

## Original Paper

# Analysis and Extension of Noisy-target Training for Unsupervised Target Signal Enhancement

Takuya Fujimura<sup>1\*</sup> and Tomoki Toda<sup>2</sup>

<sup>1</sup>*Graduate School of Informatics, Nagoya University, Nagoya, Japan*

<sup>2</sup>*Information Technology Center, Nagoya University, Nagoya, Japan*

---

### ABSTRACT

Deep neural network-based target signal enhancement (TSE) is usually trained in a supervised manner using clean target signals. However, collecting clean target signals is costly and such signals are not always available. Thus, it is desirable to develop an unsupervised method that does not rely on clean target signals. Among various studies on unsupervised TSE methods, Noisy-target Training (NyTT) has been established as a fundamental method. NyTT simply replaces clean target signals with noisy ones in the typical supervised training, and it has been experimentally shown to achieve TSE. Despite its effectiveness and simplicity, its mechanism and detailed behavior are still unclear. In this paper, to advance NyTT and, thus, unsupervised methods as a whole, we analyze NyTT from various perspectives. We experimentally demonstrate the mechanism of NyTT, the desirable conditions, and the effectiveness of utilizing noisy signals in situations where a small number of clean target signals are available. Furthermore, we propose an improved version of NyTT based on its properties and explore its capabilities in the dereverberation and declipping tasks, beyond the denoising task.

---

\*Corresponding author: [fujimura.takuya@g.sp.m.is.nagoya-u.ac.jp](mailto:fujimura.takuya@g.sp.m.is.nagoya-u.ac.jp)

*Keywords:* Target signal enhancement, unsupervised learning, denoising, dereverberation, declipping

## 1 Introduction

Target signal enhancement (TSE) is a technique to extract a target signal from a noisy observation. In various speech communication systems, such as online meetings, hearing aids, and automatic speech recognition (ASR) systems, this technique has been employed to extract human speech [41, 66, 27, 11]. It has also been applied to various types of target signal beyond speech, including music [17, 6, 45] and environmental sounds [25, 54, 14]. This TSE technique can be classified into multi-channel and single-channel methods. Multi-channel methods extract the target signal by leveraging the spatial information obtained from multiple microphones [7, 59, 58]. However, the physical size of the microphone array can sometimes limit its application. Consequently, single-channel methods, which use a single microphone and perform TSE based on differences in acoustic features between the target signal and other noise signals, also play a crucial role in TSE applications. Classical signal processing-based single-channel methods [46, 40] have been widely adopted owing to their simplicity and low computational cost; however, their enhancement performance is often insufficient. In contrast, recent single-channel TSE methods have achieved significant performance improvements by incorporating deep neural networks (DNNs) [62, 36, 50, 28, 61, 45].

Most single-channel DNN-based TSE methods rely on supervised learning with clean target signals, which we refer to as Clean-target Training (CTT) (Figure 1(a)). In CTT, we input a noisy signal into a DNN and train it to predict the corresponding clean target signal. CTT is an appropriate strategy, and various improvements have been made, including modifications to model architectures (e.g., convolutional networks [36], recurrent networks [35], and Transformers [50, 28]), loss functions (e.g., mean-squared-error (MSE) and signal-to-noise ratio (SNR) [9]), and signal representations to which TSE is applied (e.g., amplitude spectrograms [34], complex spectrograms [62, 61], waveforms [44, 36, 50], and both spectrograms and waveforms [45]). Furthermore, memory-efficient [56, 51] and real-time [4] architectures have also been explored. Despite these improvements, CTT still has one major problem: collecting clean target signals is costly. Typically, clean target signals are recorded in controlled settings, such as an anechoic chamber, to prevent degradation from environmental noise and reverberation. Consequently, the recording process is costly and time-consuming, limiting the amount of training data. Moreover, although it is theoretically possible to achieve the TSE of any target signals, such as animals and vehicle sounds, it is often not feasible to record such clean target signals. Therefore, the types of target signal used in CTT are realistically limited.

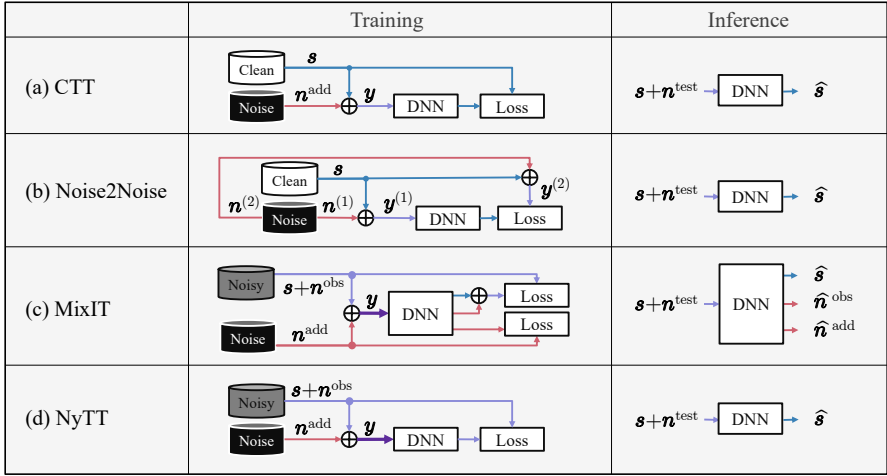


Figure 1: Comparison of NyTT and its related methods. Blue, red, and purple lines represent clean target signals, noise signals, and mixtures of target and noise signals, respectively. Deep purple lines represent *more noisy* signals. The loss function receives estimated signals from DNNs and computes a reconstruction loss between the estimates and the target signals.

To alleviate this limitation, unsupervised<sup>1</sup> TSE methods have been studied [64, 22, 12, 24, 1, 13]. PULSE [22] is an unsupervised TSE method based on positive-unlabeled (PU) learning [8] and uses noisy target signals and additional noise signals for its training. PU learning is a machine learning technique that enables the classification of positive and negative examples using positive and unlabeled training data. On the basis of this technique, PULSE classifies local patches of amplitude spectrograms into noise (positive) or target signal (negative) components. For training, patches of noise signals are used as positive data, while patches of noisy signals are used as unlabeled data since a noisy signal contains both noise and target signal patches. During inference, PULSE performs TSE by applying a mask to the input amplitude spectrogram, filtering out noise (positive) patches. Another method, MetricGAN-U, is based on a generative adversarial network (GAN) and uses only noisy signals for its training. In MetricGAN-U, the discriminator is trained to predict a signal quality metric, whereas the generator is trained to maximize the evaluation from the discriminator. This achieves the unsupervised TSE by employing a non-intrusive metric, which does not use a clean target signal as a reference, as the metric that the discriminator mimics. Although PULSE and MetricGAN-U have demonstrated their TSE capabilities,

<sup>1</sup>In the context of TSE, methods that do not require clean target signals are considered unsupervised, even if the training is performed in a supervised manner [64].

they cannot directly inherit advancements made in the CTT framework owing to their specialized training algorithms. For example, PULSE restricts TSE models to the time–frequency (T–F) masking approach because it relies on the classification of spectrogram patches, even though time-domain models have also been developed within the CTT framework [36, 50, 28, 45]. Moreover, MetricGAN-U requires a non-intrusive evaluation metric for its training, but it is not always available. Although a pre-trained DNN-based evaluation metric predictor can be employed as the non-intrusive metric, as demonstrated in the experiments in [12], constructing this predictor still requires clean target signals. Thus, it does not serve as an essential solution, especially when developing TSE systems for new types of target signal.

In contrast to the aforementioned unsupervised TSE methods, another approach utilizes noisy signals in the same training algorithms as CTT [64, 24, 13]. Noisy-target Training (NyTT) [13] is the basic method in this approach (Figure 1(d)). NyTT utilizes a noisy signal as the target signal instead of a clean one. It trains a DNN to predict the noisy target from a signal synthesized by mixing the noisy target with additional noise. During inference, the enhanced signal can be obtained directly from the DNN by inputting an unprocessed noisy signal. NyTT has been experimentally shown to achieve TSE without clean target signals and has the same training algorithm as CTT, which allows us to easily inherit advancements made in the CTT framework. However, the exact mechanism and desirable conditions of NyTT have not been clarified, hindering further advancements.

In this paper, we aim to advance the field of unsupervised TSE by analyzing the fundamental method, NyTT, from various perspectives and deepening our understanding of NyTT.<sup>2</sup> Through this analysis, we clarify 1) the mechanism of NyTT, 2) the desirable conditions of NyTT, and 3) the effectiveness of utilizing noisy signals in situations where a small number of clean target signals are available. Additionally, 4) we propose an improved version of NyTT based on its properties, demonstrating its potential to achieve performance comparable to CTT by iteratively improving the quality of the noisy target signals. Finally, 5) we demonstrate that NyTT can also handle dereverberation and declipping tasks, inheriting the broad applicability of CTT.

The rest of this paper is organized as follows: Section 2 provides details of NyTT and its closely related works. In Section 3, we outline the motivation and contents of our analyses. In Sections 4, 5, and 6, we present experimental results in denoising, dereverberation, and declipping tasks, respectively. Finally, in Section 7, we conclude this paper.

