



Remixing-based Unsupervised Source Separation from Scratch

Kohei Saijo¹, Tetsuji Ogawa¹

¹Department of Communications and Computer Engineering, Waseda University, Tokyo, Japan

saijo@pcl.cs.waseda.ac.jp

Abstract

We propose an unsupervised approach for training separation models from scratch using *RemixIT* and *Self-Remixing*, which are recently proposed self-supervised learning methods for refining pre-trained models. They first separate mixtures with a teacher model and create *pseudo-mixtures* by shuffling and remixing the separated signals. A student model is then trained to separate the pseudo-mixtures using either the teacher's outputs or the initial mixtures as supervision. To refine the teacher's outputs, the teacher's weights are updated with the student's weights. While these methods originally assumed that the teacher is pre-trained, we show that they are capable of training models from scratch. We also introduce a simple remixing method to stabilize training. Experimental results demonstrate that the proposed approach outperforms mixture invariant training, which is currently the only available approach for training a monaural separation model from scratch.

Index Terms: monaural source separation, unsupervised learning, remixing, deep learning

1. Introduction

Remarkable progress has been made in source separation thanks to advances in neural networks [1–3]. Typically, models are trained with supervised learning, exploiting a large number of the mixture and ground truth pairs. However, obtaining these pairs in real-world environments is challenging, so they are generally synthesized using simulation toolkits [4]. Unfortunately, the separation performance in real environments is often degraded due to mismatches with real-recorded mixtures, such as differences in reverberation, noise type, or channel mismatch.

To address this issue, unsupervised learning methods directly train models using real-recorded mixtures without ground truths [5, 6]. In monaural source separation, mixture invariant training (MixIT) [7] has been highly successful in a variety of separation tasks by training models to separate a mixture of mixtures (MoM), which is created by summing up multiple mixtures [8–10]. However, due to the mismatch between MoMs and the actual mixtures in that MoMs contain more sources, models often suffer from the *over-separation* problem.

To address the over-separation issue, several attempts have been made to refine MixIT pre-trained models in an unsupervised manner. Teacher-student MixIT [11] trains a new student model with fewer output channels, using the MixIT pre-trained model as a teacher. On the other hand, RemixIT [12, 13] trains a student model to estimate the teacher's outputs from a *pseudo-mixture* generated by shuffling and remixing the teacher's outputs. While Teacher-student MixIT uses a static teacher, RemixIT iteratively updates the teacher's weights using the student's weights, leading to improved performance. Self-

Remixing [14] has a similar training process to RemixIT but uses initial mixtures as supervision to ensure training stability.

In contrast to the multi-stage training methods that refine MixIT pre-trained models, we propose a novel approach to train models from scratch using RemixIT and Self-Remixing. While refining pre-trained models has shown success, training from scratch with them has not been explored because of the anticipated high level of distortion in pseudo-mixtures generated by randomly initialized models, as well as the possibility of falling into trivial solutions that do not separate sources. However, our empirical findings suggest that pseudo-mixtures generated by randomly initialized models can be treated as MoMs, and the training can be done through MixIT-like optimization problems, indicating that RemixIT and Self-Remixing can be effective even when starting from scratch. Furthermore, we introduce a carefully designed remixing algorithm to overcome the trivial solution issue. At the start of training, there is a significant mismatch between pseudo-mixtures and actual mixtures since pseudo-mixtures become MoMs. However, such a mismatch gradually decreases as the teacher model is refined, leading to better performance than using static MoMs.

Our main contributions can be summarized as follows: *i*) We show how to train models from scratch using RemixIT and Self-Remixing, without the need for pre-training. Additionally, we propose a remixing method to enhance training stability. *ii*) We investigate the impact of the remixing algorithm on the separation performance, which has received little attention until now. Our analysis reveals that the design of the remixing algorithm plays a crucial role in the final separation performance, especially regarding the word error rate.

2. Methods

Our proposal is to train separation models from scratch in an unsupervised manner using RemixIT or Self-Remixing. We start by providing an overview of related methods, including MixIT [7] and mixture permutation invariant training (MixPIT) [15]. Then, we delve into the details of RemixIT and Self-Remixing and explain how they work from scratch. To enhance the training stability, we introduce a simple remixing method.

