\text{ICC}}^{(m_r \rightarrow m)} = \frac{1}{N} \sum_n \mathcal{L}(\hat{s}_n^{(m \rightarrow m)}, \hat{s}_\pi^{(m_r \rightarrow m)}), \quad (10)$$

where  $\pi$  is the optimal permutation minimizing the loss [5, 6]. The final combined loss is obtained with ICC loss weight  $\gamma$ :

$$\mathcal{L}_{\text{ERAS}}^{(m_r \rightarrow m)} = \mathcal{L}_{\text{RAS+ISMS}}^{(m_r \rightarrow m)} + \gamma \mathcal{L}_{\text{ICC}}^{(m_r \rightarrow m)}. \quad (11)$$

#### 3.3. Two-stage training

While we introduced methods for stabilizing the training and improving the separation performance in Sections 3.1 and 3.2, respectively, we found that two-stage training on the ISMS loss and the ICC loss is helpful to guarantee both stability and strong performance.

**ISMS loss:** Although the ISMS loss is used for stabilizing the training, we found that it limits the separation performance. Table 2 shows the ISMS loss values computed on various signals<sup>5</sup>. The result implies that the ISMS loss may penalize well-separated signals too much compared to other signals. In particular, the ISMS loss could interfere with accurate magnitude estimation because it favors spectrograms whose bins all have the same magnitude. This analysis motivates us to improve the separation performance by removing the ISMS loss after the model gets to a certain level of separation performance. Specifically, we first train the model with  $\beta > 0.0$  and then fine-tune it with  $\beta = 0.0$ .

**ICC loss:** In preliminary experiments, we found that the ICC loss improves the separation performance but does not contribute to stabilization of training. We again use a two-stage approach, where we first pre-train the model with  $\gamma = 0.0$  for several epochs and then switch to  $\gamma > 0.0$ .

## 4. Experiments

#### 4.1. Experimental setup

We use two datasets for our experiments, both contain speech from WSJ [27, 28] and a sampling rate of 8 kHz.

<sup>4</sup>This is achieved by including both channels in the same mini-batch, i.e., the mini-batch is  $[x^{(L)}, x^{(R)}, \dots]$ .

<sup>5</sup>The input signals are randomly cut into 4-second segments to imitate the training setup. Note that when frequency-permuting the clean sources, the loss on SMS-WSJ is higher than on NF-WHAMR! because SMS-WSJ mixtures are partially overlapped. As we use 4-second segments, one of the sources is often zero-padded and frequency permutation significantly increases the variance of the log-magnitude spectrogram for zero-padded frames. Indeed, the loss on the fully-overlapped SMS-WSJ\* is much lower.

Table 3: Effectiveness of adding ICC loss and removing ISMS loss in ERAS training.  $\beta$  and  $\gamma$  are the ISMS loss and ICC loss weights.

ID	Method	$\beta$	$\gamma$	NF-WHAMR!			SMS-WSJ				
				SI-SNR	SDR	PESQ	SI-SNR	SDR	PESQ		
A1	ERAS	0.3	0.0	13.1	13.9	3.27	11.3	12.0	2.96		
A2	A1@20epochs $\rightarrow$ ERAS	<b>0.0</b>	0.0	13.8	14.7	3.65	13.7	14.4	3.54		
A3	A1@20epochs $\rightarrow$ ERAS	0.3	<b>0.1</b>	13.2	13.9	3.32	11.7	12.4	3.07		
A4	A1@20epochs $\rightarrow$ ERAS	<b>0.0</b>	<b>0.1</b>	<b>14.7</b>	<b>15.6</b>	<b>3.75</b>	<b>13.9</b>	<b>14.5</b>	<b>3.55</b>		
Supervised				-	-	15.7	16.6	3.87	15.8	16.4	3.89

Table 4: Number of training successes / failures among 25 trials with different initial seed values, when changing ISMS loss weight  $\beta$ .  $\alpha^{(m_r)}$  is loss weight on reference microphone.

