


## Full Length Article

## SuperM2M: Supervised and mixture-to-mixture co-learning for speech enhancement and noise-robust ASR

Zhong-Qiu Wang 

Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, 518055, Guangdong, PR China

## ARTICLE INFO

## Keywords:

Neural speech enhancement  
Noise-robust ASR

## ABSTRACT

The current dominant approach for neural speech enhancement is based on supervised learning by using simulated training data. The trained models, however, often exhibit limited generalizability to real-recorded data. To address this, this paper investigates training enhancement models directly on real target-domain data. We propose to adapt mixture-to-mixture (M2M) training, originally designed for speaker separation, for speech enhancement, by modeling multi-source noise signals as a single, combined source. In addition, we propose a co-learning algorithm that improves M2M with the help of supervised algorithms. When paired close-talk and far-field mixtures are available for training, M2M realizes speech enhancement by training a deep neural network (DNN) to produce speech and noise estimates in a way such that they can be linearly filtered to reconstruct the close-talk and far-field mixtures. This way, the DNN can be trained directly on real mixtures, and can leverage close-talk and far-field mixtures as a weak supervision to enhance far-field mixtures. To improve M2M, we combine it with supervised approaches to co-train the DNN, where mini-batches of real close-talk and far-field mixture pairs and mini-batches of simulated mixture and clean speech pairs are alternately fed to the DNN, and the loss functions are respectively (a) the mixture reconstruction loss on the real close-talk and far-field mixtures and (b) the regular enhancement loss on the simulated clean speech and noise. We find that, this way, the DNN can learn from real and simulated data to achieve better generalization to real data. We name this algorithm SuperM2M (supervised and mixture-to-mixture co-learning). Evaluation results on the CHiME-4 dataset show its effectiveness and potential.

## 1. Introduction

Deep learning has dramatically advanced speech enhancement (Wang & Chen, 2018). The current dominant approach is based on supervised learning, where clean speech is synthetically mixed with noises in simulated reverberant conditions to create paired clean speech and noisy-reverberant mixtures for training neural speech enhancement models in a supervised, discriminative way to predict the clean speech from its paired mixture (Wang & Chen, 2018). Although showing strong performance in matched simulated test conditions (Chen, Wang, Yoho, Wang, & Healy, 2016; Chen et al., 2020; Chetupalli & Habets, 2023; Ephrat et al., 2018; Kavalerov et al., 2019; Luo, Chen, & Yoshioka, 2020; Luo & Mesgarani, 2019; Nachmani, Adi, & Wolf, 2020; Pons, Liu, Pascual, & Serra, 2024; Quan & Li, 2024; Saijo et al., 2023; Tan, Wang, & Wang, 2022; Tesch & Gerkmann, 2021; Wang et al., 2023; Wang & Wang, 2020; Wang, Wang, & Wang, 2020, 2021; Xu, Rao, Chng, & Li, 2020; Zeghidour & Grangier, 2021; Zhang, Jung, Watanabe, & Qian, 2024; Zhang et al., 2023; Zhang, Xu, et al., 2021; Zheng et al., 2023; Zmolikova et al., 2023), the trained models often exhibit limited generalizability to real-recorded data (Cornell et al., 2023; Cox et al.,

2023; Haeb-Umbach et al., 2019; Leglaive et al., 2023; Pandey & Wang, 2020; Tzinis, Adi, et al., 2022; Tzinis, Wisdom, Remez, & R. Hershey, 2022; Wang & Chen, 2018; Yang, Pandey, & Wang, 2024; Zhang et al., 2024, 2023; Zhang, Shi, Li, Watanabe, & Qian, 2021), largely due to mismatches between simulated training and real-recorded test conditions.

A possible way to improve the generalizability, we think, is to have the model see, and learn to model, real-recorded target-domain mixtures during training. This, however, cannot be applied in a straightforward way, since the clean speech at each sample of the real mixtures cannot be annotated or computed in an easy way. As a result, there lacks a good sample-level supervision for real mixtures, unlike simulated mixtures where a sample-level supervision is readily available.

During data collection, multiple far-field microphones are usually utilized to record target speakers. In UNSSOR (Wang & Watanabe, 2023), our recent algorithm proposed for unsupervised speaker separation, we find that the mixture signal recorded by each far-field

E-mail address: [wang.zhongqiu41@gmail.com](mailto:wang.zhongqiu41@gmail.com).<https://doi.org/10.1016/j.neunet.2025.107408>

Received 15 June 2024; Received in revised form 31 January 2025; Accepted 13 March 2025

Available online 26 March 2025

0893-6080/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

microphone can be leveraged as a *weak supervision* for training DNNs to separate speakers. The idea is that each far-field mixture can be utilized as a constraint to regularize speaker estimates. That is, the speaker estimates, produced by using the mixtures captured at a subset of microphones as input, should be capable of being leveraged to reconstruct the mixture captured by each microphone.

On the other hand, during data collection, besides using far-field microphones to record target speech, a close-talk microphone is often placed near the target speaker to collect its close-talk speech (e.g., in the CHiME (Barker, Watanabe, Vincent, & Trmal, 2018), AMI (Carletta et al., 2006), AliMeeting (Yu et al., 2022), and MISP (Wu et al., 2024) setup).<sup>1</sup> Although the close-talk microphone can also pick up non-target signals, the recorded close-talk mixture typically has a much higher signal-to-noise ratio (SNR) of the target speaker than any far-field mixtures. In our recent mixture to mixture (M2M) algorithm (Wang, 2024a), which builds upon UNSSOR, we find that, besides far-field mixtures, the close-talk mixture can also be leveraged as a *weak supervision* for training DNNs to separate mixed speakers.

To leverage the weak-supervision in far-field and close-talk mixtures for separation, two difficulties need to be solved. First, far-field mixtures contain multiple sources and are not clean, and due to the contamination of the other sources (Barker et al., 2018; Carletta et al., 2006; Watanabe et al., 2020), close-talk mixtures are often not clean enough. Second, each speaker's image in the close-talk mixture is not time-aligned with its image in each far-field mixture, and each speaker's images in different far-field mixtures are also not time-aligned with each other. As a result, close-talk and far-field mixtures cannot be naively used as the training targets for training speaker separation models. To overcome the two difficulties, we have recently proposed UNSSOR in a conference publication (Wang & Watanabe, 2023) and M2M in a letter submission (Wang, 2024a). The idea is that, at each training step, we can (a) feed a far-field mixture to a DNN to produce an estimate for each speaker; and (b) regularize the speaker estimates such that they can be linearly filtered via multi-frame linear filtering to reconstruct the close-talk and far-field mixtures. This way, the first difficulty is addressed by having the filtered speaker estimates to respectively approximate (i.e., explain) the speaker images in each mixture, and the second difficulty is addressed by multi-frame linear filtering.

Although UNSSOR and M2M are capable of being trained directly on real-recorded mixtures (i.e., not requiring the availability of clean speech), they have been only trained and evaluated on simulated mixtures (Wang, 2024a; Wang & Watanabe, 2023). It is yet unknown (a) whether they are effective on real data; and (b) whether they can lead to better generalization to real data, compared with the current dominant purely-supervised approaches, which train models only on simulated data. We emphasize that these two concerns are very reasonable, since, on real data, the physical models hypothesized in UNSSOR and M2M are expected to be much less satisfied. For example, there could be microphone synchronization errors, microphone failures, different frequency responses in different microphones, signal clipping, slight speaker and array movement, non-linear filter relationships among speaker images at different microphones, distortions to target speech caused by real microphones, etc. These issues can potentially pose difficulties for UNSSOR and M2M. In addition, UNSSOR and M2M were designed for separating mixed speakers. It is unclear whether they would be effective for single-speaker speech enhancement and robust automatic speech recognition (ASR), where suppressing non-target signals (such as noises) is a major concern. Furthermore, UNSSOR and M2M assume stationary, weak Gaussian noises in their physical models.

It is unclear whether they can deal with strong, non-stationary noises, which could contain an unknown number of diffuse and directional sources.

In this context, we propose to extend UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a) for speech enhancement and robust ASR, where we train and evaluate the proposed algorithms not only on simulated data but also on real data.<sup>2</sup> Our major goal is to show whether the resulting algorithms can yield better generalization to real data than purely-supervised models trained on simulated data, a demonstration that is missing in UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a). Without this demonstration, the evidence supporting whether this un- and weakly-supervised line of research is worth investigating would be lacking, especially considering the dominance and simplicity of purely-supervised approaches based on simulated training data. We summarize the key contributions of this paper as follows:

- We are the first proposing to leverage close-talk and far-field mixtures as weak supervision for speech enhancement, a task different from speaker separation.
- Considering noise sources as a single, combined source, we formulate the training of speech enhancement models on real data as solving a blind deconvolution problem, following the formulations in UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a) designed for speaker separation.
- To account for the case when close-talk mixture is not time-synchronized with its paired far-field mixtures, we propose an on-the-fly filter tap estimation algorithm that can deal with the time-synchronization issue.
- We propose SuperM2M, a co-learning strategy which trains the same DNN model by alternating between M2M training on real data and supervised learning on simulated data. This way, M2M can benefit from massive simulated training data, especially when the real training data is scarce. In addition, we find that this strategy can help mitigate the weaknesses of UNSSOR and M2M on source permutation, frequency permutation, and source ambiguity (all of which will be detailed in Section 6.1).

We validate SuperM2M on the CHiME-4 dataset (Vincent, Watanabe, Nugraha, Barker, & Marxer, 2017), which is consisted of simulated and challenging real-recorded mixtures and is currently the major benchmark for evaluating robust ASR and speech enhancement algorithms. In our experiments, state-of-the-art ASR and enhancement performance is obtained. The evaluation results suggest that:

- SuperM2M can effectively learn from real mixtures and leverage the weak supervision afforded by real close-talk and far-field mixtures.
- The co-learning strategy can significantly improve the generalizability of purely-supervised models trained on simulated data to real data.

The evaluation results provide an evidence supporting the strong potential of our un- and weakly-supervised line of research for speech enhancement. A sound demo is provided in the link below.<sup>3</sup>

## 2. Related work

SuperM2M is related to other work in five major aspects.

<sup>1</sup> Close-talk speech is almost always recorded together with far-field speech in speech separation and recognition datasets, as it is much easier for humans to annotate word transcriptions and speaker activities based on close-talk recordings (where the target speech is very strong) than far-field recordings.