Let us denote a mini-batch of B mixtures as $\mathbf{x} \in \mathbb{R}^{B \times T}$, where T denotes the number of samples in the time domain. Each mixture in the batch, denoted as x_b ($b=1, \dots, B$), may contain up to K sources. To define a separator with N output channels, we use f_S with its parameters represented by θ_S .

2.1. MixIT and MixPIT

To generate a mixture of mixtures (MoM) \bar{x} , MixIT takes a mini-batch \mathbf{x} and adds sets B' mixtures together, denoted as $x_1, \dots, x_{B'}$. The resulting \bar{x} is then fed into f_S to produce separated signals $\hat{\mathbf{s}} \in \mathbb{R}^{N \times T}$ ($N \geq B'K$). The MixIT loss is

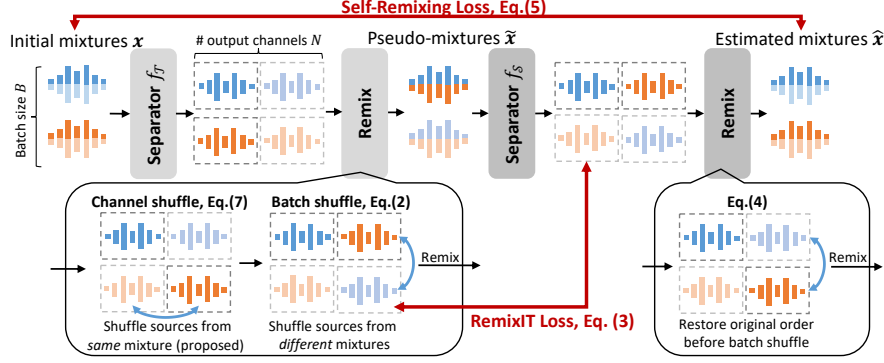


Figure 1: Overview of RemixIT and Self-Remixing. We first make pseudo-mixtures by separating mixtures with $f_{\mathcal{T}}$ and remixing outputs. We then train $f_{\mathcal{S}}$ to separate pseudo-mixtures using outputs of $f_{\mathcal{T}}$ or initial mixtures as supervision. Channel shuffle (bottom left) is introduced to stabilize Self-Remixing. Note that they can be applied regardless of B and N .

then computed between these separated signals and the individual mixtures, as described in [7]:

$$\mathcal{L}_{\text{MixIT}} = \min_{\mathbf{A}} \sum_{b'=1}^{B'} \mathcal{L}(x_{b'}, [\mathbf{A}\hat{\mathbf{s}}]_{b'}), \quad (1)$$

where \mathcal{L} is a loss function and $\mathbf{A} \in \mathbb{B}^{B' \times N}$ assigns each \hat{s}_n to the original mixtures. B' is typically set to two. While MixIT has made unsupervised monaural source separation feasible, the models often encounter an over-separation issue due to the fact that MoMs contain more sources than individual mixtures.

MixPIT is a recent development that aims to address the over-separation problem. MixPIT creates MoMs by summing the same number of mixtures as the number of sources in individual mixtures (K) and trains the models to estimate the individual mixtures directly (i.e., $B' = N = K$). Despite underperforming MixIT, as reported in [15], the success of MixPIT implies that RemixIT and Self-Remixing can also be used effectively from scratch (as detailed in Section 2.3).

2.2. RemixIT and Self-Remixing

An overview of RemixIT and Self-Remixing is provided in Fig. 1. RemixIT involves a teacher model, denoted as $f_{\mathcal{T}}$, that generates pseudo-mixtures from observed mixtures, and a student model, denoted as $f_{\mathcal{S}}$, that is trained to separate these pseudo-mixtures. The student model is trained via gradient descent using the teacher’s outputs as supervision, while the teacher’s parameters $\theta_{\mathcal{T}}$ are updated using the student’s parameters $\theta_{\mathcal{S}}$. Self-Remixing follows a similar training procedure as RemixIT but uses the initial mixtures as supervision.