---

<sup>2</sup>This paper is an extension of our previous paper [15]. Compared with the previous paper, this paper provides a more thorough discussion of related works, a more comprehensive analysis of the desirable conditions for NyTT, and an investigation of its capability in dereverberation and declipping tasks, employing multiple evaluation metrics.

## 2 NyTT and its Related Works

### 2.1 Noise2Noise

Noise2Noise [32] is an unsupervised training method originally proposed for image denoising (Figure 1(b)). In Noise2Noise, pairs of noisy signals ( $\mathbf{y}^{(1)} = \mathbf{s} + \mathbf{n}^{(1)}, \mathbf{y}^{(2)} = \mathbf{s} + \mathbf{n}^{(2)}$ ) are used as training data, where the two noisy signals  $\mathbf{y}^{(1)}$  and  $\mathbf{y}^{(2)}$  share the same clean target signal  $\mathbf{s}$  but have different noise components  $\mathbf{n}^{(1)}$  and  $\mathbf{n}^{(2)}$ . A DNN  $f(\cdot)$  is trained to minimize the following prediction error:

$$\mathcal{L}^{\text{N2N}} = \mathbb{E}_{(\mathbf{y}^{(1)}, \mathbf{y}^{(2)}) \sim \mathcal{D}} [L(f(\mathbf{y}^{(1)}; \theta), \mathbf{y}^{(2)})], \quad (1)$$

where  $\mathbb{E}[\cdot]$  is the expectation operator,  $\mathcal{D}$  is a training dataset,  $L(\cdot)$  is a loss function, and  $\theta$  is the set of parameters of the DNN  $f(\cdot)$ . Here, Noise2Noise trains the DNN to acquire one-to-one mappings between the two noisy signals  $\mathbf{y}^{(1)}$  and  $\mathbf{y}^{(2)}$ . However, multiple plausible outputs can exist for a given input, especially when there is no consistent relationship between the two noise signals  $\mathbf{n}^{(1)}$  and  $\mathbf{n}^{(2)}$ . In such a case, if the loss function is MSE, the optimal solution becomes the average of the plausible candidates. For example, we consider the optimal output  $\hat{\mathbf{y}}^{(2)}$  for a given input  $\mathbf{y}^{(1)}$ , when using MSE as the loss function. Here, the training objective is to minimize the following  $\mathcal{L}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}}^{\text{N2N}}$ :

$$\mathcal{L}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}}^{\text{N2N}} = \mathbb{E}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}} [L(\hat{\mathbf{y}}^{(2)}, \mathbf{y}^{(2)})]. \quad (2)$$

Therefore,  $\hat{\mathbf{y}}^{(2)}$  is obtained as

$$\frac{\partial}{\partial \hat{\mathbf{y}}^{(2)}} \mathcal{L}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}}^{\text{N2N}} = 0, \quad (3)$$

$$\frac{\partial}{\partial \hat{\mathbf{y}}^{(2)}} \mathbb{E}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}} [\|\hat{\mathbf{y}}^{(2)} - \mathbf{y}^{(2)}\|_2^2] = 0, \quad (4)$$

$$\hat{\mathbf{y}}^{(2)} = \mathbb{E}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}} [\mathbf{y}^{(2)}]. \quad (5)$$

This averaging effect is observed as a problem of a blurred output in super-resolution and a greyish output in autocoloring [31, 68, 21]. On the basis of this property, Noise2Noise can achieve the same denoising training as CTT without requiring clean target signals since the averaging effect can remove zero-mean noise in the output signal (i.e.,  $\hat{\mathbf{y}}^{(2)} = \mathbb{E}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}} [\mathbf{s} + \mathbf{n}^{(2)}] = \mathbb{E}_{\mathbf{y}^{(2)}|\mathbf{y}^{(1)}} [\mathbf{s}]$ ).

It has been theoretically and experimentally proven that Noise2Noise can achieve denoising training using only noisy signals. In the case of images, it is relatively easy to obtain pairs of noisy signals,  $(\mathbf{y}^{(1)}, \mathbf{y}^{(2)})$ , that share the same clean image  $\mathbf{s}$  by taking consecutive shots when the subject is static. In

this case, Noise2Noise is a useful technique. However, in the case of audio, it is not possible to naturally obtain such pairs of noisy signals. Instead, we must synthesize them using a clean audio signal  $\mathbf{s}$ . Therefore, in the audio TSE task, Noise2Noise does not serve as an essential solution [24, 1].

## 2.2 MixIT

MixIT [64] is an unsupervised training method for sound source separation, and it can also be used for TSE (Figure 1(c)). In MixIT, training is conducted using noisy signals  $\mathbf{x} = \mathbf{s} + \mathbf{n}^{\text{obs}}$  and additional noise signals  $\mathbf{n}^{\text{add}}$ , where  $\mathbf{n}^{\text{obs}}$  represents the noise already present in  $\mathbf{x}$  at the time of observation. MixIT trains a DNN to minimize the following prediction error  $\mathcal{L}^{\text{MixIT}}$ :

$$\mathcal{L}^{\text{MixIT}} = \mathbb{E}_{(\mathbf{x}, \mathbf{n}^{\text{add}}) \sim \mathcal{D}} [\min(\mathcal{L}_{\text{MixIT1}}, \mathcal{L}_{\text{MixIT2}})], \quad (6)$$

$$\mathcal{L}_{\text{MixIT1}} = L(\mathbf{u}_1 + \mathbf{u}_2, \mathbf{x}) + L(\mathbf{u}_3, \mathbf{n}^{\text{add}}), \quad (7)$$

$$\mathcal{L}_{\text{MixIT2}} = L(\mathbf{u}_1 + \mathbf{u}_3, \mathbf{x}) + L(\mathbf{u}_2, \mathbf{n}^{\text{add}}). \quad (8)$$

Here,  $\mathbf{u}_1$ ,  $\mathbf{u}_2$ , and  $\mathbf{u}_3$  are the outputs of the DNN  $f(\mathbf{y}; \theta)$ , where  $\mathbf{y} = \mathbf{x} + \mathbf{n}^{\text{add}}$ . MixIT can achieve TSE because  $(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3) = (\mathbf{s}, \mathbf{n}^{\text{obs}}, \mathbf{n}^{\text{add}})$  or  $(\mathbf{s}, \mathbf{n}^{\text{add}}, \mathbf{n}^{\text{obs}})$  is the optimal solution for  $\mathcal{L}^{\text{MixIT}}$ . Although  $\mathcal{L}^{\text{MixIT}}$  can also be minimized by outputting a noisy signal as  $\mathbf{u}_1$  (i.e.,  $(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3) = (\mathbf{x}, \mathbf{0}, \mathbf{n}^{\text{add}})$ ), this issue can be avoided under the assumption that  $\mathbf{y}$  does not provide any information indicating that  $\mathbf{x}$  consists of  $\mathbf{s}$  and  $\mathbf{n}^{\text{obs}}$ . Under this assumption, the DNN cannot estimate a pair of signals that compose  $\mathbf{x}$  and, therefore, cannot always accurately estimate  $\mathbf{x}$ . Consequently, the DNN is trained to separate  $\mathbf{y}$  into individual sources, as this always minimizes  $\mathcal{L}^{\text{MixIT}}$ . When this assumption does not hold (e.g., when the characteristics of  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  differ and are distinguishable), it has been observed that MixIT suffers from performance degradation [48, 38].

MixIT is one of the major unsupervised TSE methods, and several improvements have been proposed. For instance, one method mitigated the overseparation problem by introducing a penalty term for the number of active sources and the correlation between the output sources [63]. Other methods produced better separation by using a pre-trained classification model (e.g., an audio event classification or an ASR model) [63, 53] or by employing a loss function that relaxes the training difficulty [48, 38]. Furthermore, the teacher–student learning approach has also been adopted [67, 55, 23, 47, 33]. In this approach, the student model is trained using the outputs of the teacher model pre-trained by MixIT as the pseudo-target signals.

### 2.3 NyTT

NyTT [13] is an unsupervised training method designed for TSE (Figure 1(d)). In NyTT, noisy signals  $\mathbf{x}$  and additional noise signals  $\mathbf{n}^{\text{add}}$  are used as training data, and a *more noisy* signal  $\mathbf{y}$  is generated as  $\mathbf{y} = \mathbf{x} + \mathbf{n}^{\text{add}}$ . NyTT trains a DNN to minimize the following prediction error  $\mathcal{L}^{\text{NyTT}}$ :

$$\mathcal{L}^{\text{NyTT}} = \mathbb{E}_{(\mathbf{y}, \mathbf{x}) \sim \mathcal{D}} [L(f(\mathbf{y}; \theta), \mathbf{x})]. \quad (9)$$

NyTT was inspired by Noise2Noise and realizes Noise2Noise training in the TSE task by considering  $\mathbf{y} = \mathbf{s} + (\mathbf{n}^{\text{obs}} + \mathbf{n}^{\text{add}}) = \mathbf{s} + \mathbf{n}^{(1)}$  and  $\mathbf{x} = \mathbf{s} + \mathbf{n}^{\text{obs}} = \mathbf{s} + \mathbf{n}^{(2)}$  as pairs of noisy signals. The prediction error is calculated using MSE in the time domain, under the assumption that the noise has a zero-mean distribution. Despite the lack of theoretical proof, NyTT has been experimentally demonstrated to achieve TSE without clean target signals [13]. Similar to the developments in MixIT, variants incorporating teacher–student learning into NyTT framework have also been proposed to improve the performance under out-of-domain noise conditions [3, 60].