$\alpha^{(m_r)}$	NF-WHAMR!				SMS-WSJ
	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.2$	$\beta = 0.3$	$\beta = 0.3$
0.0	15 / 10	20 / 5	21 / 4	<b>25 / 0</b>	<b>25 / 0</b>
0.1	0 / 25	0 / 25	9 / 16	19 / 6	21 / 4

**WHAMR!** [29] includes two-channel two-speaker mixtures (with spacing from 15 cm to 17 cm). We use the `min` version, where the two speech signals are fully overlapped. Reverberation times range from 0.1 s to 1.0 s. Training, validation, and test sets contain 20,000 ( $\sim 30$ h), 5,000 ( $\sim 10$ h), and 3,000 ( $\sim 5$ h) mixtures, respectively. We use a noise-free version of WHAMR! (NF-WHAMR!) since the strong noise weakens our assumption that relative RIR maps signals from a channel to others, and investigation in noisy setups is left to future work.

**SMS-WSJ** [30] contains six-channel two-speaker mixtures. We use the first and second microphones with a spacing of 10 cm. The reverberation times range from 0.2 s to 0.5 s, and a weak Gaussian noise is added at an SNR from 20 dB to 30 dB. The average overlap ratio of two speech signals is around 50%. Training, validation, and test sets contain 33,561 ( $\sim 87.4$ h), 982 ( $\sim 2.5$ h), and 1,332 ( $\sim 3.4$ h) mixtures, respectively.

We use TF-GridNet [8] (from ESPNet-SE [31, 32]) as the separation model. We set  $B = 4$ ,  $D = 48$ ,  $I = 4$ ,  $J = 1$  and  $H = 256$  (notations follow Table 1 in [8]). The model takes as inputs the real and imaginary (RI) parts of the mixture spectrogram and estimates the RI parts of each source.

In all the experiments, we train the model up to 100 epochs using the Adam optimizer [33] with an initial learning rate of  $1e-3$ . The learning rate is halved if validation loss does not improve for two epochs. Gradient clipping is applied with a maximum gradient  $L_2$ -norm of 1. The batch size is 8 and the input mixture is 4 s long. Each input mixture is normalized by dividing it by its standard deviation. For STFT, we use a 32 ms long square-root Hann window with hop size of 8 ms. In FCP, we set  $K_{\text{past}} = 19$  and  $K_{\text{future}} = 1$ . As loss function  $\mathcal{L}$ , we use the  $L_1$  distance between the RI components and magnitudes of the reference and the estimate, normalized by the  $L_1$  norm of the input mixture, following [25].

For evaluation, we use the source-image (reverberant) signals as ground truth. We apply FCP to the network outputs to map them to the reference channel so that the ground truths and separated signals are time-aligned. The evaluation metrics are SI-SNR, source-to-distortion ratio (SDR) [34, 35], and perceptual evaluation of speech quality (PESQ) [36].

## 4.2. Training stability of ERAS

As discussed in Section 3.1, we aim to stabilize the training by leveraging the ISMS loss. To examine the stability of the proposed ERAS, we trained the model while changing the ISMS loss weight  $\beta$ , with 25 different initial seeds. We also investi-

gate UNSSOR-like RAS, where the RAS loss is also computed on the reference channel as in Eq. (5), with  $\alpha^{(m_r)} = 0.1$ , to investigate if such a strategy is helpful for stabilizing training. Table 4 shows the number of training successes and failures when changing  $\beta$ , where we consider training to have failed when the average SI-SNR on the validation set does not reach 5 dB after training for 10 epochs<sup>6</sup>. We set  $\gamma = 0.0$  to focus on the effectiveness of the ISMS loss.

We can clearly observe that the number of failures decreases when using higher ISMS loss weights  $\beta$ . All the 25 trials succeeded on NF-WHAMR! for  $\beta = 0.3$  and  $\alpha^{(m_r)} = 0.0$ . We confirmed that  $\beta = 0.3$  was also an effective value on SMS-WSJ. Although we observed that the UNSSOR-like RAS with  $\alpha^{(m_r)} = 0.1$  was more unstable, utilizing high  $\beta$  was still effective for training stabilization.