The teacher $f_{\mathcal{T}}$ first separates the mixtures \mathbf{x} into sources: $\tilde{\mathbf{s}} = f_{\mathcal{T}}(\mathbf{x}; \theta_{\mathcal{T}}) \in \mathbb{R}^{B \times N \times T}$. To ensure that the sources add up to the initial mixtures, $\sum_{n=1}^N \tilde{s}_{b,n} = x_b$, we enforce mixture consistency (MC) [16]. Next, we shuffle the sources within a batch in each output channel n with an $B \times B$ permutation matrix $\mathbf{\Pi}_n$, and make pseudo-mixtures $\tilde{\mathbf{x}}$ by adding them together:

$$\tilde{x}_b = \sum_{n=1}^N \tilde{s}_{b,n}^{(\mathbf{\Pi})}, \quad \tilde{s}_{b,n}^{(\mathbf{\Pi})} = [\mathbf{\Pi}_n \tilde{\mathbf{s}}_n^{\top}]_b, \quad (2)$$

where we define $^{\top}$ as the transpose of the first and the second dimension ($\mathbb{R}^{B \times N \times T} \rightarrow \mathbb{R}^{N \times B \times T}$) and $\tilde{\mathbf{s}}_n^{\top} \triangleq [\tilde{s}_{1,n}, \dots, \tilde{s}_{B,n}]$. The process described in Eq. (2) is referred to as *batch shuffle*. After remixing, $f_{\mathcal{S}}$ separates the pseudo-mixtures: $\hat{\mathbf{s}} = f_{\mathcal{S}}(\tilde{\mathbf{x}}; \theta_{\mathcal{S}})$. The loss for RemixIT is computed in a permutation-invariant way [2] between the separated signals of $f_{\mathcal{T}}$ and $f_{\mathcal{S}}$:

$$\mathcal{L}_{\text{RemixIT}}^{(b)} = \min_{\mathbf{P}} \frac{1}{N} \sum_{n=1}^N \mathcal{L}(\tilde{s}_{b,n}^{(\mathbf{\Pi})}, [\mathbf{P}_b \hat{\mathbf{s}}_b]_n), \quad (3)$$

where \mathbf{P}_b is an $N \times N$ permutation matrix. We then align \hat{s}_b in the same order as $\tilde{s}_b^{(\mathbf{\Pi})}$ using the optimal permutation matrix

$\bar{\mathbf{P}}_b: \hat{s}_b^{(\bar{\mathbf{P}})} = \bar{\mathbf{P}}_b \hat{s}_b$. Next, we restore the order of sources to their original sequence before the batch shuffle in Eq. (2):

$$\hat{s}_{b,n}^{(\mathbf{\Pi}^{-1})} = [\mathbf{\Pi}_n^{-1} \hat{s}_n^{(\bar{\mathbf{P}})\top}]_b, \quad (4)$$

where $\hat{s}_n^{(\bar{\mathbf{P}})\top} \triangleq [\hat{s}_{1,n}^{(\bar{\mathbf{P}})}, \dots, \hat{s}_{B,n}^{(\bar{\mathbf{P}})}]$. The Self-Remixing loss is computed using the initial mixtures as supervision:

$$\mathcal{L}_{\text{Self-Remixing}}^{(b)} = \mathcal{L}\left(x_b, \sum_{n=1}^N \hat{s}_{b,n}^{(\mathbf{\Pi}^{-1})}\right). \quad (5)$$

While $f_{\mathcal{S}}$ is trained via gradient descent using Eqs.(3) or (5), the teacher’s parameters $\theta_{\mathcal{T}}$ are updated using the student’s parameters $\theta_{\mathcal{S}}$. In this work, we update $\theta_{\mathcal{T}}$ in the exponential moving average (EMA) fashion at every epoch end:

$$\theta_{\mathcal{T}}^{(j+1)} = \alpha \theta_{\mathcal{T}}^{(j)} + (1 - \alpha) \theta_{\mathcal{S}}^{(j)}, \quad (6)$$

where $\alpha \in [0, 1]$ and j is the epoch index. We set $\alpha = 0.8$.

2.3. Why do RemixIT/Self-Remixing work from scratch?

In this work, we demonstrate that RemixIT and Self-Remixing can work even when $f_{\mathcal{T}}$ is randomly initialized, contrary to the original assumption that $f_{\mathcal{T}}$ is pre-trained.