Although NyTT and MixIT stem from different conceptual foundations, their resulting training algorithms are similar. The primary difference is which output is selected as the enhanced signal during the inference. MixIT involves a separation task, where  $\mathbf{y}$  is separated into several sound sources, and one of them is selected as the enhanced signal. In contrast, NyTT has only one slot and uses the output as the enhanced signal. Therefore, NyTT can be viewed as the simplified version of MixIT, where the DNN is trained with only  $L(\mathbf{u}_1 + \mathbf{u}_2, \mathbf{x})$  and the enhanced signal is created by  $\mathbf{u}_1 + \mathbf{u}_2$  during the inference. Both MixIT and NyTT are major unsupervised TSE methods, and they provide a greater flexibility than specialized unsupervised training algorithms. However, as mentioned above, NyTT has a simpler and more flexible architecture than MixIT. The simple architecture enables us to make improvements and expansions easily, and it is more suitable for analysis. For these reasons, we have chosen NyTT as the target of our analysis.

One limitation of NyTT is that it requires the zero-mean noise assumption and the use of MSE as conditions for Noise2Noise. Conversely, if NyTT does not require these conditions, it can be easily applied to wider ranges of tasks and loss functions. In Section 4.2, we demonstrate that NyTT indeed works without these conditions, and in Sections 5 and 6, we further show its effectiveness in dereverberation and declipping tasks.

### 3 Motivation and Content of the Investigation

#### 3.1 Validity of the Interpretation of NyTT

NyTT has been proposed, inspired by Noise2Noise, which utilizes noisy signals as target signals on the basis of the averaging effect of MSE loss function. On the other hand, NyTT can also be interpreted as being trained to remove  $\mathbf{n}^{\text{add}}$  from the *more noisy* signal  $\mathbf{y}$ , performing TSE by removing noise components corresponding to  $\mathbf{n}^{\text{add}}$  from the noisy input signal. Therefore, we investigate whether NyTT can be interpreted as Noise2Noise or not, through 1) the analysis of the signals processed in NyTT (Section 4.2.1) and 2) the evaluation of NyTT with a loss function that does not satisfy the conditions of Noise2Noise (Section 4.2.2). If NyTT is not Noise2Noise, the zero-mean noise assumption and the use of MSE will no longer be necessary, allowing us to use various loss functions and apply NyTT to various tasks.

#### 3.2 Improvement of NyTT Through Iteration

It has been shown that the performance of NyTT improves as the SNR of the noisy target increases [13]. On the basis of this property, we propose IterNyTT, which achieves better performance through an iterative process (Figure 2). In the first iteration of IterNyTT, we train a DNN  $f_1(\cdot)$  using NyTT with the noisy target  $\mathbf{x}$ . Next, we apply TSE to the noisy target  $\mathbf{x}$  using  $f_1(\cdot)$  and obtain the enhanced signal  $\hat{\mathbf{s}}_{i-1}$ . Then, we train another DNN  $f_2(\cdot)$  using NyTT with  $\hat{\mathbf{s}}_{i-1}$  as the noisy target signal. To prevent the degradation of the target signal component, we apply TSE to the original noisy target  $\mathbf{x}$ , not to the already enhanced signal  $\hat{\mathbf{s}}_{i-1}$ . Through this iterative process, we can improve the SNR of noisy targets and expect the performance improvement of NyTT. We investigate the effectiveness of IterNyTT in Section 4.3.




Figure 2: Overview of IterNyTT.

#### 3.3 Effects of Mismatches Between Noise Signals

In the NyTT framework, there are three types of noise signal:  $\mathbf{n}^{\text{obs}}$ ,  $\mathbf{n}^{\text{add}}$ , and noise included in the test data  $\mathbf{n}^{\text{test}}$ . In Section 4.4, we investigate the effects



of mismatches between types of noise signal on the performance of NyTT and IterNyTT. For instance, the performance of CTT is degraded when there is a mismatch between  $\mathbf{n}^{\text{add}}$  and  $\mathbf{n}^{\text{test}}$  [64, 22]. In the experiments, 1) we evaluate the performance of CTT, NyTT, and IterNyTT under mismatched conditions, and investigate the effects of each mismatch (Section 4.4.1). Additionally, considering the effects of mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ , we investigate the effects of 2) the SNRs of the noisy targets  $\mathbf{x}$  ( $\text{SNR}_{\mathbf{x}} = \log_{10} \|\mathbf{s}\|_2^2 / \|\mathbf{n}^{\text{obs}}\|_2^2$ ) and 3) the SNRs of the *more noisy* signals  $\mathbf{y}$  ( $\text{SNR}_{\mathbf{y}} = \log_{10} \|\mathbf{x}\|_2^2 / \|\mathbf{n}^{\text{add}}\|_2^2$ ) on performance.

### 3.4 Effectiveness of Utilizing Noisy Signals in a Situation Where Clean Target Signals are Available

Collecting a large number of clean target signals is challenging owing to high recording costs. In some cases, no clean target signals may be available, whereas in others, a small number of clean target signals are available, depending on the task. In Section 4.5, assuming situations where a small number of clean target signals are available, we investigate the effectiveness of utilizing a larger number of noisy signals.

### 3.5 Capabilities in the Dereverberation Task

The dereverberation task aims to restore an original signal from a reverberant signal. CTT can achieve dereverberation in the same manner as the denoising task by inputting reverberant signals and training a DNN to predict clean target signals. In Section 5, we investigate whether NyTT can also perform dereverberation as CTT does. Specifically, we evaluate the performance of CTT, NyTT, and IterNyTT, where CTT predicts a clean target signal  $\mathbf{s}$  from a reverberant signal  $\mathbf{s} * \mathbf{r}^{\text{add}}$  and NyTT predicts  $\mathbf{x} = \mathbf{s} * \mathbf{r}^{\text{obs}}$  from  $\mathbf{x} * \mathbf{r}^{\text{add}}$ , where  $\mathbf{r}^{\text{obs}}$  and  $\mathbf{r}^{\text{add}}$  are room impulse responses (RIRs), and  $*$  denotes the convolution operation.

### 3.6 Capabilities in the Declipping Task

In addition to the dereverberation task, we investigate the capabilities of NyTT in the declipping task. The declipping task aims to restore an original signal from a clipped signal where the clipping function  $f_{\text{clip}} : \mathbb{R}^T \rightarrow \mathbb{R}^T$  is defined as

$$f_{\text{clip}}(\mathbf{s}; c)[m] = \begin{cases} \mathbf{s}[m] & |\mathbf{s}[m]| < c \\ c \cdot \text{sgn}(\mathbf{s}[m]) & \text{otherwise} \end{cases}, \quad (10)$$

where  $m$  is the time index and  $c$  is a clipping threshold. In Section 6, we evaluate the performance of CTT, NyTT, and IterNyTT, where CTT predicts

a clean target signal  $\mathbf{s}$  from a clipped signal  $f_{\text{clip}}(\mathbf{s}; c^{\text{add}})$  NyTT predicts  $\mathbf{x} = f_{\text{clip}}(\mathbf{s}; c^{\text{obs}})$  from  $f_{\text{clip}}(\mathbf{x}; c^{\text{add}})$ , and  $c^{\text{add}} < c^{\text{obs}}$ .

## 4 Experimental Analysis in the Denoising Task

One promising application of unsupervised TSE methods is the extraction of environmental sounds, as clean human speech corpora have already been created through extensive community efforts [42, 16]. However, in our experiments, we used clean human speech as the target signal, as this allows us to control the quality and volume of the noisy signals by distorting the clean speech corpora. We believe that the insights gained from our experiments are equally applicable to other types of target signal.

### 4.1 Setups

In the experiments, we used several datasets as shown in Table 1, including LibriSpeech [42], CHiME3 [2], noise extracted from the training dataset of VoiceBank-DEMAND [57], and the training dataset of the DCASE 2016 Challenge Task2 dataset (DCASE) [19]. CHiME3 included background noise recorded in a bus, a cafe, a pedestrian area, and a street junction. VoiceBank-DEMAND included noise recorded in a kitchen, an office, a cafe, and a subway, along with artificially synthesized bubble and white noise. The DCASE 2016 Task2 dataset included sounds of coughing, door knocking, and telephone ringing. Thus, there are differences in the types of noise across these three datasets. The noise signals from CHiME3 were segmented every 10 s, and 7.83 h of data were split into CHiME-A and CHiME-B, and another 0.56 h of data were used as CHiME-C. 11,572 clips of VoiceBank-DEMAND were split into two subsets, DEMAND-A and DEMAND-B. The noisy target training dataset was generated by mixing 10,000 utterances of clean target signals from LibriSpeech and noise signals  $\mathbf{n}^{\text{obs}}$  at  $\text{SNR}_{\mathbf{x}}$  randomly selected from 0, 5, 10, and 15 dB. The test dataset was generated by mixing 1,000 utterances of clean target signals from LibriSpeech and CHiME-C at SNR randomly selected from 2.5, 7.5, 12.5, and 17.5 dB. During training, the input signal to the DNN was generated by mixing target signals and additional noise signals, and the  $\text{SNR}_{\mathbf{y}}$  ranges were -5 to 5 dB for NyTT, and 0, 5, 10, and 15 dB for IterNyTT after the second iteration and CTT. We evaluated the performance using the best validation epoch, where the validation was conducted with 50 pairs of input and target signals generated under the same  $\text{SNR}_{\mathbf{y}}$ ,  $\mathbf{n}^{\text{obs}}$ , and  $\mathbf{n}^{\text{add}}$  as in the training dataset.