<sup>2</sup> This long-form paper is a journal expansion of our preliminary work in UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a).

<sup>3</sup> [https://zqwang7.github.io/demos/SuperM2M\\_demo/index.html](https://zqwang7.github.io/demos/SuperM2M_demo/index.html)

## 2.1. Frontend enhancement and robust ASR

Leveraging neural speech enhancement as a frontend processing to improve the robustness of backend ASR systems to noise, reverberation and competing speech has been a long-lasting research topic (Haeb-Umbach et al., 2020, 2019). Although dramatic progress has been made in neural speech enhancement (Wang & Chen, 2018; Zheng et al., 2023), directly feeding the immediate estimate produced by DNN-based enhancement models for ASR has had limited success, largely for two reasons: (a) enhancement DNNs, which can suppress non-target signals aggressively, often incur speech distortion detrimental to ASR; and (b) enhancement DNNs are often trained on simulated data, which inevitably mismatches real data, and this mismatch further aggravates the speech distortion problem. Through years of efforts, robust ASR approaches have gradually converged to (a) leveraging DNN estimates to derive linear beamforming results for ASR (Boeddecker et al., 2018; Heymann, Drude, Chinaev, & Haeb-Umbach, 2015; Zhang, Wang, & Wang, 2017); and (b) jointly training ASR models with enhancement models (Chang, Zhang, Qian, Le Roux, & Watanabe, 2019; Heymann, Drude, Boeddecker, Hanebrink, & Haeb-Umbach, 2017; Narayanan & Wang, 2015; Wang & Wang, 2016). These two approaches aim at improving robust ASR performance. Their enhancement modules usually do not produce sufficiently accurate estimation of target speech. For example, linear beamforming is known to introduce little speech distortion but it has limited capabilities at suppressing non-target signals (especially when the number of microphones is limited and when the non-target signals are diffuse) (Gannot, Vincent, Markovich-Golan, & Ozerov, 2017). Another example is that jointly training enhancement models with ASR models often degrades the performance of the enhancement models on realistic mixtures (Masuyama, Chang, Zhang, et al., 2023).

Differently, in this paper we aim at building neural speech enhancement models whose immediate estimate *itself* can have low distortion to target speech and high reduction to non-target signals, especially on real test data. We find that, on the challenging real test data of CHiME-4 (Vincent et al., 2017), the immediate output of SuperM2M bears low distortion to target speech and high reduction to non-target signals, and feeding it directly to strong ASR models for recognition yields strong performance.

## 2.2. Generalizability of supervised models to real data

Improving the generalizability of neural speech enhancement models to real data has received decade-long efforts. The current dominant approach (Chen et al., 2016; Wang & Chen, 2018; Zhang et al., 2024, 2023) is to train supervised models on large-scale synthetic data, which is simulated in a way to cover as many variations (that could happen in real test data) as possible. However, the success has been limited, largely due to the current simulation techniques being not good enough at generating simulated mixtures as realistic as real mixtures. This can be observed from recent speech enhancement and ASR challenges. In the Clarity enhancement challenge (Cox et al., 2023), all the teams scored well on simulated data failed on real data. In CHiME-3/4 (Vincent et al., 2017), in the multi-channel cases, all the top teams use conventional beamformers (although with signal statistics estimated based on DNN estimates) as the only frontend, and in the single-channel cases, frontend enhancement often degrades ASR performance compared to not using any enhancement (assuming no joint frontend-backend training) (Chang, Maekaku, Fujita, & Watanabe, 2022). In CHiME-5,6,7 (Watanabe et al., 2020) and M2MeT (Yu et al., 2022), almost all the teams adopt guided source separation (Boeddecker et al., 2018), a signal processing algorithm, as the only frontend.

Since the current simulation techniques are not satisfactory enough, a possible way to improve the generalizability to real data, we think, is to train enhancement models directly on real data.

## 2.3. Unsupervised speech separation

To model real data, unsupervised neural speech separation algorithms (such as MixIT Wisdom et al., 2020, ReMixIT Tzinis, Adi, et al., 2022, NyTT Fujimura, Koizumi, Yatabe, & Miyazaki, 2021, Neural FCA Bando et al., 2021, RAS Aralikatti, Boeddecker, Wichern, Subramanian, & Le Roux, 2023, UNSSOR Wang & Watanabe, 2023 and USDnet Wang, 2024b), which can train separation models directly on mixtures or synthetic mixtures of mixtures, have been proposed. Due to their unsupervised nature, their performance could be limited due to not leveraging any supervision. Meanwhile, many algorithms in this stream are only evaluated on simulated data and their effectiveness on real data and for robust ASR is unclear. In contrast, we will show that SuperM2M works well on the challenging real data of CHiME-4.

## 2.4. Semi-supervised speech separation

A promising direction, suggested by Han, Wilson, Wisdom, and R. Hershey (2024), Sivaraman, Wisdom, Erdogan, and R. Hershey (2022) (and subsequent studies Hao, Xu, & Xie, 2023; Zhang, Zorilă, Doddipatla, & Barker, 2022), is to combine supervised learning on simulated data and unsupervised learning on real data for model training, forming a semi-supervised approach. The rationale is that supervised learning on massive simulated data offers an easy and feasible way for the model to learn to model speech patterns, and unsupervised learning on real data can help the model learn from real data.

SuperM2M follows this direction, but differs from Han et al. (2024), Sivaraman et al. (2022) in two major aspects. First, SuperM2M leverages M2M (Wang, 2024a), which builds upon UNSSOR (Wang & Watanabe, 2023), to model real data, while (Han et al., 2024; Sivaraman et al., 2022) uses MixIT (Wisdom et al., 2020). As is suggested in Wang (2024b), Wang and Watanabe (2023), UNSSOR based methods (a) avoid tricky (and often unrealistic) synthesis of mixtures of mixtures, which, on the other hand, increases the number of sources to separate; (b) are more flexible at multi-channel separation; and (c) can be readily configured to perform dereverberation besides separation (Wang, 2024b), while MixIT cannot. On the other hand, when close-talk mixtures are available, M2M can be readily configured weakly-supervised to leverage the weak supervision afforded by close-talk mixtures.

## 2.5. Weakly-supervised speech separation

SuperM2M, building upon UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a), can be configured to leverage close-talk mixtures as a weak supervision to enhance far-field mixtures. In the literature, there are earlier studies on weakly-supervised speech enhancement and source separation. In Stoller, Ewert, and Dixon (2018), Zhang, Yan, and Zhou (2018), discriminators, essentially source prior models trained in an adversarial way, are used to help separation models produce separation results with distributions close to clean sources. In Chang et al. (2019), separation models are jointly trained with ASR models to leverage the weak supervision of word transcriptions. In Pishdadian, Wichern, and Le Roux (2020), a pre-trained sound classifier is employed to check whether separated signals can be classified into target sound classes, thereby promoting separation. These approaches require clean sources, human annotations (e.g., word transcriptions), and source prior models (e.g., discriminators, ASR models, and sound classifiers). In comparison, M2M needs close-talk and far-field mixture pairs, which can be obtained during data collection by using close-talk in addition to far-field microphones, and it does not require source prior models. In addition, close-talk mixtures exploited in M2M can provide a *sample-level* supervision, offering much more fine-grained supervision than source prior models, word transcriptions, and segment-level sound class labels.

### 3. Problem formulation

We start with describing the hypothesized physical models, then propose to formulate speech enhancement as a blind deconvolution problem, and at last overview the proposed algorithm for speech enhancement.

#### 3.1. Physical model

In noisy-reverberant conditions with a compact far-field P-microphone array and a single target speaker wearing a close-talk microphone, the physical model for each real-recorded mixture can be formulated, in the short-time Fourier transform (STFT) domain, as follows.

At a designated reference far-field microphone  $q \in \{1, \dots, P\}$ , the real-recorded mixture is formulated as

$$Y_q(t, f) = X_q(t, f) + V_q(t, f), \quad (1)$$

where  $t$  indexes  $T$  frames,  $f$  indexes  $F$  frequency bins, and  $Y_q(t, f)$ ,  $X_q(t, f)$ , and  $V_q(t, f)$  respectively denote the STFT coefficients of the mixture, speaker image of the target speaker, and non-speech signals at time  $t$ , frequency  $f$  and microphone  $q$ . In the rest of this paper, when dropping indices  $t$  and  $f$ , we refer to the corresponding spectrograms. We emphasize that  $V_q$  could contain multiple strong, non-stationary directional as well as diffuse noises. In this paper, we model them using a single, combined source.

Let us denote the index of the close-talk microphone as 0, which is different from the set of indices  $\{1, \dots, P\}$  used for far-field microphones. This way, we can index any of the non-reference microphones<sup>4</sup> using  $p \in \{0, 1, \dots, P\}$ , with  $p \neq q$ . We then formulate the mixture captured at microphone  $p$  as

$$\begin{aligned} Y_p(t, f) &= X_p(t, f) + V_p(t, f) \\ &= \mathbf{h}_p(f)^H \bar{\mathbf{X}}_q(t, f; p) + V_p(t, f) + \varepsilon_p(t, f) \\ &= \mathbf{h}_p(f)^H \bar{\mathbf{X}}_q(t, f; p) + \mathbf{r}_p(f)^H \bar{\mathbf{V}}_q(t, f; p) + \varepsilon'_p(t, f). \end{aligned} \quad (2)$$

In row 2, following narrowband approximation (Gannot et al., 2017; Talmon, Cohen, & Gannot, 2009), we approximate the speaker image captured at microphone  $p$  (i.e.,  $X_p(\cdot, f)$ ) as a linear convolution between the speaker image captured at microphone  $q$  (i.e.,  $X_q(\cdot, f)$ ) and a linear filter  $\mathbf{h}_p(f)$ . That is,  $X_p(t, f) \approx \mathbf{h}_p(f)^H \bar{\mathbf{X}}_q(t, f; p)$ , where  $\bar{\mathbf{X}}_q(t, f; p) = [X_q(t - I_p + 1, f), \dots, X_q(t + J_p, f)] \in \mathbb{C}^{I_p + J_p}$  and  $\mathbf{h}_p(f) \in \mathbb{C}^{I_p + J_p}$ , and we use  $\varepsilon_p$  to denote the modeling error. Notice that the past and future filter taps,  $I_p$  and  $J_p$ , are dependent on the microphone index (i.e.,  $p$ ) of the speaker image  $X_p$  we want to approximate.  $\mathbf{h}_p(f)$  can be interpreted as the relative transfer function (RTF) relating the speaker image  $X_q$  (captured by the reference far-field microphone  $q$ ) to the speaker image captured at another microphone  $p$  (i.e.,  $X_p$ ). In row 3, we use the same trick to approximate non-speech signals  $V_p(t, f) \approx \mathbf{r}_p(f)^H \bar{\mathbf{V}}_q(t, f; p)$ , where  $\bar{\mathbf{V}}_q(t, f; p) = [V_q(t - I_p + 1, f), \dots, V_q(t + J_p, f)] \in \mathbb{C}^{I_p + J_p}$  and  $\mathbf{r}_p(f) \in \mathbb{C}^{I_p + J_p}$ , and  $\varepsilon'_p$  absorbs the incurred modeling error and  $\varepsilon_p$ .