We analyzed the output of a model initialized randomly¹ using the Conformer architecture described in Section 3.2, and noisy two-speaker mixtures described in Section 3.1. Our empirical results indicated that the outputs after enforcing MC tended to be similar to the input mixture, albeit with different scales. Specifically, the average scale-invariant signal-to-distortion ratio (SISDR) [18] of the outputs compared to the input mixture was 18.4 dB, indicating that they were somewhat distorted, but the deviation from the initial mixtures was small.

Based on the assumption that the outputs of $f_{\mathcal{T}}$ can be considered as mixtures, the pseudo-mixtures created by remixing them can be treated as MoMs. Given this assumption, the loss function used in RemixIT is the same as MixPIT’s because RemixIT trains the models to separate MoMs into each mixture. In addition, Self-Remixing also promotes the separation of MoMs to reconstruct the initial mixtures. Therefore, the objectives of RemixIT and Self-Remixing are similar to MixPIT’s. These analyses suggest that RemixIT and Self-Remixing can be effective even when starting from random initialization.

Note that while our analysis was conducted using the Conformer architecture, depending on the architecture, there may be cases where the outputs of $f_{\mathcal{T}}$ deviate from the input mixtures and the above assumption does not hold. However, it is possible to ensure that the assumption holds, e.g., by initializing the output layer of $f_{\mathcal{T}}$ with zeros. This results in all outputs having the

¹We used the default initialization strategy of Pytorch 1.12.1 [17].

same values (e.g., simply zeros or 0.5 with the sigmoid activation) and the MC layer scales the sources to be $\frac{1}{N}$ mixtures, so pseudo-mixtures become MoMs. While models were randomly initialized in this work, we confirmed that RemixIT and Self-Remixing still worked when we initialized the final linear layer of the Conformer with zeros.

2.4. Design of Remixing

We introduce our proposed remixing algorithm which has been specifically designed to facilitate training from scratch.

Channel shuffle (CS): In the previous subsection, we discussed how RemixIT and Self-Remixing can operate without pre-training. However, we discovered that Self-Remixing can encounter a trivial solution in that sources are not separated at all because the Self-Remixing loss can be minimized even when the model does not separate sources. Specifically, the model outputs the input mixture from one output channel and zeros from the others, e.g., $\tilde{s}_{b,1} = x_b$ and $\tilde{s}_{b,n \neq 1} = 0$. Since the batch shuffle (as seen in Eq. (2)) is conducted in each output channel, the model may easily fall into such a trivial solution. To address this problem, we propose a solution called *channel shuffle* (CS), which shuffles sources that are separated from the same mixture before the batch shuffle:

$$\tilde{s}_b \leftarrow \mathbf{\Lambda}_b \tilde{s}_b \in \mathbb{R}^{N \times T}, \quad (7)$$

where $\mathbf{\Lambda}_b$ is an $N \times N$ permutation matrix. We show this simple method greatly improves the stability of Self-Remixing.

Avoiding remixing of sources from the same mixture: When the permutation matrix $\mathbf{\Pi}_n$ in Eq.(2) is randomly set, pseudo-mixtures sometimes contain the sources separated from the same mixture. However, such pseudo-mixtures interfere with training when applying RemixIT from scratch. Let us consider the case where $f_{\mathcal{T}}$ with $N = 3$ is randomly initialized and the pseudo-mixture is composed of the following three sources $\tilde{s}_1 = \frac{1}{3}x_1, \tilde{s}_2 = \frac{1}{3}x_2, \tilde{s}_3 = \frac{1}{3}x_2$ (i.e., $\tilde{x} = \frac{1}{3}x_1 + \frac{2}{3}x_2$). In this case, RemixIT loss encourages f_S to separate $\frac{2}{3}x_2$ into two $\frac{1}{3}x_2$; however, this objective does not necessarily improve separation performance. As a result, when applying RemixIT, we opt to select the permutation matrix $\mathbf{\Pi}_n$ under the condition that sources from the same mixture are not remixes. Note that Self-Remixing is still effective even when the sources from the same mixture are remixes because it encourages f_S to separate \tilde{x} into $\frac{1}{3}x_1$ and $\frac{2}{3}x_2$ to reconstruct initial mixtures, x_1 and x_2 .