The DNN architecture was CNN-BLSTM used in [37]. The input feature was the log-amplitude spectrogram, and the network estimated a complex-valued T-F mask. For the short-time Fourier transform (STFT) parameters,

Table 1: Datasets used in the denoising task

Signal type	Original Dataset	Train set		Test set
Clean signal $\mathbf{s}$ (Utterances)	LibriSpeech [42]	(10,000)		(1,000)
Noise $\mathbf{n}$ (Volume [h])	CHiME3 [2]	CHiME-A (3.92)	CHiME-B (3.92)	CHiME-C (0.56)
	VoiceBank-DEMAND [57]	DEMAND-A (4.70)	DEMAND-B (4.69)	
	DCASE 2016 Task2 [19]	DCASE (0.07)		

the frame shift, window size, and DFT size were set to 128, 512, and 512 samples, respectively, using the Hamming window with a sampling frequency of 16 kHz. We trained the DNN for 1,500 epochs with a mini-batch size of 50, using the Adam optimizer [26] with a fixed learning rate of 0.0001. For the loss function, we used MSE calculated in the time domain. As the metrics, we used the scale-invariant signal-to-distortion ratio (SI-SDR) [30], the perceptual evaluation of speech quality (PESQ) [20], and the short-time objective intelligibility (STOI) [52].

#### 4.2 Validity of Interpretation of $NyTT$

In this section, we trained the DNN using CHiME-A and CHiME-B as  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$ , respectively.

##### 4.2.1 Analysis of signals processed in $NyTT$

If  $NyTT$  is Noise2Noise, the output signal corresponding to the *more noisy* signal  $\mathbf{y}$  should be the estimate of the clean target signal  $\mathbf{s}$ . To investigate this, we analyzed the output signals when we input  $\mathbf{y}$  to the DNN. To analyze the output signals during training, we generated *more noisy* signals by using the training dataset. Additionally, to analyze the output signals for unseen *more noisy* signals, we generated *more noisy* signals by mixing noisy signals from the test dataset and CHiME-C at SNRs ranging from -5 to 5 dB.

Table 2 shows the speech quality of 1,000 utterances of the noisy targets  $\mathbf{x}$ , the *more noisy* signals  $\mathbf{y}$ , and the corresponding output signals  $f(\mathbf{y}; \theta)$ , for both the training and test datasets. Considering that the SI-SDR, PESQ, and STOI of clean speech are  $\infty$  dB, 4.64, and 1.00, respectively, the output signals are closer in quality to the noisy targets than the clean target signals. Figure 3 provides examples of spectrograms from the test dataset, further illustrating that the output signal is better interpreted as the estimate of the noisy target rather than the clean target signal.

Table 2: Quality of singles processed in NyTT.  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  were CHiME-A and CHiME-B, respectively.

Data	Metric	Noisy target	More noisy	Output
Train set	SI-SDR	<b>7.33</b>	-2.13	6.16
	PESQ	1.35	1.07	<b>1.37</b>
	STOI	<b>0.838</b>	0.666	0.776
Test set	SI-SDR	<b>9.67</b>	-1.58	7.76
	PESQ	<b>1.48</b>	1.08	1.47
	STOI	<b>0.874</b>	0.696	0.811

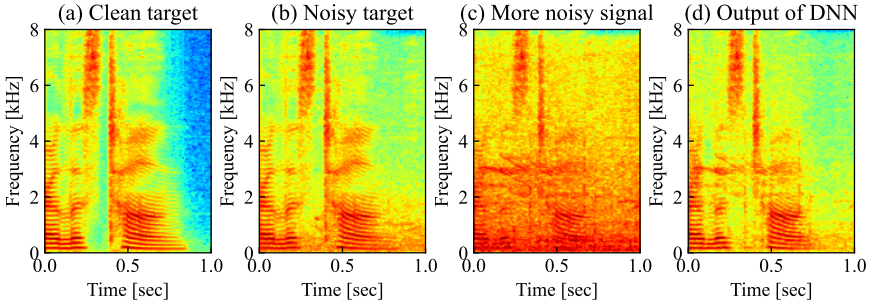


Figure 3: Spectrograms of the test dataset. (a) Clean target  $\mathbf{s}$ , (b) noisy target  $\mathbf{x}$ , (c) *more noisy* signal  $\mathbf{y}$ , and (d) output of a DNN  $f(\mathbf{y}; \theta)$ .  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  used for the training were CHiME-A and CHiME-B, respectively.

#### 4.2.2 Evaluation of NyTT with loss functions that do not satisfy the conditions of Noise2Noise

If NyTT is Noise2Noise,  $\mathbf{n}^{\text{obs}}$  must have a zero-mean distribution, and therefore, the MSE of the loss function should be calculated in the time domain. To analyze the significance of this assumption in NyTT, we compared the performance of NyTT with MSE in the time domain (**Time**) and that with MSE in the amplitude spectrogram domain (**Spec**). If NyTT strictly adheres to the Noise2Noise framework, **Spec** should not be able to perform TSE, as the zero-mean distribution for  $\mathbf{n}^{\text{obs}}$  cannot be satisfied in the amplitude spectrogram domain. In this experiment, we estimated real-valued T-F masks for the amplitude spectrograms and transformed the spectrograms to the time-domain signals using the phase of the unprocessed noisy signals.

Table 3 shows the evaluation results, demonstrating that **Spec** can improve speech quality, despite the lack of the zero-mean assumption for  $\mathbf{n}^{\text{obs}}$ . The results in Sections 4.2.1 and 4.2.2 indicate that NyTT achieves TSE by reducing the noise component corresponding to  $\mathbf{n}^{\text{add}}$  rather than the Noise2Noise framework. Therefore, the zero-mean distribution assumption for  $\mathbf{n}^{\text{obs}}$  and

Table 3: Comparison of loss functions.  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  were CHiME-A and CHiME-B, respectively.

	Unprocessed	Time	Spec
SI-SDR	9.67	<b>15.89</b>	14.94
PESQ	1.48	<b>2.33</b>	2.10
STOI	0.874	<b>0.928</b>	0.923

the use of the MSE loss function are unnecessary, making NyTT a more flexible training strategy (this conclusion is also supported by the experimental results in Sections 5 and 6).

### 4.3 Effectiveness of IterNyTT

To verify the effectiveness of IterNyTT, we evaluated its performance over five iterations. In this experiment,  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  were CHiME-A and CHiME-B, respectively.

Figure 4 illustrates the SI-SDR of the noisy targets, along with the SI-SDR, PESQ, and STOI of the processed results for the test dataset at each iteration of IterNyTT. The figure shows that IterNyTT improves the quality of the noisy targets, and thus, the performance on the test dataset approaches that of CTT as the number of iterations increases. Additionally, we observe that the SI-SDR of the noisy target is significantly improved at the second iteration and the performance on the test dataset is also improved at that time.

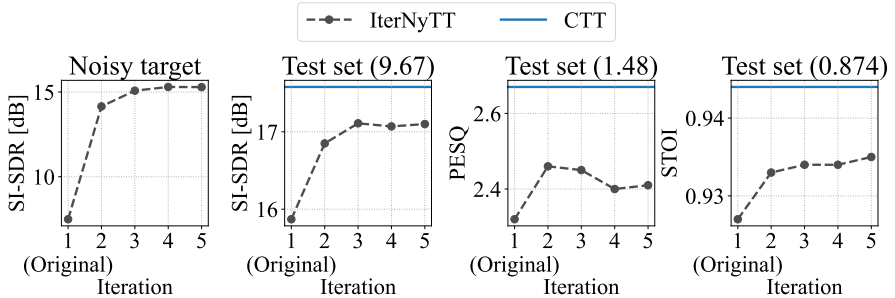


Figure 4: Changes in SI-SDR of the target signals and evaluation results for the test dataset through IterNyTT. The first iteration of IterNyTT is equivalent to the original NyTT. Values in parentheses indicate the evaluation results of unprocessed input signals.  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  were CHiME-A and CHiME-B, respectively.

#### 4.4 Effects of Noise Mismatches

To investigate the effects of mismatches between noise signals in NyTT (i.e.,  $\mathbf{n}^{\text{obs}}$ ,  $\mathbf{n}^{\text{add}}$ , and  $\mathbf{n}^{\text{test}}$ ), we simulated mismatched conditions using CHiME-A, DEMAND-A, and DCASE as  $\mathbf{n}^{\text{obs}}$ , and CHiME-B and DEMAND-B as  $\mathbf{n}^{\text{add}}$ .

To clarify the noise mismatches, we visualize the distribution of each noise dataset in Figure 5. The plots in Figure 5 were created by extracting features using a pre-trained audio event classification model, VGGish [18], and projecting them into a two-dimensional space using UMAP [39]. For visibility, we randomly selected 200 samples of 2 s noise signals from each dataset. The figure shows that each noise dataset forms a distinct cluster, indicating that they have different characteristics. We can also see that DCASE has particularly different characteristics from the other datasets.