## 4.3. Performance of ERAS

In this section, we investigate performance when removing the ISMS loss and adding the proposed ICC loss in a two-stage training as described in Section 3.3. After training model A1 with  $\beta = 0.3$ , we use its 20-th epoch checkpoint to initialize other models, which are fine-tuned for 80 epochs after changing loss weights as shown in Table 3. When fine-tuning, we warm up the learning rate for 4000 training steps from 0 to  $1e-3$ .

Results of A1 and A2 show that removing the ISMS loss is effective, particularly on SMS-WSJ. This is expected as the ISMS loss value on SMS-WSJ can be high even with perfect separation (cf. Section 3.3). On NF-WHAMR!, PESQ, which favors more accurate magnitude estimation, showed large improvements, in line with the fact that the ISMS loss leads to inaccurate magnitude estimation as discussed in Section 3.3.

The result of A3 shows that the ICC loss is also beneficial. Removing the ISMS loss while adding the ICC loss (A4) leads to the best performance. While there is still a gap with supervised learning, the proposed ERAS improved both training stability and separation quality.

## 5. Conclusion

We proposed ERAS, a stable and high-performing unsupervised speech separation method for training a monaural separation model using multi-channel mixtures. Inspired by the analysis that implies the ISMS loss could be effective as the regularizer to avoid undesirable solutions, we achieved stable training by setting the ISMS-loss weight high. We also proposed ICC loss and introduced a two-stage training strategy to ensure both high stability and performance. In future work, we will further improve ERAS by imposing more constraints such as spatial consistency [15, 37]. Our source code is available online<sup>7</sup>.

<sup>6</sup>We believe that this definition is reasonable because the SI-SNR reaches around 10 dB after 10 epochs when the training is successful.