3. Experiments

3.1. Datasets

WSJ-mix: We synthesized two-speaker mixtures with reverberation and noise using speeches from WSJ0 [19] and WSJ1 [20] and noises from CHiME3 [21] at a sampling rate of 8kHz. The configuration was similar to that of SMS-WSJ [22] but with changes in reverberation times, and noise type and level. Pyroomacoustics [4] was used for simulation, and the reverberation times were chosen from 0.2 s to 1.0 s. Additionally, all source image locations were jittered by up to 8cm to avoid the sweeping echo problem [23]. Note that the condition is different from our previous paper [14]. The SNR of the noise ranged between 10 dB to 20 dB. The dataset consisted of 33561 (~87.4h), 982 (~2.5h), and 1332 (~3.6h) mixtures for the training, validation, and test sets, respectively. To evaluate word error rates (WERs), we used the ASR backend provided in [22]. **FUSS:** We utilized the anechoic version of the free universal sound separation (FUSS) dataset [24] to evaluate the performance in USS tasks [25]. The dataset comprised mixtures with one to four sources, drawn from 357 classes of audio sources.

All mixtures were ten seconds long and sampled at 16kHz. The training, validation, and test sets contained 20000, 1000, and 1000 mixtures, respectively.

3.2. Separation model

As the separation model, we utilized Conformer [26] (implemented based on [27]). The model contained about 21.6M parameters and was composed of 16 Conformer encoder layers with four attention heads, 256 attention dimensions, and 1024 feed-forward network dimensions. We replaced the batch normalization [28] with the group normalization [29] with eight groups. It took log-magnitude (or magnitude) spectrograms in the STFT domain as inputs and produced real-valued TF masks in WSJ-mix (FUSS). The FFT size was 512 and the window size and the hop length were 400 and 160, respectively.

3.3. Compared methods

We compared the performance of several methods, including **MixIT**, **MixIT+Sparsity**, **RemixIT**, **Self-Remixing**, and their supervised counterparts, **Sup. RemixIT** and **Sup. Self-Remixing**. MixIT is implemented based on [30]. In MixIT and MixIT+Sparsity, N was set to six in WSJ-mix and eight in FUSS, and we ensured MC. In the other methods, N was set to three in WSJ-mix and four in FUSS. MixIT+Sparsity uses the sparsity loss (Eq. (6) in [8]) that promotes sparsity of the outputs and prevents over-separation. Note that the sparsity loss does not necessarily improve separation performance, as demonstrated in [8]. We trained RemixIT and Self-Remixing from scratch using Eqs. (3) and (5), respectively. For the supervised versions, we used ground-truth signals instead of the outputs of $f_{\mathcal{T}}$ to create pseudo-mixtures.

We excluded MixPIT from the baseline because prior studies have demonstrated that MixIT outperforms MixPIT [15]. Additionally, we did not consider remix-cycle-consistent learning [31, 32], which is similar to Self-Remixing but shares parameters between $\theta_{\mathcal{T}}$ and θ_S and computes the gradient through two separation and remixing processes because we did not achieve good results when training it from scratch.

3.4. Training details

As the signal-level loss function \mathcal{L} , we used the negative thresholded SNR between the reference y and the estimate \hat{y} :

$$\mathcal{L}(y, \hat{y}) = -10 \log_{10} \frac{\|y\|^2}{\|y - \hat{y}\|^2 + \tau \|y\|^2}, \quad (8)$$

where $\tau = 10^{-3}$ is a threshold that clamps the SNR at 30 dB. In supervised methods, we used the SNR loss that can handle zero-references (Eq. (2) of [24]), instead of Eq. (8).

When training with WSJ-mix, the batch size was 32 and the input was 7 seconds long. Models were trained for 600 epochs. When training with FUSS, the batch size was 16 and the input was 10 seconds long. Models were trained for 400 epochs. In both experiments, we used the AdamW optimizer [33] with the weight decay of $1e-2$. The learning rate was increased linearly from 0 to $2e-4$ over the first 5000 training steps, kept constant at $2e-4$ for 100 epochs, and then decayed by 0.98 for every three epochs until it reached $2e-5$. The gradients were clipped with a maximum norm of 5. In MixIT+Sparsity, we initially trained the model for 450 (300) epochs using only the MixIT loss. We then fine-tuned the model with the sparsity loss for 150 (100) epochs in WSJ-mix (FUSS), with the weight of sparsity loss set to 4 (23). In supervised methods, all-zero pseudo-mixtures were occasionally generated, but we did not use them for training.