To facilitate understanding, we provide several comments:

- For simplicity, we use  $I_p - 1$  past and  $J_p$  future taps for both  $\bar{\mathbf{X}}_q(t, f; p)$  and  $\bar{\mathbf{V}}_q(t, f; p)$ , although it might be better to use different taps for different sources.
- For the filter taps  $I_p$  and  $J_p$ , we have a subscript  $p$  to indicate that we can use different number of filter taps for each non-reference microphone  $p$ . It may be a good idea to use different filter taps for the non-reference far-field microphones and the close-talk microphone, especially if they are placed on different devices or not synchronized with each other.

- $\mathbf{r}_p(f)^H \bar{\mathbf{V}}_q(t, f; p)$  could be a crude approximation of  $V_p(t, f)$ , as there could be multiple directional and diffuse noise sources in  $V$ , rather than a single directional source like in  $X_q$ .
- In the speaker image captured by the close-talk microphone (i.e.,  $X_0$ ), the direct-path signal of the target speaker is often much stronger than its reverberation. Therefore,  $X_0$  can be largely viewed as the dry source signal. In this case,  $\mathbf{h}_0(f)$  can be interpreted as a deconvolutional filter that can reverse the speaker image  $X_q$  back to the speech source signal.

#### 3.2. Formulating speech enhancement as blind deconvolution

As is suggested by our preliminary studies, UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a), both close-talk and far-field mixtures contain weak supervision for speaker separation. They (Wang, 2024a; Wang & Watanabe, 2023) exploit the weak supervision by formulating speaker separation as solving a blind deconvolution problem.

Similar to speaker separation, speech enhancement can be viewed as two-source separation, with one source being speech and the other being all non-speech sources combined. With this understanding, following UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a), we can also formulate speech enhancement as solving a blind deconvolution problem, and adapt UNSSOR and M2M for speech enhancement. In detail, the problem, formulated below, finds sources,  $X_q(\cdot, \cdot)$  and  $V_q(\cdot, \cdot)$ , and filters,  $\mathbf{h}_p(\cdot)$  and  $\mathbf{r}_p(\cdot)$ , that are most consistent with the physical models in (1) and (2):

$$\begin{aligned} \underset{\substack{X_q(\cdot, \cdot), V_q(\cdot, \cdot), \\ \mathbf{h}_p(\cdot), \mathbf{r}_p(\cdot)}}{\operatorname{argmin}} \quad & \left( \sum_{t, f} |Y_q(t, f) - X_q(t, f) - V_q(t, f)|^2 + \right. \\ & \left. \sum_{p=0, p \neq q} \sum_{t, f} |Y_p(t, f) - \mathbf{h}_p(f)^H \bar{\mathbf{X}}_q(t, f; p) - \mathbf{r}_p(f)^H \bar{\mathbf{V}}_q(t, f; p)|^2 \right), \end{aligned} \quad (3)$$

where  $|\cdot|$  computes magnitude. As is suggested in UNSSOR (Wang & Watanabe, 2023), this problem is non-convex and difficult to be solved, since the speech source, noise source, and linear filters are all unknown and need to be estimated. It is known not solvable if no prior knowledge is assumed about the sources or filters (Levin, Weiss, Durand, & T. Freeman, 2011). In our preliminary work, UNSSOR (Wang & Watanabe, 2023), which models source priors via unsupervised deep learning, is proposed to tackle this category of blind deconvolution problems. It is shown effective at separating reverberant multi-speaker mixtures to reverberant speaker images in simulated conditions.

#### 3.3. Overview of proposed algorithms

Building upon the preliminary successes of UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a) in speaker separation on simulated data, this paper adapts UNSSOR and M2M training for neural speech enhancement and performs training and evaluation on real-recorded data.

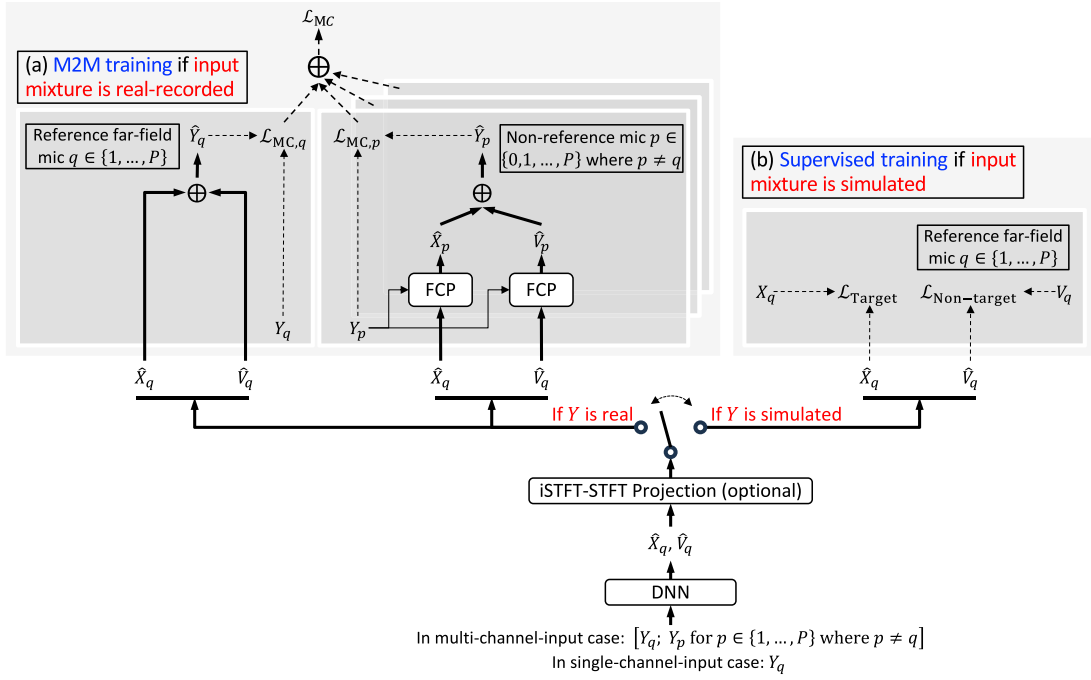
We first give UNSSOR and M2M a review in Section 4. We then propose to improve them by addressing potential synchronization issues between close-talk and far-field microphones in Section 5. Next, in Section 6 we point out several weaknesses of UNSSOR and M2M, and propose to address them by combining UNSSOR and M2M training on real data with supervised learning on simulated data.

### 4. M2M review, and adapting it for speech enhancement

Since we formulate speech enhancement as two-source separation, we can readily adapt UNSSOR and M2M (originally designed for speaker separation) for speech enhancement. This section reviews UNSSOR and M2M training, but in the context of adapting them for speech enhancement.

Fig. 1(a) illustrates M2M training for speech enhancement. The DNN takes in far-field mixtures as input and produces an estimate  $\hat{X}_q$  for the

<sup>4</sup> By non-reference microphones, we mean the close-talk microphone and  $P - 1$  non-reference far-field microphones.



**Fig. 1.** Illustration of SuperM2M, which consists of (a) M2M training on real close-talk and far-field mixture pairs (described in Section 4); and (b) supervised training on simulated far-field mixtures (described in Section 6). M2M trains the DNN to separate far-field mixtures to two sources that can be linearly filtered to reconstruct far-field and close-talk mixtures. In SuperM2M, we alternately feed in mini-batches of real mixtures and mini-batches of simulated mixtures, and train the same DNN by alternating between M2M training (if the input mixtures are real) and supervised learning (if the input mixtures are simulated).

target speaker and an estimate  $\hat{V}_q$  for non-target signals. Each estimate is then linearly filtered via forward convolutive prediction (Wang, Wichern, & Le Roux, 2021a) to optimize a so-called mixture-constraint loss (Wang & Watanabe, 2023), which encourages the filtered estimates to add up to the close-talk mixture and each far-field mixture, thereby exploiting the weak supervision afforded by close-talk and far-field mixtures for enhancement. This section describes the DNN setup, loss functions, and FCP filtering.

#### 4.1. DNN configurations

The DNN is trained to perform complex spectral mapping (Tan et al., 2022; Wang et al., 2020; Wang, Wang, & Wang, 2021), where the real and imaginary (RI) components of far-field mixtures are stacked as input features for the DNN to predict the RI components of  $\hat{X}_q$  and  $\hat{V}_q$ . The DNN setup is described later in Section 7.5, and the loss function next. We can optionally apply iSTFT-STFT projection to  $\hat{X}_q$  and  $\hat{V}_q$  before loss computation (see later Section 7.7.2 for details).

#### 4.2. Mixture-constraint loss

Following UNSSOR (Wang & Watanabe, 2023), M2M (Wang, 2024a) and the objective in (3), we design the following mixture-constraint (MC) loss, which regularizes the DNN estimates  $\hat{X}_q$  and  $\hat{V}_q$  to have them respectively approximate  $X_q$  and  $V_q$ , by checking whether they can be utilized to reconstruct the recorded mixtures:

$$\mathcal{L}_{MC} = \mathcal{L}_{MC,q} + \sum_{p=0, p \neq q}^P \alpha_p \times \mathcal{L}_{MC,p}, \quad (4)$$

where the two loss terms respectively follow the ones in (3) and will be detailed next, and  $\alpha_p$ , a weighting term, is set to 1.0 for the close-talk microphone and to  $1/(P-1)$  for non-reference far-field microphones.