Figure 5: Distribution of each noise dataset. Although UMAP features were calculated using all noise datasets, they were separately plotted for visibility. All figures have the same axes.

##### 4.4.1 Effects of mismatches on the performance

We investigated the effects of mismatches between noise signals on the performance of CTT, NyTT, and IterNyTT. In this experiment, we set the number of iterations for IterNyTT to three.

Table 4 shows the evaluation results of CTT, NyTT, and IterNyTT for each combination of  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$ . The table also includes the evaluation results of IterNyTT using different noise datasets for the first and second iterations (training for the TSE of the noisy targets) and the third iteration (training for the TSE of the test dataset). We analyze these results from three perspectives: mismatches between a)  $\mathbf{n}^{\text{add}}$  and  $\mathbf{n}^{\text{test}}$ , b)  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ , and c)  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$ , where  $\mathbf{n}^{\text{test}}$  was CHiME-C.

First, we focus on the impact of the mismatch between  $\mathbf{n}^{\text{add}}$  and  $\mathbf{n}^{\text{test}}$  on the performance of NyTT. For example, we analyze the performance of NyTT

Table 4: Evaluation results on the test dataset. The SI-SDR, PESQ, and STOI of the unprocessed noisy signals were 9.67 dB, 1.48, and 0.874, respectively.  $\mathbf{n}^{\text{test}}$  was CHiME-C. When IterNyTT used two noise datasets as  $\mathbf{n}^{\text{add}}$ , the first noise dataset in parentheses was used in the first and second iterations of IterNyTT, whereas the second noise dataset in parentheses was used in the third iteration of IterNyTT. In CTT, the choice of  $\mathbf{n}^{\text{obs}}$  is not involved.

$\mathbf{n}^{\text{obs}}$	$\mathbf{n}^{\text{add}}$	SISDR			PESQ			STOI		
		CTT	NyTT	IterNyTT	CTT	NyTT	IterNyTT	CTT	NyTT	IterNyTT
CHiME-A	CHiME-B	<b>17.58</b>	15.87	17.11	<b>2.67</b>	2.32	2.45	<b>0.944</b>	0.927	0.934
	DEMAND-B	<b>15.04</b>	10.27	9.80	<b>2.16</b>	1.57	1.51	<b>0.926</b>	0.882	0.877
	(CHiME-B, DEMAND-B)	-	-	10.67	-	-	1.59	-	-	0.881
	(DEMAND-B, CHiME-B)	-	-	14.64	-	-	1.96	-	-	0.915
DEMAND-A	CHiME-B	<b>17.58</b>	16.59	17.20	<b>2.67</b>	2.53	2.48	<b>0.944</b>	0.936	0.937
	DEMAND-B	<b>15.04</b>	13.56	14.11	<b>2.16</b>	1.98	2.06	<b>0.926</b>	0.905	0.917
	(CHiME-B, DEMAND-B)	-	-	13.85	-	-	1.94	-	-	0.903
	(DEMAND-B, CHiME-B)	-	-	17.33	-	-	2.56	-	-	0.940
DCASE	CHiME-B	<b>17.58</b>	16.75	17.06	<b>2.67</b>	2.47	2.45	<b>0.944</b>	0.937	0.937
	DEMAND-B	<b>15.04</b>	14.20	13.54	<b>2.16</b>	2.06	1.91	<b>0.926</b>	0.914	0.908

when  $\mathbf{n}^{\text{obs}}$  is CHiME-A. In this case, NyTT achieves SI-SDRs of 15.87 dB when  $\mathbf{n}^{\text{add}}$  is CHiME-B and 10.27 dB when  $\mathbf{n}^{\text{add}}$  is DEMAND-B. Similarly, even when  $\mathbf{n}^{\text{obs}}$  is DEMAND-A or DCASE, and even when the metric is PESQ or STOI, we can consistently see that NyTT performs better when there is no mismatch between  $\mathbf{n}^{\text{add}}$  and  $\mathbf{n}^{\text{test}}$ , as in CTT.

Second, we focus on the impact of the mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$  on the performance of NyTT. For example, we analyze the performance of NyTT when  $\mathbf{n}^{\text{add}}$  is DEMAND-B. In this case, NyTT achieves SI-SDRs of 10.27 dB when  $\mathbf{n}^{\text{obs}}$  is CHiME-A, 13.56 dB when  $\mathbf{n}^{\text{obs}}$  is DEMAND-A, and 14.20 dB when  $\mathbf{n}^{\text{obs}}$  is DCASE. Similarly, even when  $\mathbf{n}^{\text{add}}$  is CHiME-B, and even when the metric is PESQ or STOI, we can consistently see that NyTT performs better when there is a mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ . In particular, we can see that NyTT achieves its best performance when  $\mathbf{n}^{\text{obs}}$  is DCASE, which has distinctly different characteristics from CHiME-C of  $\mathbf{n}^{\text{test}}$ , as shown in Figure 5. Since the DNN is trained to include  $\mathbf{n}^{\text{obs}}$  in the output signals in NyTT, it is expected that noise will remain in the output signals when there is no mismatch between  $\mathbf{n}^{\text{test}}$  and  $\mathbf{n}^{\text{obs}}$ .

Third, we focus on the impact of the mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  on the performance of IterNyTT. For example, we analyze the performance of IterNyTT when  $\mathbf{n}^{\text{obs}}$  is CHiME-A. In this case, IterNyTT achieves SI-SDRs of 17.11 dB when  $\mathbf{n}^{\text{add}}$  is CHiME-B and 14.64 dB when  $\mathbf{n}^{\text{add}}$  is (DEMAND-B, CHiME-B). When  $\mathbf{n}^{\text{add}}$  is (DEMAND-B, CHiME-B), in the first and second iterations, the mismatch between  $\mathbf{n}^{\text{obs}}$  (=CHiME-A) and  $\mathbf{n}^{\text{add}}$  (=DEMAND-B) prevents IterNyTT from effectively removing  $\mathbf{n}^{\text{obs}}$  from the noisy targets. Thus, there is no performance improvement on the test dataset in the third iteration. Additionally, IterNyTT achieves a higher SI-SDR when  $\mathbf{n}^{\text{add}}$  is (CHiME-B, DEMAND-B) (10.67 dB) than when  $\mathbf{n}^{\text{add}}$  is DEMAND-B (9.80 dB). Moreover,

IterNyTT shows little performance improvement when  $\mathbf{n}^{\text{obs}}$  is DCASE. Similarly, even when  $\mathbf{n}^{\text{obs}}$  is DEMAND-A, and even when the metric is PESQ or STOI, we can consistently see that IterNyTT performs better when using  $\mathbf{n}^{\text{add}}$  matched with  $\mathbf{n}^{\text{obs}}$  in the first and second iterations. The impact of the mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$  on IterNyTT is consistent with that of the mismatch between  $\mathbf{n}^{\text{add}}$  and  $\mathbf{n}^{\text{test}}$  on NyTT.

Summing up the above results, we derive the desirable condition for the NyTT framework, as shown in Figure 6. a) NyTT achieves high performance when there is no mismatch between  $\mathbf{n}^{\text{add}}$  and  $\mathbf{n}^{\text{test}}$ , as in CTT, b) NyTT achieves high performance when there is a mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ , and c) IterNyTT improves the performance when there is no mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{add}}$ . Additionally, these results are consistent with the interpretation of NyTT in Section 4.2.




Figure 6: Desirable condition in NyTT framework.

#### 4.4.2 Difference in the impact of $\text{SNR}_x$ with and without the mismatch between $\mathbf{n}^{\text{obs}}$ and $\mathbf{n}^{\text{test}}$

NyTT experiences performance degradation when  $\text{SNR}_x$  is low [13]. However, the impact of  $\text{SNR}_x$  on performance is expected to vary depending on the mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ . To investigate this, we evaluated the performance of NyTT using CHiME-A, DEMAND-A, and DCASE as  $\mathbf{n}^{\text{obs}}$ , and CHiME-B as  $\mathbf{n}^{\text{add}}$ , and by adjusting  $\text{SNR}_x$  to -5, 0, 5, 10, 15, and 20 dB. Additionally, we evaluated the performance of CTT for the case where  $\text{SNR}_x$  is  $\infty$  dB. In this experiment,  $\text{SNR}_y$  was set to range from -5 to 5 dB for both CTT and NyTT.

Figure 7 shows the SI-SDR, PESQ, and STOI of the processed results for the test dataset at each  $\text{SNR}_x$ . Overall, the performance of NyTT degrades as  $\text{SNR}_x$  decreases. The impact depends on mismatches; significant degradation occurs when  $\mathbf{n}^{\text{obs}}$  is CHiME-A, which has no mismatch with  $\mathbf{n}^{\text{test}}$ . In contrast, when there is a mismatch, the performance remains relatively high even at a low  $\text{SNR}_x$ . From the results of this experiment, we confirmed that the impact of  $\text{SNR}_x$  significantly depends on the mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ .






Figure 7: Relationship between  $\text{SNR}_x$  and the evaluation results of NyTT. Values in parentheses indicate the evaluation results of unprocessed input signals.  $\mathbf{n}^{\text{add}}$  was CHiME-B and  $\text{SNR}_y$  ranged from -5 to 5 dB.