<sup>7</sup><https://github.com/merlresearch/reverberation-as-supervision>

## 6. References

- [1] S. Settle, J. Le Roux, T. Hori, S. Watanabe, and J. R. Hershey, “End-to-end multi-speaker speech recognition,” in *Proc. ICASSP*, 2018.
- [2] R. Haeb-Umbach, J. Heymann, L. Drude, S. Watanabe, M. Delcroix, and T. Nakatani, “Far-field automatic speech recognition,” *Proc. IEEE*, vol. 109, no. 2, pp. 124–148, 2020.
- [3] K. Kinoshita, M. Delcroix, S. Araki, and T. Nakatani, “Tackling real noisy reverberant meetings with all-neural source separation, counting, and diarization system,” in *Proc. ICASSP*, 2020.
- [4] D. Raj, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo *et al.*, “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *Proc. SLT*, 2021.
- [5] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *Proc. ICASSP*, 2016.
- [6] 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. ICASSP*, 2017.
- [7] Y. Luo and N. Mesgarani, “Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [8] Z.-Q. Wang, S. Cornell, S. Choi, Y. Lee, B.-Y. Kim, and S. Watanabe, “TF-GridNet: Integrating full-and sub-band modeling for speech separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 2023.
- [9] W. Zhang, J. Shi, C. Li, S. Watanabe, and Y. Qian, “Closing the gap between time-domain multi-channel speech enhancement on real and simulation conditions,” in *Proc. WASPAA*. IEEE, 2021.
- [10] L. Drude, D. Hasenklever, and R. Haeb-Umbach, “Unsupervised training of a deep clustering model for multichannel blind source separation,” in *Proc. ICASSP*, 2019.
- [11] 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. ICASSP*, 2020.
- [12] K. Saijo and T. Ogawa, “Remix-cycle-consistent learning on adversarially learned separator for accurate and stable unsupervised speech separation,” in *Proc. ICASSP*, 2022.
- [13] —, “Unsupervised training of sequential neural beamformer using coarsely-separated and non-separated signals,” in *Proc. Interspeech*, 2022.
- [14] H. Munakata, R. Takeda, and K. Komatani, “Training data generation with DOA-based selecting and remixing for unsupervised training of deep separation models,” in *Proc. Interspeech*, 2022.
- [15] K. Saijo and R. Scheibler, “Spatial loss for unsupervised multi-channel source separation,” in *Proc. Interspeech*, 2022.
- [16] Y. Bando, K. Sekiguchi, Y. Masuyama, A. A. Nugraha, M. Fontaine, and K. Yoshii, “Neural full-rank spatial covariance analysis for blind source separation,” *IEEE Signal Process. Lett.*, vol. 28, pp. 1670–1674, 2021.
- [17] Y. Bando, Y. Masuyama, A. A. Nugraha, and K. Yoshii, “Neural fast full-rank spatial covariance analysis for blind source separation,” *arXiv preprint arXiv:2306.10240*, 2023.
- [18] S. Wisdom, E. Tzinis, H. Erdogan, R. Weiss, K. Wilson, and J. Hershey, “Unsupervised sound separation using mixture invariant training,” *Proc. NeurIPS*, vol. 33, pp. 3846–3857, 2020.
- [19] E. Tzinis, Y. Adi, V. K. Ithapu, B. Xu, P. Smaragdis, and A. Kumar, “RemixIT: Continual self-training of speech enhancement models via bootstrapped remixing,” *IEEE J. Sel. Top. Signal Process.*, 2022.
- [20] K. Saijo and T. Ogawa, “Self-Remixing: Unsupervised speech separation via separation and remixing,” in *Proc. ICASSP*, 2023.
- [21] —, “Remixing-based unsupervised source separation from scratch,” in *Proc. Interspeech*, 2023.
- [22] E. Tzinis, S. Venkataramani, and P. Smaragdis, “Unsupervised deep clustering for source separation: Direct learning from mixtures using spatial information,” in *Proc. ICASSP*, 2019.
- [23] P. Seetharaman, G. Wichern, J. Le Roux, and B. Pardo, “Bootstrapping single-channel source separation via unsupervised spatial clustering on stereo mixtures,” in *Proc. ICASSP*, 2019.
- [24] R. Aralikatti, C. Boeddeker, G. Wichern, A. Subramanian, and J. Le Roux, “Reverberation as supervision for speech separation,” in *Proc. ICASSP*, 2023.
- [25] Z.-Q. Wang and S. Watanabe, “UNSSOR: Unsupervised neural speech separation by leveraging over-determined training mixtures,” *arXiv preprint arXiv:2305.20054*, 2023.
- [26] Z.-Q. Wang, G. Wichern, and J. Le Roux, “Convolutional prediction for monaural speech dereverberation and noisy-reverberant speaker separation,” *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 29, pp. 3476–3490, 2021.
- [27] J. S. Garofolo *et al.*, *CSR-I (WSJ0) Complete LDC93S6A*, Linguistic Data Consortium, Philadelphia, 1993, Web download.
- [28] Linguistic Data Consortium, and NIST Multimodal Information Group, *CSR-II (WSJ1) Complete LDC94S13A*, Linguistic Data Consortium, Philadelphia, 1994, Web download.
- [29] M. Maciejewski, G. Wichern, E. McQuinn, and J. Le Roux, “WHAMR!: Noisy and reverberant single-channel speech separation,” in *Proc. ICASSP*, 2020.
- [30] L. Drude, J. Heitkaemper, C. Boeddeker, and R. Haeb-Umbach, “SMS-WSJ: Database, performance measures, and baseline recipe for multi-channel source separation and recognition,” *arXiv preprint arXiv:1910.13934*, 2019.
- [31] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, “ESPnet: End-to-end speech processing toolkit,” in *Proc. Interspeech*, 2018.
- [32] C. Li, J. Shi, W. Zhang, A. S. Subramanian, X. Chang, N. Kamo, M. Hira, T. Hayashi, C. Boeddeker, Z. Chen *et al.*, “ESPnet-SE: End-to-end speech enhancement and separation toolkit designed for asr integration,” in *Proc. SLT*, 2021.
- [33] P. K. Diederik and B. Jimmy, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.
- [34] E. Vincent, R. Gribonval, and C. Févotte, “Performance measurement in blind audio source separation,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [35] R. Scheibler, “SDR — Medium rare with fast computations,” in *Proc. ICASSP*, May 2022.
- [36] 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, 2001.
- [37] R. Falcon-Perez, G. Wichern, F. G. Germain, and J. Le Roux, “Location as supervision for weakly supervised multi-channel source separation of machine sounds,” in *Proc. WASPAA*. IEEE, 2023.