$\mathcal{L}_{MC,q}$ , following the first term in (3), is the MC loss at the reference far-field microphone  $q$ :

$$\mathcal{L}_{MC,q} = \sum_{t,f} \mathcal{F}(Y_q(t, f), \hat{X}_q(t, f) + \hat{V}_q(t, f))$$

$$= \sum_{t,f} \mathcal{F}(Y_q(t, f), \hat{X}_q(t, f) + \hat{V}_q(t, f)), \quad (5)$$

where, in row 2, the DNN estimates  $\hat{X}_q$  and  $\hat{V}_q$  are utilized to reconstruct the mixture  $Y_q$  via  $\hat{Y}_q = \hat{X}_q + \hat{V}_q$ , and  $\mathcal{F}(\cdot, \cdot)$ , to be described in (7), is a distance function.

Similarly,  $\mathcal{L}_{MC,p}$ , following the second term in (3), is the MC loss at each non-reference microphone  $p$ :

$$\begin{aligned} \mathcal{L}_{MC,p} &= \sum_{t,f} \mathcal{F}(Y_p(t, f), \hat{X}_p(t, f) + \hat{V}_p(t, f)) \\ &= \sum_{t,f} \mathcal{F}(Y_p(t, f), \hat{X}_p(t, f) + \hat{V}_p(t, f)) \\ &= \sum_{t,f} \mathcal{F}(Y_p(t, f), \hat{\mathbf{h}}_p(f)^H \bar{\mathbf{X}}_q(t, f; p) + \hat{\mathbf{r}}_p(f)^H \bar{\mathbf{V}}_q(t, f; p)), \end{aligned} \quad (6)$$

where  $\hat{X}_p(t, f) = \hat{\mathbf{h}}_p(f)^H \bar{\mathbf{X}}_q(t, f; p)$ , with  $\bar{\mathbf{X}}_q(t, f; p) = [\hat{X}_q(t - I_p + 1, f), \dots, \hat{X}_q(t + J_p, f)]^T \in \mathbb{C}^{I_p + J_p}$  and  $\hat{\mathbf{h}}_p(f) \in \mathbb{C}^{I_p + J_p}$ , and  $\hat{V}_p(t, f) = \hat{\mathbf{r}}_p(f)^H \bar{\mathbf{V}}_q(t, f; p)$ , with  $\bar{\mathbf{V}}_q(t, f; p) = [\hat{V}_q(t - I_p + 1, f), \dots, \hat{V}_q(t + J_p, f)]^T \in \mathbb{C}^{I_p + J_p}$  and  $\hat{\mathbf{r}}_p(f) \in \mathbb{C}^{I_p + J_p}$ .  $\hat{\mathbf{h}}_p(f)$  and  $\hat{\mathbf{r}}_p(f)$  are both estimated filters to be described in Section 4.3.

Following Wang, Wichern, and Le Roux (2021c),  $\mathcal{F}(\cdot, \cdot)$  computes a loss on the estimated RI components and their magnitude:

$$\mathcal{F}(Y_r(t, f), \hat{Y}_r(t, f)) = \frac{1}{\sum_{t', f'} |Y_r(t', f')|} \mathcal{G}(Y_r(t, f), \hat{Y}_r(t, f)), \quad (7)$$

$$\begin{aligned} \mathcal{G}(Y_r(t, f), \hat{Y}_r(t, f)) &= |\mathcal{R}(Y_r(t, f)) - \mathcal{R}(\hat{Y}_r(t, f))| \\ &\quad + |\mathcal{I}(Y_r(t, f)) - \mathcal{I}(\hat{Y}_r(t, f))| \\ &\quad + ||Y_r(t, f)| - |\hat{Y}_r(t, f)||, \end{aligned} \quad (8)$$

where  $r \in \{0, 1, \dots, P\}$  indexes all the close-talk and far-field microphones,  $|\cdot|$  computes magnitude,  $\mathcal{R}(\cdot)$  and  $\mathcal{I}(\cdot)$  respectively extract RI components, and the denominator in (7) balances the losses at different microphones and across training mixtures.

Notice that the DNN can use all or a subset of the far-field microphone signals as the input and for loss computation. For example,

a monaural enhancement model can be trained by just using the reference microphone signal as the input but computing the loss on all the microphone signals.

#### 4.3. FCP for filter estimation

To compute  $\mathcal{L}_{MC}$ , we need to first compute the linear filters (i.e., RTFs) in (6). Following UNSSOR (Wang & Watanabe, 2023), we leverage FCP (Wang, Wichern, & Le Roux, 2021a, 2021b) for filter estimation, based on the DNN estimates and observed mixtures.

Assuming that the target speaker is non-moving within each utterance, we estimate, e.g., the filter  $\hat{\mathbf{h}}_p(f)$  in (6), by solving the following problem:

$$\hat{\mathbf{h}}_p(f) = \underset{\mathbf{h}_p(f)}{\operatorname{argmin}} \sum_t \frac{|Y_p(t, f) - \mathbf{h}_p(f)^H \bar{\mathbf{X}}_q(t, f; p)|^2}{\hat{\lambda}_p(t, f)}, \quad (9)$$

where  $\hat{\lambda}_p$  to be described in (10), is a weighting term. The objective in (9) is quadratic, where a closed-form solution can be readily computed. We use the same method in (9) (i.e., linearly projecting DNN estimate to observed mixture) to compute all the other filters, and then plug the closed-form solutions to compute the  $\mathcal{L}_{MC}$  loss and train the DNN.

In (9),  $\hat{\lambda}$  is a weighting term balancing the importance of each T-F unit, as different T-F units usually have diverse energy levels. Following Wang, Wichern, and Le Roux (2021a), it is defined as

$$\hat{\lambda}_r(t, f) = \xi \times \max(|Y_r(t, f)|^2 + |Y_r(t, f)|^2), \quad (10)$$

where  $r$  indexes all the microphones,  $\xi$  (tuned to  $10^{-2}$  in this study) floors the weighting term, and  $\max(\cdot)$  extracts the maximum value of a power spectrogram. We compute  $\hat{\lambda}$  differently for different microphones, as the energy level of each source can be very different at close-talk and far-field microphones, and deployed microphones, even if placed close to each other, often produce very different gain levels in real-world scenarios.

#### 5. On-the-fly estimation of future taps $J_0$ for close-talk microphone

The FCP filter taps used in M2M training (i.e.,  $I_p$  and  $J_p$  for  $p \in \{0, 1, \dots, P\}$  and  $p \neq q$ ) are hyperparameters to tune. Their ideal values are likely different for different utterances, while it is tricky and cumbersome to tune each one of them individually for each utterance. For simplicity, in UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a), the value of each filter tap is configured shared for all the training utterances.

This strategy should be improved when dealing with real-recorded data, such as CHiME-4 (Vincent et al., 2017). In CHiME-4, we observe that the far-field microphones are reasonably synchronized as they are placed on, and processed by, the same device, but the close-talk microphone, placed on a different device, is not accurately synchronized with the far-field microphones. Their time-misalignment can be as large as 50 ms. On the other hand, the distance between the far-field microphone array and the close-talk microphone is unknown and can vary from utterance to utterance.

In this context, for simplicity, for all the training utterances, we set the past filter taps the same for all the non-reference microphones (i.e., all the  $I_p$  for  $p \in \{0, 1, \dots, P\}$  and  $p \neq q$  are configured the same) and set the future tap  $J_p$  to 1 for all the non-reference far-field microphones, while we propose to, at each training step, estimate  $J_0$  (i.e., the future taps for the close-talk microphone) for each training utterance in the mini-batch by solving the problem below. The problem below, following the  $\mathcal{L}_{MC,0}$  loss in (6), enumerates a set of future taps and finds the one that leads to the best approximation of the close-talk mixture based on very short FCP filters:

$$J_0 = \underset{J'_0 \in \Omega}{\operatorname{argmin}} \sum_{t, f} \mathcal{F} \left( Y_0(t, f), \hat{\mathbf{h}}_0(f)^H \bar{\mathbf{X}}_q(t, f; 0) + \hat{\mathbf{r}}_0(f)^H \bar{\mathbf{V}}_q(t, f; 0) \right), \quad (11)$$

where  $\bar{\mathbf{X}}_q(t, f; 0) = [\hat{X}_q(t + J'_0 - K + 1, f), \dots, \hat{X}_q(t + J'_0, f)]^T \in \mathbb{C}^K$ ,  $\bar{\mathbf{V}}_q(t, f; 0) = [\hat{V}_q(t + J'_0 - K + 1, f), \dots, \hat{V}_q(t + J'_0, f)]^T \in \mathbb{C}^K$ ,  $K$  is set to a small value (3 in this study) so that the filter is short and the amount of computation spent on solving this problem is small,  $\Omega = \{0, 1, \dots, R\}$  denotes a set of future taps to enumerate (with  $R$  tuned to 8 for CHiME-4 to account for potentially large errors in synchronization),  $J'_0 \in \Omega$  denotes an enumerated candidate future tap, and  $\hat{\mathbf{h}}_0(f)$  and  $\hat{\mathbf{r}}_0(f)$  are computed in the same way as in (9).

Note that we run (11) at each training step to estimate  $J_0$  for each training utterance in the mini-batch. We stop gradients for the operations in (11) in the forward pass, and hence no back-propagation is performed for the operations in (11). The estimated  $J_0$  is then used for computing the  $\mathcal{L}_{MC,0}$  loss in (6).

#### 6. SuperM2M

We first point out the weaknesses of M2M training, and then propose SuperM2M to address the weaknesses. Next, we discuss the necessity of close-talk mixtures in SuperM2M.

##### 6.1. Weaknesses of M2M training

M2M is a weakly-supervised speech enhancement algorithm that can learn from the weak supervision afforded by close-talk and far-field mixtures. It can also be viewed, with a grain of salt, as an unsupervised enhancement algorithm, where the DNN is trained to produce two source estimates that can be linearly filtered to best *explain* (i.e., reconstruct) the close-talk and far-field mixtures. In this regard, the resulting enhancement system needs to deal with three tricky issues.

First, the source estimates could be permuted randomly. That is, they could respectively correspond to speech and noise, or the opposite, since the two estimates and their linearly-filtered results are only constrained to sum up to the mixtures.