To normalize each mixture, we subtracted its mean and di-

Table 1: Evaluation results on WSJ-mix test set.

Method	SISDR	STOI	PESQ	WER
Unprocessed	-0.4	0.684	1.82	82.9%
A1 MixIT	8.8	0.847	2.54	42.3%
A2 + Sparsity	8.6	0.843	2.51	44.8%
A3 RemixIT w/o CS	10.8	0.890	2.84	47.4%
A4 RemixIT	10.3	0.878	2.75	43.3%
A5 Self-Remixing	10.3	0.877	2.69	50.1%
A6 Self-Remixing [†]	10.3	0.878	2.74	39.7%
Sup. RemixIT w/o CS	10.9	0.896	3.01	30.9%
Sup. RemixIT	10.6	0.889	2.93	33.4%
Sup. Self-Remixing	10.6	0.889	2.93	33.7%




Figure 2: Validation SISDR of f_S of RemixIT and Self-Remixing in early stage of training.

vided it by its standard deviation. For MoM \bar{x} in MixIT, we followed the same process. Although [14] reported the instability of RemixIT in source separation tasks, we found that normalization improves the stability of the model. As a result, we successfully trained models on reverberant noisy two-speaker mixtures similar to the dataset used in [14]. It is worth noting that we only normalized the initial mixtures x and did not normalize the pseudo-mixtures \tilde{x} because normalizing \tilde{x} resulted in lower performance in our experiments.

3.5. Results on WSJ-mix

The results for SISDR [34], short-time objective intelligibility (STOI) [35], perceptual evaluation of speech quality (PESQ) [36], and WER on WSJ-mix are shown in Table 1. To evaluate the model, we used the averaged parameters from five checkpoints that yielded the highest SISDR on the validation set. Note that [†] denotes instances where we allowed remixing of sources from the same mixture (e.g., Self-Remixing[†] in A6).

The results show that both RemixIT and Self-Remixing performed well from scratch and outperformed MixIT in terms of speech metrics. However, the WERs of A3–A5 were worse than that of MixIT. As discussed in [14], remixing-based methods are trained with distorted pseudo-mixtures, which can result in models outputting distorted signals. This implies that the distortion caused by using pseudo-mixtures is more harmful to speech recognition than that caused by using MoMs for training. Despite this, the substantial improvement in the WER from A5 to A6 suggests that allowing remixing of sources from the same mixture can alleviate distortion in separated signals.

Fig. 2 displays the validation SISDR in the early stage of training. It has been confirmed that CS stabilizes Self-Remixing. This outcome, along with the observed WER improvement in A6, highlights the significance of the remixing algorithm. It impacts both the final separation performance and training stability. In future work, we plan to explore a more thoughtfully designed remixing approach to enhance the performance of both Self-Remixing and RemixIT.

3.6. Results on FUSS

The evaluation results on the FUSS test set are listed in Table 2, which includes the following evaluation metrics [8]: **1S**, repre-

Table 2: Evaluation results on FUSS test set. Checkpoints that gave best MSi (B^*) and TRF (C^*) were evaluated.

Method	1S	2Si	3Si	4Si	MSi	TRF
B1 MixIT	9.5	9.8	14.4	15.9	13.2	12.2
B2 + Sparsity	15.2	12.3	15.5	16.6	14.7	14.8
B3 RemixIT	42.1	12.3	11.7	8.2	10.8	19.1
B4 Self-Remixing	29.4	14.7	16.2	13.6	14.8	18.7
B5 Self-Remixing [†]	31.9	14.9	15.9	13.7	14.8	19.4
Sup. Self-Remixing	32.9	14.1	15.8	13.5	14.5	19.4
C1 MixIT	9.5	9.9	14.3	15.8	13.2	12.2
C2 + Sparsity	18.0	12.3	15.5	16.5	14.7	15.6
C3 RemixIT	63.6	4.6	2.1	0.1	2.4	18.6
C4 Self-Remixing	34.4	14.8	15.9	13.1	14.6	19.8
C5 Self-Remixing [†]	34.9	15.0	15.6	13.4	14.7	20.0
Sup. Self-Remixing	49.2	13.7	14.7	12.0	13.5	23.0