#### 4.4.3 Difference in the impact of $\text{SNR}_y$ with and without the mismatch between $\mathbf{n}^{\text{obs}}$ and $\mathbf{n}^{\text{test}}$

Considering that NyTT trains a DNN to estimate noisy targets by removing  $\mathbf{n}^{\text{add}}$  from *more noisy* signals,  $\text{SNR}_y$  also affects the performance. For example, when there is no mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ , and  $\text{SNR}_y$  is high, the effect of reducing noise is less significant than the adverse effect of the residual noise in the output signal. To investigate this, we evaluated the performance of NyTT using CHiME-A and DCASE as  $\mathbf{n}^{\text{obs}}$ , and CHiME-B as  $\mathbf{n}^{\text{add}}$ , and by varying the  $\text{SNR}_y$  range to  $[-10, -5)$ ,  $[-5, 0)$ ,  $[0, 5)$ ,  $[5, 10)$ , and  $[10, 15)$  dB. We also evaluated the performance of CTT. In this experiment, we set  $\text{SNR}_x$  to 5 dB for NyTT.

Figure 8 shows the SI-SDR, PESQ, and STOI of the processed results for the test dataset at each  $\text{SNR}_y$  range, where the performance of CTT varies depending on the mismatch of SNR between the training and test datasets. When  $\mathbf{n}^{\text{obs}}$  is DCASE, which has a mismatch with  $\mathbf{n}^{\text{test}}$ , the performance of NyTT has a similar tendency to that of CTT. On the other hand, when  $\mathbf{n}^{\text{obs}}$  is CHiME-A, which has no mismatch with  $\mathbf{n}^{\text{test}}$ , the performance of NyTT degrades when  $\text{SNR}_y$  exceeds 5 dB, and this trend differs from that of CTT. Thus, when there is no mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ , NyTT effectively acquires the TSE feature by setting  $\text{SNR}_y$  to a moderately low level, not too low.

#### 4.5 Effectiveness of Utilizing Noisy Signals in a Situation Where Clean Target Signals are Available

In this experiment, we used 100 utterances from LibriSpeech as Clean-100 and 900 utterances as Clean-900. We generated Noisy-900 by mixing Clean-900 with CHiME-A at an SNR of 5 dB. We investigated the effectiveness of using Noisy-900. Moreover, we investigated the effectiveness of using




Figure 8: Relationship between  $\text{SNR}_y$  and the evaluation results of NyTT. Values in parentheses indicate the evaluation results of unprocessed input signals.  $\mathbf{n}^{\text{add}}$  was CHiME-B and  $\text{SNR}_x$  was 5 dB.

**EnhNoisy-900**, which was generated by applying TSE to **Noisy-900** using the CTT model trained on **Clean-100**. In this experiment, we used CHiME-B as  $\mathbf{n}^{\text{add}}$ , and we used 50 clean utterances from LibriSpeech for the validation of both CTT and NyTT. Note that there was no mismatch between  $\mathbf{n}^{\text{obs}}$  and  $\mathbf{n}^{\text{test}}$ , and  $\text{SNR}_x$  was 5 dB, resulting in a challenging condition for NyTT. Since the volumes of the speech datasets were different, we carefully trained a DNN for enough epochs, ensuring that the best epoch remained unchanged for the last 300 epochs.

Table 5 shows the evaluation results, demonstrating that the combined use of **Clean-100** and **Noisy-900** achieves a higher SI-SDR than using either dataset separately. Furthermore, the performance of the combined use of **Clean-100** and **EnhNoisy-900** approaches that of the ideal situation where both **Clean-100** and **Clean-900** are available. We can also expect that the performance will improve with the use of noisy targets recorded under better conditions (i.e., higher  $\text{SNR}_x$  or  $\mathbf{n}^{\text{obs}}$  mismatched with  $\mathbf{n}^{\text{test}}$ ). These results indicate that leveraging a large number of noisy signals is beneficial, even when a small number of clean target signals are available.

Table 5: Evaluation results of the combined use of the clean and noisy signals.  $\mathbf{n}^{\text{obs}}$  was CHiME-A,  $\mathbf{n}^{\text{add}}$  was CHiME-B, and  $\text{SNR}_x$  was 5 dB. SI-SDR, PESQ, and STOI of the unprocessed noisy signals were 9.67 dB, 1.48, and 0.874, respectively.  $\mathbf{n}^{\text{test}}$  was CHiME-C. We assume that **Clean-900** is not available.

Training dataset	SI-SDR	PESQ	STOI	Epoch
Clean-100	13.82	2.03	0.910	3,816
Noisy-900	13.68	1.88	0.906	1,302
Clean-100, Noisy-900	14.35	2.00	0.914	2,350
Clean-100, EnhNoisy-900	16.45	2.33	0.935	11,626
Clean-100, Clean-900	17.13	2.56	0.942	11,593

## 5 Experimental Analysis in the Dereverberation Task

### 5.1 Setups

In the experiments, we used RIR simulated by utilizing Pyroomacoustics [49]. The room width and depth were randomly selected from 4 to 8 m, the height was randomly selected from 2 to 6 m, and the distance between the microphone and the sound source was set to 1 m. The reverberation time  $RT_{60}$  ranged from 0.20 to 1.10 s. We divided the range into 0.05 s segments and generated 170 RIRs for each interval. The 170 RIRs were split into 80, 80, and 10 samples, and which were used as RIR-A, RIR-B, and RIR-C, respectively. The total number of RIR-A, RIR-B, and RIR-C were 1,440, 1,440, and 180, respectively. Each of these three datasets covers the same  $RT_{60}$  range. RIR-A, RIR-B, and RIR-C were used as  $\mathbf{r}^{\text{obs}}$ ,  $\mathbf{r}^{\text{add}}$ , and  $\mathbf{r}^{\text{test}}$ , respectively. The clean target signals were 10,000 utterances from LibriSpeech and the reverberant target signals were generated by convolving the clean target signals with RIR-A. The test dataset of reverberant signals was generated by convolving 1,000 utterances from LibriSpeech with RIR-C. The sampling frequency was 16 kHz.

The DNN was Conv-TasNet [36] implemented in the Asteroid toolkit [43], and the loss function was SNR. We trained the DNN for 850 epochs with a mini-batch size of 12, using the Adam optimizer with a fixed learning rate of 0.0001. For the validation of both CTT and NyTT, we used 50 clean utterances of LibriSpeech and generated reverberant signals using the same  $\mathbf{r}^{\text{add}}$  as in the training. As the metrics, we used speech-to-reverberation modulation energy ratio (SRMR) [10] in addition to SI-SDR, PESQ, and STOI.

### 5.2 Effectiveness of NyTT in the Dereverberation Task

We conducted experimental evaluations of NyTT in the dereverberation task. Table 6 shows the evaluation results under different  $RT_{60}$  conditions of  $\mathbf{r}^{\text{obs}}$ . First, we observe that NyTT achieves higher scores than the unprocessed signals in most cases, demonstrating its capability in this task. This result also demonstrates that NyTT is not Noise2Noise, since the degradation is not even caused by additive noise. Additionally, we can see that the performance improves with shorter  $RT_{60}$  values and degrades with longer  $RT_{60}$  values. This trend is consistent with the results of the denoising task, where higher quality ( $\text{SNR}_x$ ) leads to better performance.

### 5.3 Effectiveness of IterNyTT in the Dereverberation Task

To verify the effectiveness of IterNyTT in the dereverberation task, we evaluated the performance over five iterations under different  $RT_{60}$  conditions for  $\mathbf{r}^{\text{obs}}$ . Figure 9 illustrates the SI-SDR of the reverberant targets, along

Table 6: Evaluation results of CTT and NyTT in the dereverberation task.

Method	RT <sub>60</sub> of $\mathbf{r}^{\text{obs}}$ [sec]	SI-SDR	PESQ	STOI	SRMR
Unprocessed	-	-5.32	1.59	0.834	4.84
CTT	0.0	<b>3.83</b>	<b>2.23</b>	<b>0.918</b>	<b>8.81</b>
NyTT	[0.20, 0.50)	1.69	1.95	0.902	7.59
NyTT	[0.50, 0.80)	0.51	1.74	0.882	6.01
NyTT	[0.80, 1.10)	-1.39	1.57	0.840	4.85




Figure 9: Changes in SI-SDR of the target signals and evaluation results on the test dataset through IterNyTT under different RT<sub>60</sub> conditions of  $\mathbf{r}^{\text{obs}}$ . The first iteration of IterNyTT is equivalent to the original NyTT. Values in parentheses indicate the evaluation results of unprocessed input signals.

with the SI-SDR, PESQ, STOI, and SRMR of the processed results for the test dataset at each iteration of IterNyTT. From this figure, we observe an overall trend of improved performance of IterNyTT. When RT<sub>60</sub> of  $\mathbf{r}^{\text{obs}}$  was [0.8, 1.1), IterNyTT in the first iteration does not perform well, and the performance is not improved in the subsequent iterations. We can also see that IterNyTT works stably when the RT<sub>60</sub> of  $\mathbf{r}^{\text{obs}}$  is [0.5, 0.8), whereas it becomes unstable when the RT<sub>60</sub> of  $\mathbf{r}^{\text{obs}}$  is [0.2, 0.5). Although there are cases where IterNyTT does not work well, especially when the target signals are of very low quality, we can conclude that IterNyTT is generally effective even in the dereverberation task.