Second, the source estimates could suffer from frequency permutation (Sawada, Ono, Kameoka, Kitamura, & Saruwatari, 2019), a common problem that needs to be dealt with in many frequency-domain unsupervised separation algorithms such as independent vector analysis, spatial clustering, and UNSSOR (Wang & Watanabe, 2023). Since FCP is performed in each frequency independently from the others, even though speech and noise sources are accurately separated in each frequency, the separation results of each source at different frequencies are not guaranteed to be grouped into the same output spectrogram.

Third, since, in realistic cases, the noise component  $V$  usually consists of an unknown number of directional and diffuse sources, in unsupervised separation the model would lack an idea to produce one estimate exactly corresponding to target speech and the other exactly corresponding to all the noise sources combined. In other words, sources are ambiguous to the model. It is possible that, even if one estimate contains the target speech plus some noise sources and the other estimate absorbs the rest noise sources, the mixture-constraint loss can still be very low. The fundamental causes of this problem are that, in unsupervised setups, (a) the model lacks an exact concept about what the target source should be like; and (b) the hypothesized number sources (in this paper, 2) is not guaranteed to match the actual number of sources (i.e., speech source plus an unknown number of noise sources) in every training mixture.

These issues do not exist in supervised approaches, as the oracle simulated speech and noise signals used in supervised approaches can penalize the DNN estimates to naturally avoid source and frequency permutation, and resolve source ambiguity. This motivates us to combine M2M training with supervised learning, leading to SuperM2M, which is described next.

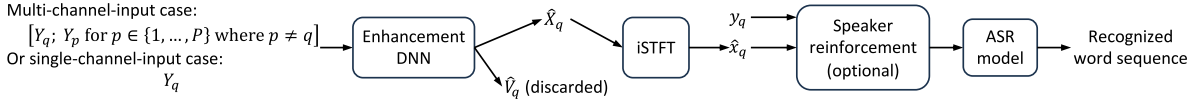


Fig. 2. Robust ASR pipeline, where enhanced speech  $\hat{x}_q = \text{iSTFT}(\hat{X}_q)$  is fed to backend ASR models for recognition. No joint training is performed. An optional speaker reinforcement module (Zorilá & Doddipatla, 2022), which adds a scaled version of the input mixture signal  $y_q$  to  $\hat{x}_q$ , can be included.

## 6.2. Supervised and mixture-to-mixture co-learning

The previous subsection points out that M2M suffers from source and frequency permutation, and source ambiguity. On the other hand, although M2M training can be performed on real mixtures, there may not be many paired close-talk and far-field real-recorded mixtures available, as collecting real data is effort-consuming. In comparison, supervised models can be readily trained on massive simulated mixtures, as one can easily simulate as many mixtures as one considers sufficient. In addition, they do not suffer from source and frequency permutation, and source ambiguity.

In this context, we propose to train the same DNN model with both M2M training and supervised learning to combine their strengths. We name the algorithm SuperM2M. See Fig. 1 for an illustration, where the supervised learning part is shown in Fig. 1(b). Notice that the DNN in M2M training is designed to directly produce target and non-target estimates. This makes M2M training capable of being easily integrated with supervised training, where the models are usually designed to directly produce target estimates.

In detail, at each training step, we sample either a mini-batch of real close-talk and far-field mixture pairs or a mini-batch of simulated far-field mixtures for DNN training. The loss for the mini-batch of real data is  $\mathcal{L}_{\text{MC}}$  in (4), and the loss for the mini-batch of simulated data is

$$\mathcal{L}_{\text{SIMU},q} = \mathcal{L}_{\text{Target},q} + \mathcal{L}_{\text{Non-target},q}, \quad (12)$$

$$\mathcal{L}_{\text{Target},q} = \frac{1}{\sum_{t,f} |Y_q(t,f)|} \sum_{t,f} \mathcal{G}(X_q(t,f), \hat{X}_q(t,f)), \quad (13)$$

$$\mathcal{L}_{\text{Non-target},q} = \frac{1}{\sum_{t,f} |Y_q(t,f)|} \sum_{t,f} \mathcal{G}(V_q(t,f), \hat{V}_q(t,f)), \quad (14)$$

where  $\mathcal{G}(\cdot, \cdot)$  is defined in (8),  $X_q$  and  $V_q$  are obtained through simulation, and the denominator balances the loss values with the ones in M2M training.

## 6.3. Necessity of close-talk mixtures

So far, we hypothesize that, during training, a paired close-talk mixture is always available for far-field mixtures. It is leveraged as a weak supervision for training by optimizing a mixture-constraint loss (i.e.,  $\mathcal{L}_{\text{MC},0}$  in (6)) defined on it.

When close-talk mixtures are not available, we find that we can still train enhancement models successfully via SuperM2M, where, in the M2M part, the DNN is trained to only recover far-field mixtures, meaning that M2M training is unsupervised.<sup>5</sup> This is a desirable property, as this means that we only need a set of real-recorded far-field multi-channel mixtures (which are easier to record than paired far-field and close-talk mixtures), and together with a set of simulated mixtures, we can train an enhancement system via SuperM2M, which could generalize better to real mixtures than purely-supervised models trained only on the simulated mixtures.

<sup>5</sup> When close-talk mixtures are not available, M2M (Wang, 2024a) regresses to UNSSOR (Wang & Watanabe, 2023). In our paper, we prefer to still call our algorithms SuperM2M, rather than SuperUNSSOR, just to avoid creating too many new names.

## 7. Experimental setup

Our main goal is to show that SuperM2M can generalize better to real data than purely supervised models trained on simulated data. We follow the robust ASR pipeline in Fig. 2 for evaluation, not using any joint frontend-backend training.

We do not use  $\hat{X}_q$  to derive linear beamforming results for ASR (Haeb-Umbach et al., 2020, 2019; Wang & Chen, 2018), although this has been extremely popular, as we would like to validate whether the enhanced speech  $\hat{X}_q$  itself is close to target speech and whether  $\hat{X}_q$  itself can yield better ASR performance. We do not jointly train enhancement models with ASR models, as this requires knowledge of ASR models and would not accurately reflect the accuracy of  $\hat{X}_q$  itself. We aim at building enhancement models that can produce enhanced speech with low distortion and high reduction to non-target signals. This way, the enhancement models could improve the robustness of many subsequent applications not limited to ASR.

In a nutshell, our main goal is to show, through SuperM2M, whether  $\hat{X}_q$  itself would be better on real test data. We validate SuperM2M on CHiME-4 (Vincent et al., 2017), a dataset consisting of simulated mixtures and real-recorded close-talk and far-field mixture pairs. To further show the effectiveness and potential of SuperM2M, a minor goal is to show whether SuperM2M can lead to state-of-the-art ASR performance on CHiME-4. The rest of this section describes the CHiME-4 dataset, miscellaneous system configurations, comparison systems, evaluation metrics, and several tricks to improve robust ASR performance.

### 7.1. CHiME-4 dataset

CHiME-4 (Vincent et al., 2017) is a major corpus for evaluating robust ASR and speech enhancement algorithms. It is recorded by using a tablet mounted with 6 microphones, with the second microphone on the rear and the others facing front. The signals are recorded in four representative environments (including cafeteria, buses, pedestrian areas, and streets), where reverberation and directional, diffuse, transient and non-stationary noises naturally exist. During data collection, the target speaker hand-holds the tablet in a designated environment, and reads text prompts shown on the screen of the tablet. The target speaker wears a close-talk microphone so that the close-talk mixture can be recorded at the same time along with far-field mixtures recorded by the microphones on the tablet. The number of simulated and real-recorded utterances is listed in Table 1. It should be noted that, in CHiME-4, the room reverberation is weak, and the major challenge is in how to remove the multi-source non-stationary noise signals.

In the real data of CHiME-4, we observe synchronization errors between the close-talk microphone and far-field microphone array. Other issues, such as microphone failures, signal clipping, speaker and array movement, and diverse gain levels even if microphones are placed close to each other, happen frequently. In real-world products, these are typical problems, which increase the difficulties of speech enhancement and ASR. They need to be robustly dealt with by frontend enhancement systems.

Depending on the number of microphones that can be used for recognition, there are three official ASR tasks in CHiME-4, including 1-, 2- and 6-channel tasks. In the 1-channel task, only one of the front microphones can be used for testing; in the 2-channel task, only two of the front microphones can be used; and in the 6-channel task, all the six microphones can be used. For the 1- and 2-channel tasks, the microphones that can be used for ASR for each utterance are selected by the challenge organizers to avoid microphone failures. The selected microphones can vary from utterance to utterance.

**Table 1**  
Number of utterances in CHiME-4 (all are six-channel).

Type	Training Set	Validation Set	Test Set
SIMU	7,138 (~ 15.1 h)	1,640 (~ 2.9 h) (410 in each environ.)	1,320 (~ 2.3 h) (330 in each environ.)
REAL	1,600 (~ 2.9 h)	1,640 (~ 2.7 h) (410 in each environ.)	1,320 (~ 2.2 h) (330 in each environ.)

## 7.2. Evaluation setup — robust ASR

We check whether SuperM2M can improve ASR performance by feeding its enhanced speech to ASR models for decoding, following the pipeline in Fig. 2. We consider two ASR models.

- The first one is Whisper Large v2<sup>6</sup> (Radford et al., 2023), pre-trained on massive data. We use its text normalizer to normalize hypothesis and reference text before computing WER.
- The second one is trained on the official CHiME-4 mixtures plus the clean signals in WSJ0 by using the public recipe (Chang et al., 2022)<sup>7</sup> in ESPnet. It is an encoder-decoder transformer-based model, trained on WavLM features (Chen et al., 2022) and using a transformer language model in decoding. It is the current strongest ASR model on CHiME-4. It is a representative model in leveraging self-supervised learning based representations for ASR and robust ASR (Chang et al., 2022; Du, Zhang, Fang, Wu, & Yang, 2023; Zhu, Zhang, Zhang, & Dai, 2023; Zhu et al., 2022).

Note that the WERs computed by ESPnet should not be directly compared with the ones by Whisper due to different text normalization. As a side result, to compare the ASR model in ESPnet with the Whisper model, we retrained an ESPnet model by using the same text normalization as the Whisper ASR model. We will report the evaluation results of using this ASR model later in Section 8.9.