sented the SISDR for single-source mixtures; kSi , representing the SISDR improvement (SISDRi) for k -source mixtures ($k = 2, 3, 4$); **MSi**, representing the average SISDRi for multi-source mixtures; and **TRF**, representing the total reconstruction fidelity with average performance on all test data. For evaluation, we utilized the averaged model parameters from the five checkpoints that resulted in the highest MSi (B^*) and TRF (C^*) on the validation set. It is worth noting that, unlike [24], we did not discard estimate-reference pairs with low-energy estimated sources following [37], as we believe that errors in low-energy estimates should be taken into account.

The results indicate that Self-Remixing achieved higher TRF than MixIT+Sparsity while providing almost the same MSi. Unlike MixIT, which suffers from low 1S due to the inclusion of at least two sources in MoMs, Self-Remixing attains high performance in both 1S and MSi by utilizing both single-source and multi-source pseudo-mixtures.

In RemixIT, the separation performance initially improved, but then it deteriorated. This can be attributed to the fact that many mixtures have fewer sources than the number of output channels, and as a result, f_T often outputs sources that are close to zero. In such a case, the RemixIT loss encourages f_S to output zeros, leading to *under-separation*. As the teacher’s weights are updated with the student’s weights, under-separation occurs more frequently in f_T , resulting in high 1S but low MSi.

Interestingly, the unsupervised Self-Remixing resulted in slightly higher MSi than the supervised one. We attribute this to the over-separation of f_T . Specifically, when f_T separates three-source mixtures, for instance, over-separation can lead to the output of four sources instead of the expected three sources and one zero signal. As a result, the average number of sources contained in pseudo-mixtures becomes higher than that of supervised learning, leading to higher MSi and lower 1S.

4. Conclusion

We proposed to train separators from scratch with RemixIT and Self-Remixing. We showed that their training objectives are similar to MixPIT’s at the start of the training and thus they are capable of training without any pre-training. The experiments demonstrated the effectiveness of RemixIT and Self-Remixing even when training from scratch, as well as the importance of remixing method. In the future, we will explore the applicability of the proposed method to other separation models [3, 38].

5. Acknowledgments

The research was supported by NII CRIS collaborative research program operated by NII CRIS and LINE Corporation.