## 6 Experimental Analysis in the Declipping Task

### 6.1 Setups

In the experiments, the clean target signals were 10,000 utterances of LibriSpeech, and we generated the clipped target signals  $\mathbf{x}$  by clipping them with an SNR <sub>$\mathbf{x}$</sub>  of 3, 7, or 15 dB. The clipped signals of the test dataset were generated by clipping 1,000 utterances of LibriSpeech with the SNR randomly selected from 1, 3, 7, and 15 dB. During the training, for both CTT and NyTT, the clipping threshold was determined from the SNR <sub>$\mathbf{y}$</sub>  randomly selected from 1 to 9 dB.

The DNN was a causal Demucs [65, 5], and the loss function was a weighted sum of the L1 waveform and multi-resolution STFT losses, as in [29, 65]. The weights for the L1 waveform and multi-resolution STFT losses were set to 10 and 0.1, respectively. We trained the DNN for 400 epochs with a mini-batch size of 12, using the Adam optimizer with a fixed learning rate of 0.0001. For the validation, we used 50 clean utterances of LibriSpeech and generated clipped signals with the SNR randomly selected from 1, 3, 7, and 15 dB. As the metrics, we used SI-SDR, PESQ, and STOI.

### 6.2 Effectiveness of NyTT in the Declipping Task

We conducted experimental evaluations of NyTT in the declipping task. Table 7 shows the evaluation results under different  $\text{SNR}_{\mathbf{x}}$  conditions. As in the denoising and dereverberation tasks, we can see that NyTT is effective in the declipping task, NyTT works without satisfying the Noise2Noise conditions, and the performance of NyTT improves with higher target signal quality.

Table 7: Evaluation results of CTT and NyTT in the declipping task.

Method	$\text{SNR}_{\mathbf{x}}$ [dB]	SI-SDR	PESQ	STOI
Unprocessed	-	6.41	1.89	0.866
CTT	$\infty$	<b>16.59</b>	<b>3.53</b>	<b>0.965</b>
NyTT	15	15.27	3.21	0.959
NyTT	7	12.44	2.65	0.941
NyTT	3	10.09	2.28	0.915

### 6.3 Effectiveness of IterNyTT in the Declipping Task

To verify the effectiveness of IterNyTT in the declipping task, we evaluated the performance over five iterations under different  $\text{SNR}_{\mathbf{x}}$  conditions. Figure 10 illustrates the SI-SDR of the clipped targets, along with the SI-SDR, PESQ, and STOI of the processed results for the test dataset at each iteration of IterNyTT. Here, we can again observe a consistent trend: IterNyTT improves performance, although its effectiveness is affected by the quality of the target signals. Specifically, when  $\text{SNR}_{\mathbf{x}}$  is 15 dB, IterNyTT achieves performance comparable to that of CTT.

## 7 Conclusion

In this study, we conducted comprehensive experimental analyses of NyTT to elucidate its detailed properties. Our experiments revealed the following key findings: 1) NyTT can be interpreted as a training method to estimate the noisy target  $\mathbf{x}$  by removing  $\mathbf{n}^{\text{add}}$ , rather than strictly adhering to

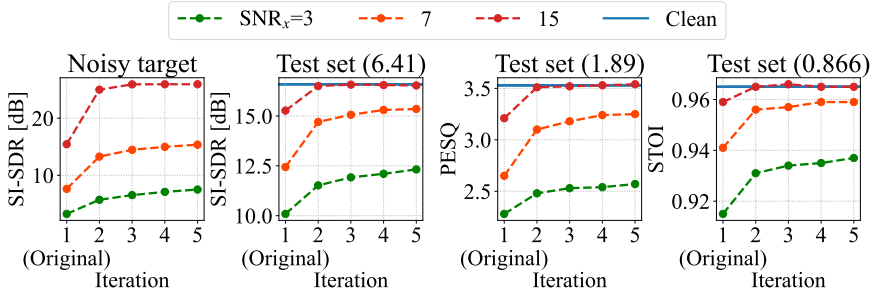


Figure 10: Changes in SI-SDR of the target signals and evaluation results on the test dataset through IterNyTT under different  $\text{SNR}_x$  conditions. The first iteration of IterNyTT is equivalent to the original NyTT. Values in parentheses indicate the evaluation results of unprocessed input signals.

the Noise2Noise framework. This indicates that the Noise2Noise conditions (i.e., the zero-mean distribution assumption for  $\mathbf{n}^{\text{obs}}$  and the use of the MSE loss function) are not necessary, demonstrating the flexibility of NyTT. 2) IterNyTT improved performance by enhancing the quality of noisy target signals, demonstrating its potential to achieve performance comparable to that of CTT. 3) By investigating the effects of noise mismatches, we derived desirable noise conditions. 4) Even when a small number of clean target signals were available, the combined use of noisy and clean target signals improved performance. 5) NyTT was also effective in the dereverberation and declipping tasks. Furthermore, both NyTT and IterNyTT exhibited similar behaviors across the denoising, dereverberation, and declipping tasks, implying their general applicability.

## Acknowledgements

This work was partly supported by JST CREST Grant Number JPMJCR19A3, JSPS KAKENHI Grant Numbers JP25KJ1439 and JP20H00102, and JST SPRING Grant Number JPMJSP2125.

## References

- [1] N. Alamdari, A. Azarang, and N. Kehtarnavaz, “Improving deep speech denoising by noisy2noisy signal mapping”, *Applied Acoustics*, 172, 2021, 107631.

- [2] 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, 504–11.
- [3] L.-W. Chen, Y.-F. Cheng, H.-S. Lee, Y. Tsao, and H.-M. Wang, “A Training and Inference Strategy Using Noisy and Enhanced Speech as Target for Speech Enhancement without Clean Speech”, in *Proc. Interspeech*, 2023, 2473–7, DOI: [10.21437/Interspeech.2023-1548](https://doi.org/10.21437/Interspeech.2023-1548).
- [4] A. Defossez, G. Synnaeve, and Y. Adi, “Real time speech enhancement in the waveform domain”, in *Proc. Interspeech*, 2020, 3291–5.
- [5] A. Défossez, G. Synnaeve, and Y. Adi, “Real Time Speech Enhancement in the Waveform Domain”, in *Proc. Interspeech*, 2020, 3291–5, DOI: [10.21437/Interspeech.2020-2409](https://doi.org/10.21437/Interspeech.2020-2409).
- [6] A. Défossez, N. Usunier, L. Bottou, and F. Bach, “Music source separation in the waveform domain”, *arXiv preprint arXiv:1911.13254*, 2019.
- [7] G. DeMuth, “Frequency domain beamforming techniques”, in *Proc. ICASSP*, Vol. 2, 1977, 713–5.
- [8] M. C. Du Plessis, G. Niu, and M. Sugiyama, “Analysis of learning from positive and unlabeled data”, *Advances in neural information processing systems*, 27, 2014.
- [9] H. Erdogan and T. Yoshioka, “Investigations on Data Augmentation and Loss Functions for Deep Learning Based Speech-Background Separation”, in *Proc. Interspeech*, 2018, 3499–503, DOI: [10.21437/Interspeech.2018-2441](https://doi.org/10.21437/Interspeech.2018-2441).
- [10] T. H. Falk, C. Zheng, and W.-Y. Chan, “A non-intrusive quality and intelligibility measure of reverberant and dereverberated speech”, *IEEE Transactions on Audio, Speech, and Language Processing*, 18(7), 2010, 1766–74.
- [11] I. Fedorov, M. Stamenovic, C. Jensen, L.-C. Yang, A. Mandell, Y. Gan, M. Mattina, and P. N. Whatmough, “TinyLSTMs: Efficient neural speech enhancement for hearing aids”, in *Proc. Interspeech*, 2020, 4054–8.
- [12] S.-W. Fu, C. Yu, K.-H. Hung, M. Ravanelli, and Y. Tsao, “MetricGAN-U: Unsupervised speech enhancement/dereverberation based only on noisy/reverberated speech”, in *Proc. ICASSP*, 2022, 7412–6.
- [13] T. Fujimura, Y. Koizumi, K. Yatabe, and R. Miyazaki, “Noisy-target Training: A Training Strategy for DNN-based Speech Enhancement without Clean Speech”, in *Proc. EUSIPCO*, 2021, 436–40.
- [14] T. Fujimura and R. Scheibler, “Multi-channel separation of dynamic speech and sound events”, in *Proc. Interspeech*, 2023, 3749–53.
- [15] T. Fujimura and T. Toda, “Analysis Of Noisy-Target Training For Dnn-Based Speech Enhancement”, in *Proc. ICASSP*, 2023, 1–5, DOI: [10.1109/ICASSP49357.2023.10096481](https://doi.org/10.1109/ICASSP49357.2023.10096481).