## 7.3. Evaluation setup — speech enhancement

We evaluate the enhancement performance of SuperM2M on the simulated test data of CHiME-4. We consider 1- and 6-channel enhancement. In the 1-channel case, SuperM2M uses the fifth microphone (CH5) signal as input, and the target direct-path signal at CH5 is used as the reference for evaluation. In the 6-channel case, SuperM2M uses all the microphone signals as input to predict the target speech at CH5.

The evaluation metrics include wide-band perceptual evaluation of speech quality (WB-PESQ) (Rix, Beerends, Hollier, & Hekstra, 2001), short-time objective intelligibility (STOI) (H. Taal, C. Hendriks, Heusdens, & Jensen, 2011), signal-to-distortion ratio (SDR) (Vincent, Gri-bonval, & Févotte, 2006), and scale-invariant SDR (SISDR) (Le Roux, Wisdom, Erdogan, & Hershey, 2019). They are widely-adopted metrics in speech enhancement, which can evaluate the quality, intelligibility, and accuracy of the magnitude and phase of enhanced speech.

## 7.4. Training setup

For monaural enhancement, we train SuperM2M using all the (7, 138+1, 600)×6 monaural signals. For 2-channel enhancement, at each training step we sample 2 microphones from the front microphones as input, and train the DNN to predict the target speech at the first of the selected microphones. For 6-channel enhancement, we train SuperM2M using the 7, 138 + 1, 600 six-channel signals. The DNN stacks all the six microphones as input to predict the target speech at CH5.

For simplicity, we do not filter out microphone signals with any microphone failures in the DNN input and loss. We would expect SuperM2M to learn to deal with the failures, as it can be trained on real mixtures.

## 7.5. Miscellaneous configurations

For STFT, the window size is 32 ms, hop size 8 ms, the square root of Hann window is used as the analysis window, and the synthesis window is designed based on the analysis window to achieve perfect reconstruction.

TF-GridNet (Wang et al., 2023), which has shown strong separation performance in major benchmarks in supervised speech separation, is used as the DNN architecture. We consider two setups. Using the symbols defined in Table I of Wang et al. (2023), the first one (denoted as TFGridNetv1) sets its hyper-parameters to  $D = 100$ ,  $B = 4$ ,  $I = 2$ ,  $J = 2$ ,  $H = 200$ ,  $L = 4$  and  $E = 2$ , and the second one (denoted as TFGridNetv2) to  $D = 128$ ,  $B = 4$ ,  $I = 1$ ,  $J = 1$ ,  $H = 200$ ,  $L = 4$  and  $E = 4$ . Please do not confuse these symbols with the ones in this paper. The models have ~6.3 and ~5.4 million parameters respectively. The v1 model uses approximately half of the computation and memory of v2,<sup>8</sup> and is utilized for faster experimentation.

We train all the enhancement models on 8-second segments using a mini-batch size of 1. At each training step, if the sampled utterance is simulated, we use supervised learning, and if it is real-recorded, we use M2M training. Adam is employed for optimization. The learning rate starts from 0.001 and is halved if the validation loss is not improved in 2 epochs.

## 7.6. Comparison systems

We consider the same DNN model trained only on the simulated data of CHiME-4 via supervised learning as the major baseline for comparison. We use exactly the same configurations as that in SuperM2M for training. We denote this baseline as **Supervised**, to differentiate it with **SuperM2M**. Since CHiME-4 is a public and popular dataset, many existing models can be used directly for comparison.

## 7.7. Tricks to improve robust ASR performance

To show the effectiveness of SuperM2M, we also check whether it can lead to state-of-the-art ASR performance on CHiME-4. This subsection describes several tricks that are known to improve robust ASR performance and are commonly used in existing studies.

### 7.7.1. SNR augmentation for simulated training mixtures

At each training step, we optionally modify the SNR of the target speech in the CHiME-4 simulated training mixture, on the fly, by  $u$  dB, with  $u$  uniformly sampled from the range  $[-10, +5]$  dB. In our experiments, we find this technique often producing slightly better enhancement and ASR, but not critical. Note that we do not change the combinations of speech and noise files to create new mixtures, and we just change the SNR of the existing mixtures in CHiME-4. No other data augmentation is used for enhancement.

<sup>8</sup> This is because in v1, setting  $I = 2$  and  $J = 2$  cuts the sequence length to be modeled by the BLSTMs in TFGridNet by 50%, compared with setting  $I = 1$  and  $J = 1$  in v2. In this case, v1 would approximately use half of the computation and memory of v2 for the following reasons. First, in TFGridNet, the cross-frame self-attention module is much less costly than the sub-band temporal module and the intra-frame full-band module, both of which use BLSTMs, the most computationally-expensive block in each of the two modules. Second, although the input dimension to the BLSTMs is 200 in v1 while 128 in v2, the hidden dimensions of the BLSTMs in v1 and v2 are both set to 200 and most computation inside BLSTMs is not spent between the input tensor and the hidden tensor. We highlight that this approximation is rough and is verified only empirically.

<sup>6</sup> <https://huggingface.co/openai/whisper-large-v2>

<sup>7</sup> [https://github.com/espnet/espnet/blob/master/egs2/chime4/asr1/conf/tuning/train\\_asr\\_transformer\\_wavlm\\_lr1e-3\\_specaug\\_accum1\\_preenc128\\_warmup20k.yaml](https://github.com/espnet/espnet/blob/master/egs2/chime4/asr1/conf/tuning/train_asr_transformer_wavlm_lr1e-3_specaug_accum1_preenc128_warmup20k.yaml)

Table 2

SuperM2M vs. purely-supervised models on CHiME-4 (#input mics: 6; ASR model: Whisper Large v2).

System	Approach	Training data	DNN arch.	$J_0$	#mics in $\mathcal{L}_{MC}$	SNR aug.	iSTFT- STFT proj.?	Spk. reinf. $\gamma$ (dB)	SIMU Test Set (CH5)				Official CHiME-4 Test Utterances			
									SISDR (dB)↑	SDR (dB)↑	WB- PESQ↑	STOI↑	Val. WER (%)↓		Test WER (%)↓	
													SIMU	REAL	SIMU	REAL
0	Mixture	–	–	–	–	–	–	–	7.5	7.5	1.27	0.870	7.43	4.96	10.97	7.69
1a	Supervised	S	TFGridNetv1	–	–	✗	✗	–	22.2	22.6	3.08	0.984	3.55	28.90	3.93	53.23
1b	Supervised	S	TFGridNetv2	–	–	✓	✓	–	<b>22.8</b>	<b>23.2</b>	<b>3.36</b>	<b>0.988</b>	3.43	58.15	3.79	79.08
2a	Supervised	S	iNeuBe Wang et al. (2020)	–	–	–	–	–	22.0	22.4	–	0.986	–	–	–	–
2b	Supervised	S	SpatialNet Quan and Li (2024)	–	–	–	–	–	22.1	22.3	2.88	0.983	–	–	–	–
2c	Supervised	S	USES Zhang et al. (2023)	–	–	–	–	–	–	20.6	3.16	0.983	–	–	4.20	78.10
2d	Supervised	S	USES2 Zhang et al. (2024)	–	–	–	–	–	–	18.8	2.94	0.979	–	–	4.60	12.10
3	SuperM2M	S+R	TFGridNetv1	–	6	✗	✗	–	22.3	22.6	3.12	0.984	3.52	3.93	3.93	4.46
4a	SuperM2M	S+R	TFGridNetv1	1	6+1	✗	✗	–	22.3	22.5	3.06	0.984	3.57	3.89	3.97	4.04
4b	SuperM2M	S+R	TFGridNetv1	2	6+1	✗	✗	–	22.2	22.6	3.11	0.984	3.51	3.84	4.12	4.16
4c	SuperM2M	S+R	TFGridNetv1	3	6+1	✗	✗	–	22.2	22.5	3.07	0.984	3.56	3.88	4.05	4.09
4d	SuperM2M	S+R	TFGridNetv1	4	6+1	✗	✗	–	22.4	22.7	3.11	0.985	3.48	3.97	3.98	4.07
4e	SuperM2M	S+R	TFGridNetv1	5	6+1	✗	✗	–	22.3	22.6	3.07	0.985	3.52	3.97	4.05	4.16
4f	SuperM2M	S+R	TFGridNetv1	6	6+1	✗	✗	–	22.3	22.7	3.10	0.985	3.61	3.94	3.98	4.21
4g	SuperM2M	S+R	TFGridNetv1	7	6+1	✗	✗	–	22.2	22.5	3.10	0.984	3.53	3.99	4.01	4.24
4h	SuperM2M	S+R	TFGridNetv1	8	6+1	✗	✗	–	22.4	22.7	3.14	0.985	3.51	3.93	3.97	4.19
5	SuperM2M	S+R	TFGridNetv1	est.	6+1	✗	✗	–	22.3	22.6	3.07	0.985	3.59	3.92	4.08	<b>3.97</b>
6a	UNSSOR Wang and Watanabe (2023)	R	TFGridNetv1	–	6	✗	✗	–	7.7	8.1	1.30	0.870	4.82	4.16	6.20	4.84
6b	UNSSOR Wang and Watanabe (2023)	S+R	TFGridNetv1	–	6	✗	✗	–	4.5	4.6	1.32	0.881	4.81	4.14	5.42	4.88
6c	M2M Wang (2024a)	R	TFGridNetv1	est.	6+1	✗	✗	–	9.8	10.3	1.50	0.902	4.94	3.95	6.76	4.32
6d	M2M Wang (2024a)	S+R	TFGridNetv1	est.	6+1	✗	✗	–	4.5	4.6	1.33	0.882	4.63	3.95	5.38	4.32
7a	SuperM2M	S+R	TFGridNetv2	est.	6+1	✗	✗	–	22.5	22.9	3.15	0.985	3.51	3.90	3.87	4.14
7b	SuperM2M	S+R	TFGridNetv2	est.	6+1	✗	✓	–	22.7	23.2	3.24	0.986	3.43	3.84	3.87	3.98
7c	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	–	<b>22.8</b>	<b>23.2</b>	3.22	0.987	<b>3.38</b>	<b>3.81</b>	3.77	4.04
8a	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	10	–	–	–	–	3.61	3.85	3.84	4.24
8b	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	15	–	–	–	–	3.59	3.87	<b>3.74</b>	4.11
8c	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	20	–	–	–	–	3.54	3.86	<b>3.74</b>	4.10
9	Close-talk Mixture	–	–	–	–	–	–	–	–	–	–	–	–	3.76	–	3.91

### 7.7.2. iSTFT-STFT projection

We apply inverse STFT (iSTFT) followed by STFT operations to the DNN estimates  $\hat{X}_q$  and  $\hat{V}_q$  before loss computation, i.e.,  $\hat{X}_q := \text{STFT}(\text{iSTFT}(\hat{X}_q))$  and  $\hat{V}_q := \text{STFT}(\text{iSTFT}(\hat{V}_q))$ . See Fig. 1 for an illustration. This often yields slight improvement, as the losses now penalize the RI components and magnitudes extracted from re-synthesized signals, which are the final system output used for human hearing and downstream tasks (Pandey & Wang, 2019; Wang, Wichern, & Le Roux, 2021c; Wisdom et al., 2019). Notice that, in Fig. 2, ASR features are extracted from re-synthesized signals.