6. References

- [1] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. ICASSP*, 2016, pp. 31–35.
- [2] 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, pp. 241–245.
- [3] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [4] R. Scheibler, E. Bezzam, and I. Dokmanić, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," in *Proc. ICASSP*, 2018, pp. 351–355.
- [5] L. Drude, D. Hasenklever, and R. Haeb-Umbach, "Unsupervised training of a deep clustering model for multichannel blind source separation," in *Proc. ICASSP*, 2019, pp. 695–699.
- [6] K. Saijo and R. Scheibler, "Spatial loss for unsupervised multi-channel source separation," in *Proc. Interspeech*, 2022, pp. 241–245.
- [7] S. Wisdom, E. Tzinis, H. Erdogan, R. Weiss, K. Wilson, and J. Hershey, "Unsupervised sound separation using mixture invariant training," *Advances in Neural Information Processing Systems*, vol. 33, pp. 3846–3857, 2020.
- [8] S. Wisdom, A. Jansen, R. J. Weiss, H. Erdogan, and J. R. Hershey, "Sparse, efficient, and semantic mixture invariant training: Taming in-the-wild unsupervised sound separation," in *Proc. WASPAA*, 2021, pp. 51–55.
- [9] T. Denton, S. Wisdom, and J. R. Hershey, "Improving bird classification with unsupervised sound separation," in *Proc. ICASSP*, 2022, pp. 636–640.
- [10] A. Sivaraman, S. Wisdom, H. Erdogan, and J. R. Hershey, "Adapting speech separation to real-world meetings using mixture invariant training," in *Proc. ICASSP*. IEEE, 2022, pp. 686–690.
- [11] J. Zhang, C. Zorilá, R. Doddipatla, and J. Barker, "Teacher-Student MixIT for Unsupervised and Semi-Supervised Speech Separation," in *Proc. Interspeech*, 2021, pp. 3495–3499.
- [12] E. Tzinis, Y. Adi, V. K. Ithapu, B. Xu, and A. Kumar, "Continual self-training with bootstrapped remixing for speech enhancement," in *Proc. ICASSP*. IEEE, 2022, pp. 6947–6951.
- [13] 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 Journal of Selected Topics in Signal Processing*, 2022.
- [14] K. Saijo and T. Ogawa, "Self-remixing: Unsupervised speech separation via separation and remixing," *arXiv preprint arXiv:2211.10194*, 2022.
- [15] E. Karamathı and S. Kirbız, "Mixcycle: Unsupervised speech separation via cyclic mixture permutation invariant training," *IEEE Signal Processing Letters*, 2022.
- [16] S. Wisdom, J. R. Hershey, K. Wilson, J. Thorpe, M. Chinen, B. Patton, and R. A. Saurous, "Differentiable consistency constraints for improved deep speech enhancement," in *Proc. ICASSP*. IEEE, 2019, pp. 900–904.
- [17] A. Paszke *et al.*, "Pytorch: An imperative style, high-performance deep learning library," *Advances in neural information processing systems*, vol. 32, 2019.
- [18] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr-half-baked or well done?" in *Proc. ICASSP*, 2019, pp. 626–630.
- [19] J. S. Garofolo *et al.*, *CSR-I (WSJ0) Complete LDC93S6A*, Linguistic Data Consortium, Philadelphia, 1993, web Download.
- [20] Linguistic Data Consortium, and NIST Multimodal Information Group, *CSR-II (WSJ1) Complete LDC94S13A*, Linguistic Data Consortium, Philadelphia, 1994, web Download.
- [21] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third 'CHiME' speech separation and recognition challenge: Dataset, task and baselines," in *Proc. ASRU*, 2015, pp. 504–511.
- [22] 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," in *arXiv preprint arXiv:1910.13934*, 2019.
- [23] E. De Sena, N. Antonello, M. Moonen, and T. van Waterschoot, "On the modeling of rectangular geometries in room acoustic simulations," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 4, pp. 774–786, 2015.
- [24] S. Wisdom, H. Erdogan, D. P. Ellis, R. Serizel, N. Turpault, E. Fonseca, J. Salamon, P. Seetharaman, and J. R. Hershey, "What's all the fuss about free universal sound separation data?" in *Proc. ICASSP*. IEEE, 2021, pp. 186–190.
- [25] I. Kavalerov, S. Wisdom, H. Erdogan, B. Patton, K. Wilson, J. Le Roux, and J. R. Hershey, "Universal sound separation," in *Proc. WASPAA*, 2019, pp. 175–179.
- [26] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented Transformer for Speech Recognition," in *Proc. Interspeech*, 2020, pp. 5036–5040.
- [27] 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. ICASSP*, 2021, pp. 5749–5753.
- [28] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. ICML*, 2015, pp. 448–456.
- [29] Y. Wu and K. He, "Group normalization," in *Proc. ECCV*, 2018, pp. 3–19.
- [30] M. Pariente *et al.*, "Asteroid: The PyTorch-Based Audio Source Separation Toolkit for Researchers," in *Proc. Interspeech*, 2020, pp. 2637–2641.
- [31] K. Saijo and T. Ogawa, "Remix-cycle-consistent learning on adversarially learned separator for accurate and stable unsupervised speech separation," in *Proc. ICASSP*, 2022.
- [32] —, "Unsupervised training of sequential neural beamformer using coarsely-separated and non-separated signals," in *Proc. Interspeech*, 2022, pp. 251–255.
- [33] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *Proc. ICLR*, 2018.
- [34] R. Scheibler, "SDR — Medium rare with fast computations," in *Proc. ICASSP*, May 2022.
- [35] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [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, pp. 749–752.
- [37] E. Tzinis, Z. Wang, X. Jiang, and P. Smaragdis, "Compute and memory efficient universal sound source separation," *Journal of Signal Processing Systems*, vol. 94, no. 2, pp. 245–259, 2022.
- [38] 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," *arXiv preprint arXiv:2211.12433*, 2022.