- [16] J. S. Garofolo, L. F. Lamel, W. M. Fisher, J. G. Fiscus, D. S. Pallett, N. L. Dahlgren, and V. Zue, “TIMIT Acoustic-Phonetic Continuous Speech Corpus LDC93S1”, Web Download. Philadelphia: Linguistic Data Consortium, 1993.
- [17] R. Hennequin, A. Khelif, F. Voituret, and M. Moussallam, “Spleeter: a fast and efficient music source separation tool with pre-trained models”, *Journal of Open Source Software*, 5(50), 2020, 2154.
- [18] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson, “CNN architectures for large-scale audio classification”, in *Proc. ICASSP*, 2017, 131–5.
- [19] “IEEE DCASE 2016 Challenge”, in <http://www.cs.tut.fi/sgn/arg/dcase2016/>, 2016.
- [20] International Telecommunication Union, “Wideband extension to recommendation P. 862 for the assessment of wideband telephone networks and speech codecs”, *ITU-T Recommendation P. 862.2*, 2007.
- [21] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks”, in *Proc. CVPR*, 2017, 1125–34.
- [22] N. Ito and M. Sugiyama, “Audio Signal Enhancement with Learning from Positive and Unlabeled Data”, in *Proc. ICASSP*, 2023, 1–5.
- [23] E. Karamatl and S. Krbz, “MixCycle: Unsupervised Speech Separation via Cyclic Mixture Permutation Invariant Training”, *IEEE Signal Processing Letters*, 29, 2022, 2637–41.
- [24] M. M. Kashyap, A. Tambwekar, K. Manohara, and S. Natarajan, “Speech Denoising Without Clean Training Data: A Noise2Noise Approach”, in *Proc. Interspeech*, 2021, 2716–20.
- [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, 175–9.
- [26] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization”, in *Proc. ICLR*, 2015, 1–15.
- [27] K. Kinoshita, T. Ochiai, M. Delcroix, and T. Nakatani, “Improving noise robust automatic speech recognition with single-channel time-domain enhancement network”, in *Proc. ICASSP*, 2020, 7009–13.
- [28] Y. Koizumi, S. Karita, S. Wisdom, H. Erdogan, J. R. Hershey, L. Jones, and M. Bacchiani, “DF-Conformer: Integrated architecture of ConvTasNet and Conformer using linear complexity self-attention for speech enhancement”, in *Proc. WASPAA*, 2021, 161–5.
- [29] Y. Kwon and J.-W. Choi, “Speech-Decliping Transformer with Complex Spectrogram and Learnerble Temporal Features”, *arXiv preprint arXiv:2409.12416*, 2024.



- [30] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR –half-baked or well done?”, in *Proc. ICASSP*, 2019, 626–30.
- [31] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, “Photo-realistic single image super-resolution using a generative adversarial network”, in *Proc. CVPR*, 2017, 4681–90.
- [32] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, “Noise2Noise: Learning image restoration without clean data”, in *Proc. ICML*, 2018, 2965–74.
- [33] L. Li and S. Seki, “Remixed2Remixed: Domain adaptation for speech enhancement by Noise2Noise learning with Remixing”, in *Proc. ICASSP*, 2024, 806–10.
- [34] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, “Speech enhancement based on deep denoising autoencoder”, in *Proc. Interspeech*, Vol. 2013, 2013, 436–40.
- [35] Y. Luo, Z. Chen, and T. Yoshioka, “Dual-path rnn: efficient long sequence modeling for time-domain single-channel speech separation”, in *Proc. ICASSP*, 2020, 46–50.
- [36] Y. Luo and N. Mesgarani, “Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation”, *IEEE/ACM TASLP*, 27(8), 2019, 1256–66.
- [37] M. Kawanaka, Y. Koizumi, R. Miyazaki, and K. Yatabe, “Stable training of DNN for speech enhancement based on perceptually-motivated black-box cost function”, in *Proc. ICASSP*, 2020, 7524–8.
- [38] M. Maciejewski, J. Shi, S. Watanabe, and S. Khudanpur, “Training noisy single-channel speech separation with noisy oracle sources: A large gap and a small step”, in *Proc. ICASSP*, 2021, 5774–8.
- [39] L. McInnes, J. Healy, N. Saul, and L. Grossberger, “UMAP: Uniform Manifold Approximation and Projection”, *The Journal of Open Source Software*, 3(29), 2018, 63 pages.
- [40] N. Wiener, *Extrapolation, interpolation and smoothing of stationary time series with engineering applications*, 1949.
- [41] A. Narayanan and D. Wang, “Ideal ratio mask estimation using deep neural networks for robust speech recognition”, in *Proc. ICASSP*, 2013, 7092–6.
- [42] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an asr corpus based on public domain audio books”, in *Proc. ICASSP*, 2015, 5206–10.
- [43] M. Pariente, S. Cornell, J. Cosentino, S. Sivasankaran, E. Tzinis, J. Heitkaemper, M. Olvera, F.-R. Stöter, M. Hu, J. M. Martín-Doñas, D. Ditter, A. Frank, A. Deleforge, and E. Vincent, “Asteroid: the PyTorch-based audio source separation toolkit for researchers”, in *Proc. Interspeech*, 2020, 2637–41.

- [44] S. Pascual, A. Bonafonte, and J. Serra, “SEGAN: Speech enhancement generative adversarial network”, in *Proc. Interspeech*, 2017, 3642–6.
- [45] S. Rouard, F. Massa, and A. Défossez, “Hybrid Transformers for Music Source Separation”, in *Proc. ICASSP*, 2023, 1–5.
- [46] S. F. Boll, “Suppression of acoustic noise in speech using spectral subtraction”, *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 27(2), 1979, 113–20.
- [47] K. Saijo and T. Ogawa, “Self-Remixing: Unsupervised Speech Separation VIA Separation and Remixing”, in *Proc. ICASSP*, 2023, 1–5.
- [48] K. Saito, S. Uhlich, G. Fabbro, and Y. Mitsufuji, “Training speech enhancement systems with noisy speech datasets”, *arXiv preprint arXiv:2105.12315*, 2021.
- [49] R. Scheibler, E. Bezzam, and I. Dokmani, “Pyroomacoustics: A python package for audio room simulation and array processing algorithms”, in *Proc. ICASSP*, 2018, 351–5.
- [50] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, “Attention is all you need in speech separation”, in *Proc. ICASSP*, 2021, 21–5.
- [51] C. Subakan, M. Ravanelli, S. Cornell, F. Lepoutre, and F. Grondin, “Resource-efficient separation transformer”, *arXiv preprint arXiv:2206.09507*, 2022.
- [52] 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*, 19(7), 2011, 2125–36.
- [53] V. A. Trinh and S. Braun, “Unsupervised speech enhancement with speech recognition embedding and disentanglement losses”, in *Proc. ICASSP*, 2022, 391–5.
- [54] N. Turpault, S. Wisdom, H. Erdogan, J. Hershey, R. Serizel, E. Fonseca, P. Seetharaman, and J. Salamon, “Improving sound event detection in domestic environments using sound separation”, *arXiv preprint arXiv:2007.03932*, 2020.
- [55] 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-STSP*, 2022.
- [56] E. Tzinis, Z. Wang, and P. Smaragdis, “Sudo rm-rf: Efficient networks for universal audio source separation”, in *Proc. MLSP*, 2020, 1–6.
- [57] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, “Investigating RNN-based speech enhancement methods for noise-robust Text-to-Speech.”, in *Proc. SSW*, 2016, 146–52.
- [58] H. L. Van Trees, *Optimum array processing: Part IV of detection, estimation, and modulation theory*, 2004.

- [59] B. D. Van Veen and K. M. Buckley, “Beamforming: a versatile approach to spatial filtering”, *IEEE ASSP Magazine*, 5(2), 1988, 4–24.
- [60] S. Wang, H. Guan, and Y. Long, “QMixCAT: Unsupervised Speech Enhancement Using Quality-guided Signal Mixing and Competitive Alternating Model Training”, in *Proc. Interspeech*, 2024, 642–6, DOI: [10.21437/Interspeech.2024-639](https://doi.org/10.21437/Interspeech.2024-639).
- [61] 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 TASLP*, 31, 2023, 3221–36.
- [62] D. S. Williamson, Y. Wang, and D. Wang, “Complex Ratio Masking for Monaural Speech Separation”, *IEEE/ACM TASLP*, 24(3), 2016, 483–92.
- [63] 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, 51–5.
- [64] S. Wisdom, E. Tzinis, H. Erdogan, R. Weiss, K. Wilson, and J. Hershey, “Unsupervised sound separation using mixture invariant training”, *Advances in NuerIPS*, 33, 2020, 3846–57.
- [65] J. Yi, J. Koo, and K. Lee, “DDD: A Perceptually Superior Low-Response-Time DNN-Based Declipper”, in *Proc. ICASSP*, 2024, 801–5.
- [66] T. Yoshioka, N. Ito, M. Delcroix, A. Ogawa, K. Kinoshita, M. Fujimoto, C. Yu, W. J. Fabian, M. Espi, T. Higuchi, S. Araki, and T. Nakatani, “The NTT CHiME-3 system: Advances in speech enhancement and recognition for mobile multi-microphone devices”, in *Proc. ASRU*, 2015, 436–43, DOI: [10.1109/ASRU.2015.7404828](https://doi.org/10.1109/ASRU.2015.7404828).
- [67] J. Zhang, C. Zorila, R. Doddipatla, and J. Barker, “Teacher-student MixIT for unsupervised and semi-supervised speech separation”, in *Proc. Interspeech*, 2021, 3495–9.
- [68] R. Zhang, P. Isola, and A. A. Efros, “Colorful image colorization”, in *Proc. ECCV*, 2016, 649–66.