### 7.7.3. Run-time speaker reinforcement for robust ASR

At run time, in default we feed  $\hat{x}_q$  for ASR. Alternatively, we employ a technique named *speaker reinforcement* (Zoril  & Doddipatla, 2022), where  $\hat{x}_q$  is re-mixed with the input mixture  $y_q$  at an energy level of  $\gamma$  dB before recognition. See Fig. 2 for an illustration. That is,  $\hat{x}_q + \eta \times y_q$ , where  $\eta \in \mathbb{R}_{>0}$  and  $\gamma = 10 \times \log_{10}(\|\hat{x}_q\|^2 / \|\eta \times y_q\|^2)$ . We find this technique usually effective for ASR, as the re-mixed input mixture can alleviate distortion to target speech.

## 8. Evaluation results

This section reports our evaluations results on CHiME-4. Following earlier studies (Wang, 2024a; Wang & Watanabe, 2023), we set  $I_p$  to 20 for all the non-reference far-field microphones and the close-talk microphone, and  $J_p$  to 1 for all the non-reference far-field microphones.  $J_0$ , the future filter tap for the close-talk microphone, can be tuned to a value shared by all the training mixtures, or it can be estimated, on the fly, for each training mixture by using the method described in Section 5. All the six far-field microphones are used for computing  $\mathcal{L}_{MC}$  in (4). We emphasize that we spent minimal amount of effort on hyper-parameter tuning. Other hyper-parameter setup could lead to better performance.

### 8.1. SuperM2M vs. Purely-supervised models

Tables 2 and 3 respectively report 6- and 1-channel enhancement performance on the fifth microphone of the CHiME-4 simulated test data, and robust ASR performance on the official CHiME-4 test utterances (by using the Whisper Large v2 model for recognition). The two tables consist of exactly the same set of experiments, and differ only in the number of input microphones to the enhancement models.

On the simulated test data, purely supervised models (in system 1a-1b) trained on the official simulated training data (denoted as S in the ‘‘Training data’’ column) produce large improvement over unprocessed mixtures (e.g., in system 1a of Table 2, 22.2 vs. 7.5 dB SISDR, and in 1a of Table 3, 17.1 vs. 7.5 dB SISDR), and achieve enhancement performance better than existing supervised models such as iNeuBe (based on TCN-DenseUNet) (Lu et al., 2022; Wang et al., 2020), SpatialNet (Quan & Li, 2024), USES (Zhang et al., 2023) and USES2 (Zhang et al., 2024) in both 1- and 6-channel cases. However, the ASR performance is much worse than directly using unprocessed mixtures for ASR, especially in multi-channel cases (e.g., in system 1a and 0 of Table 2, 53.23% vs. 7.69% WER on REAL test data, and in system 1a and 0 of Table 3, 10.20% vs. 7.69% WER on REAL test data). This degradation has been widely observed by existing studies (Haeb-Umbach et al., 2020, 2019), largely due to the mismatches between simulated training and real-recorded test conditions, and the speech distortion incurred by enhancement.

In system 7c of Table 2 and 7b of Table 3, we respectively show the results of our best performing SuperM2M models (without employing the speaker reinforcement technique described in Section 7.7.3) for 6- and 1-channel cases. By training on the official simulated and real data combined (denoted as S+R), SuperM2M obtains clearly better ASR results on the real test data than the purely-supervised models (in system 1a-1b) and unprocessed mixtures (in system 0), and the enhancement

performance on the simulated test data remains strong. These results indicate that SuperM2M can effectively learn from real data, has better generalizability to real data, and can perform enhancement with low distortion to target speech and high reduction to non-target signals.

Next, we present some ablations results of SuperM2M.

### 8.2. Effects of including loss on close-talk mixture in $\mathcal{L}_{MC}$

First, in system 3, we report the results of not including close-talk microphone in the loss function. That is, in (4) we do not include the  $\mathcal{L}_{MC,0}$  loss, when training SuperM2M. We observe clear improvement in ASR performance on the real test data over the purely-supervised model in system 1a (i.e., 4.46% vs. 53.23% in Table 2 and 7.02% vs. 10.20% in Table 3), indicating that close-talk mixtures are not must-have for our system.

Next, we include close-talk mixtures for training. In system 4a-4h, the number of future taps for the close-talk microphone,  $J_0$ , is tuned to a value shared by all the training utterances, and we enumerate  $J_0$  from 1 all the way up to 8. In system 5, we instead estimate  $J_0$ , on the fly, for each training mixture during training, using the technique described in Section 5. Compared with system 4a-4h, system 5 obtains better ASR performance on the real test data, indicating the effectiveness of estimating  $J_0$ .

In both Tables 2 and 3, comparing system 5 and 3, we observe that including close-talk mixtures for training SuperM2M produces better ASR performance on the real test data of CHiME-4 (e.g., 3.97% vs. 4.46% in Table 2 and 6.51% vs. 7.02% in Table 3).

### 8.3. SuperM2M vs. UNSSOR and M2M

In system 6a-6d of Tables 2 and 3, we respectively report the results of UNSSOR (Wang & Watanabe, 2023) and M2M (Wang, 2024a). UNSSOR is trained in an unsupervised way directly on far-field mixtures by using the  $\mathcal{L}_{MC}$  loss in (4) but without including  $\mathcal{L}_{MC,0}$  (i.e., the loss on close-talk mixture). M2M is the same as SuperM2M but without the supervised learning branch. We can train UNSSOR and M2M on the real data alone (denoted as R), or on the simulated and real data combined (i.e., S+R) considering that SuperM2M is trained on S+R. When M2M is trained on S+R, the loss function on the simulated data does not include the loss on close-talk mixture, since CHiME-4 does not provide simulated close-talk mixtures.

Since the source estimates produced by UNSSOR and M2M could suffer from frequency permutation, we employ an existing algorithm (Sawada, Araki, & Makino, 2007; Vu & Haeb-Umbach, 2010)<sup>9</sup> for frequency alignment. The algorithm exploits inter-frequency correlation of estimated source posteriors, and we compute the source posteriors based on the DNN estimates  $\hat{S}_q$  and  $\hat{V}_q$  respectively via  $|\hat{S}_q|/(|\hat{S}_q|+|\hat{V}_q|)$  and  $|\hat{V}_q|/(|\hat{S}_q|+|\hat{V}_q|)$ . In addition, since the source estimates of UNSSOR and M2M exhibit source permutation, for each evaluation metric we compute a score for each estimate and select the better score.

Comparing system 6a-6d and 5, we can see that SuperM2M obtains much better performance than UNSSOR and M2M, suggesting the effectiveness of the proposed co-learning mechanism. The low performance of M2M and UNSSOR is possibly because the noise in CHiME-4 typically consists of an unknown number of diffuse and directional sources. In this case, unsupervised algorithms with two hypothesized directional sources tend to get confused about which source should be the target speech source. This problem is naturally avoided by the supervised learning mechanism in SuperM2M, which models noise sources as a single, combined source.

<sup>9</sup> See [https://github.com/fngt/pb\\_bss/blob/master/pb\\_bss/permutation\\_alignment.py](https://github.com/fngt/pb_bss/blob/master/pb_bss/permutation_alignment.py)

Table 3

SuperM2M vs. purely-supervised models on CHiME-4 (#input mics: 1; ASR model: Whisper Large v2).

System	Approach	Training data	DNN arch.	$J_0$	#mics in $\mathcal{L}_{MC}$	SNR aug.	iTSTFT-STFT proj.?	Spk. reinf. $\gamma$ (dB)	SIMU Test Set (CH5)				Official CHiME-4 Test Utterances			
									SISDR (dB)↑	SDR (dB)↑	WB-PESQ↑	STOI↑	Val. WER (%)↓		Test WER (%)↓	
													SIMU	REAL	SIMU	REAL
0	Mixture	–	–	–	–	–	–	–	7.5	7.5	1.27	0.870	7.43	4.96	10.97	7.69
1a	Supervised	S	TFGridNetv1	–	–	✗	✗	–	17.1	17.5	2.44	0.960	7.14	5.33	13.16	10.20
1b	Supervised	S	TFGridNetv2	–	–	✓	✓	–	17.1	17.5	2.44	0.961	7.58	5.19	12.44	9.03
2	Supervised	S	iNeuBe Wang et al. (2020)	–	–	–	–	–	15.1	–	–	0.954	–	–	–	–
3	SuperM2M	S+R	TFGridNetv1	–	6	✗	✗	–	16.8	17.3	2.40	0.960	7.13	5.17	11.80	7.02
4a	SuperM2M	S+R	TFGridNetv1	1	6+1	✗	✗	–	16.8	17.4	2.38	0.961	7.10	5.05	11.87	6.93
4b	SuperM2M	S+R	TFGridNetv1	2	6+1	✗	✗	–	16.9	17.5	2.45	0.962	6.99	5.19	11.67	6.87
4c	SuperM2M	S+R	TFGridNetv1	3	6+1	✗	✗	–	16.6	17.5	2.47	0.961	7.02	4.86	12.26	6.86
4d	SuperM2M	S+R	TFGridNetv1	4	6+1	✗	✗	–	16.8	17.5	2.48	<b>0.963</b>	7.10	4.97	11.75	6.87
4e	SuperM2M	S+R	TFGridNetv1	5	6+1	✗	✗	–	16.9	17.4	2.36	0.959	7.34	5.29	12.23	7.51
4f	SuperM2M	S+R	TFGridNetv1	6	6+1	✗	✗	–	16.9	17.4	2.41	0.962	7.34	5.29	12.23	7.51
4g	SuperM2M	S+R	TFGridNetv1	7	6+1	✗	✗	–	17.0	17.4	2.42	0.961	7.23	5.02	12.19	6.90
4h	SuperM2M	S+R	TFGridNetv1	8	6+1	✗	✗	–	16.6	17.4	2.41	0.961	7.26	5.03	12.05	6.95
5	SuperM2M	S+R	TFGridNetv1	est.	6+1	✗	✗	–	16.4	17.4	<b>2.56</b>	0.962	7.01	4.87	11.69	6.51
6a	UNSSOR Wang and Watanabe (2023)	R	TFGridNetv1	–	6	✗	✗	–	10.3	11.0	1.42	0.898	11.02	6.24	15.04	10.42
6b	UNSSOR Wang and Watanabe (2023)	S+R	TFGridNetv1	–	6	✗	✗	–	11.6	11.7	1.70	0.937	7.71	5.24	11.85	8.12
6c	M2M Wang (2024a)	R	TFGridNetv1	est.	6+1	✗	✗	–	11.6	12.3	1.77	0.924	10.72	5.95	14.98	8.61
6d	M2M Wang (2024a)	S+R	TFGridNetv1	est.	6+1	✗	✗	–	10.3	10.4	1.77	0.942	7.70	5.66	11.94	8.00
7a	SuperM2M	S+R	TFGridNetv2	est.	6+1	✗	✗	–	<b>17.3</b>	<b>17.7</b>	2.44	<b>0.963</b>	6.78	4.93	11.42	6.31
7b	SuperM2M	S+R	TFGridNetv2	est.	6+1	✗	✓	–	17.1	17.5	2.43	0.962	7.02	4.74	11.24	6.23
7c	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	–	16.6	17.4	2.51	<b>0.963</b>	7.05	4.78	11.34	6.02
8a	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	10	–	–	–	–	<b>5.90</b>	<b>4.42</b>	<b>8.80</b>	<b>5.80</b>
8b	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	15	–	–	–	–	6.22	4.47	9.34	5.76
8c	SuperM2M	S+R	TFGridNetv2	est.	6+1	✓	✓	20	–	–	–	–	6.40	4.60	9.87	5.86
9	Close-talk Mixture	–	–	–	–	–	–	–	–	–	–	–	–	3.76	–	3.91

Table 4

SuperM2M vs. USES (Zhang et al., 2023) and USES2 (Zhang et al., 2024) on CHiME-4 (ASR model: Whisper Large v2).

System	Cross reference	Approach	Training data	Task	SIMU Test (CH5)				CH5 of CHiME-4 Test Utterances			
					SISDR (dB)↑	SDR (dB)↑	WB-PESQ↑	STOI↑	Val. SIMU	WER (%)↓	Test REAL	WER (%)↓
0	–	Mixture	–	1-channel	7.5	7.5	1.27	0.870	5.49	5.09	5.82	6.69
1a	–	USES Zhang et al. (2023)	S	1-channel	–	–	–	–	–	–	–	11.00
1b	–	USES Zhang et al. (2023)	S+extra	1-channel	–	–	–	–	–	–	–	7.10
2	7c of Table 3	SuperM2M	S+R	1-channel	16.6	17.4	2.51	0.963	5.37	4.84	6.14	5.73
3a	–	USES Zhang et al. (2023)	S	6-channel	–	20.6	3.16	0.983	–	–	4.20	78.10
3b	–	USES Zhang et al. (2023)	S+extra	6-channel	–	19.1	2.95	0.979	–	–	4.10	85.90
3c	–	USES2 Zhang et al. (2024)	S	6-channel	–	18.8	2.94	0.979	–	–	4.60	12.10
3d	–	USES2 Zhang et al. (2024)	S+extra	6-channel	–	–	–	–	–	–	–	10.30
4	7c of Table 2	SuperM2M	S+R	6-channel	22.8	23.2	3.22	0.987	3.38	3.81	3.77	4.04

Notes: (a) The “extra” in “SIMU+extra” means extra ~ 230 hours of simulated training data (see Zhang et al. (2024, 2023) for details).

(b) The “cross reference” entry means that the other setups are the same as the ones in the referred row.

#### 8.4. Miscellaneous results

In system 7a, 7b, and 7c of Tables 2 and 3, we switch TF-GridNet from v1 to v2, perform iSTFT-STFT projection, and apply SNR augmentation. Comparing system 7a with 7b, we observe that applying iSTFT-STFT projection can improve ASR performance in most cases. From system 5 to 7c, although the ASR performance on the real test data gets better in Table 3 while worse in Table 2, in both tables better ASR performance is observed on the real validation set. We therefore use system 7c as our default system for the rest experiments of this paper.

#### 8.5. SuperM2M vs. Purely large-scale supervised training

Table 4 compares SuperM2M with a representative line of research (USES Zhang et al., 2023 and USES2 Zhang et al., 2024), which trains enhancement models in a purely-supervised way on a much larger-scale simulated data of ~245 hours, which is ~14 times the size of the CHiME-4 SIMU+REAL data (i.e., 245/(15.1 + 2.9)).

USES and USES2 have shown that, by increasing the diversity of simulated training mixtures to cover as many conditions (that could happen in real data) as possible, better enhancement can be achieved on real data. However, on CHiME-4, compared with using unprocessed mixtures directly for recognition, they do not obtain better ASR performance, likely because the simulated training data is not representative of the real data in CHiME-4.

In comparison, SuperM2M, although trained only on the official small-scale CHiME-4 simulated and real mixtures, obtains much better ASR performance on the real data of CHiME-4 in both single- and multi-channel cases than USES, USES2, and the unprocessed mixtures. This comparison does not suggest that purely supervised learning on large-scale simulated data is a bad idea, as it can offer an easy and feasible way for training DNNs to model speech patterns for enhancement, and building upon this strength, SuperM2M can very likely produce even better enhancement on real data. Rather, it sounds an alarm that purely large-scale supervised learning on simulated data has a fundamental limitation incurred by using the current simulation techniques, which usually cannot simulate mixtures in a sufficiently realistic way. In addition, it suggests the benefits of co-training the DNN on real data via M2M. This way, the DNN, during training, can see and learn from the signal characteristics of real-recorded data, and could hence generalize better to real-recorded data.

#### 8.6. Breaking out to new highs on CHiME-4 ASR tasks

Table 5 reports the ASR performance of SuperM2M based on the official ASR setup of CHiME-4, using the ASR model proposed in Chang et al. (2022) and still following the evaluation pipeline in Fig. 2. Comparing system 0b with 0a, we observe that we have successfully reproduced the ASR system proposed in Chang et al. (2022).

SuperM2M, despite not jointly trained with the ASR model, achieves a new state-of-the-art on the REAL test set in each of the 1-, 2- and 6-channel tasks, significantly outperforming the previous best obtained by IRIS (Chang et al., 2022) and MultiIRIS (Masuyama, Chang, Cornell, et al., 2023) (e.g., in the 1-channel case 3.04% vs. 3.92% WER in 2a and 1b, in the 2-channel case 1.94% vs. 2.65% WER in 4a and 3b, and in the 6-channel case 1.61% vs. 1.77% WER in 6a and 5b). IRIS jointly trains a Conv-TasNet based monaural enhancement model, a WavLM based ASR feature extractor, and an encoder-decoder transformer based ASR model. MultiIRIS, building upon IRIS, replaces the Conv-TasNet module with a DNN based weighted power minimization distortionless response (WPD) beamformer. From system 1a vs. 1b, 3a vs. 3b, and 5a vs. 5b, we observe that, without joint training, IRIS and MultiIRIS often obtain clearly worse ASR performance, especially in the 1- and 2-channel cases. These results further indicate the effectiveness and potential of SuperM2M on real data, as it decouples enhancement and ASR and does not employ joint training.

#### 8.7. Effects of speaker reinforcement

In system 8a-8c of Tables 2 and 3, where the Whisper ASR model is used for recognition, we apply speaker reinforcement with the energy level  $\gamma$  between the enhancement output and input mixture tuned based on the set of {10, 15, 20} dB. Better ASR performance is observed in the 1-channel case but not in the 6-channel case, possibly because the enhanced speech is already reliable in the 6-channel case, rendering speaker reinforcement not necessary.

In Table 5, where the ASR system proposed in Chang et al. (2022) is used for recognition, applying speaker reinforcement in system 2b, 4b and 6b respectively outperforms 2a, 4a and 6a, pushing down the WER on the real test set to 2.40%, 1.84% and 1.58%.

#### 8.8. Comparison with using close-talk mixtures for ASR

In system 7 of Table 5 and system 9 of Tables 2 and 3, we provide the ASR results of using close-talk mixtures for decoding. We observe that our proposed 6-channel system obtains ASR results comparable to using close-talk mixtures for decoding. This further indicates the effectiveness of SuperM2M and our robust ASR system.

#### 8.9. Whisper vs. ESPnet ASR model under same text norm

In previous tables, the WER results of using the Whisper ASR model and using the ESPnet ASR model cannot be directly compared, due to the use of different text normalization. As a side result, we compare the ESPnet ASR model with the Whisper Large v2 model, by training and evaluating the ESPnet ASR model using the same text normalizer as Whisper. Table 6 reports the results. We observe that the ESPnet ASR model obtains clearly better WERs than Whisper on the CHiME-4 dataset.

**Table 5**

ASR results in official CHiME-4 setup (ASR model: WavLM features + encoder-decoder model (Chang et al., 2022) in ESPnet).

System	Cross reference	Approach	Frontend	Joint training	Input #mics	Spk. reinf. $\gamma$ (dB)	Official CHiME-4 Test Utterances			
							Val. WER (%)↓		Test WER (%)↓	
							SIMU	REAL	SIMU	REAL
0a	–	Mixture Chang et al. (2022)	–	–	1	–	5.93	4.03	8.25	4.47
0b	–	Mixture(reproduced)	–	–	1	–	5.93	4.07	8.29	4.47
1a	–	IRIS Chang et al. (2022)	Conv-TasNet	✗	1	–	5.96	4.37	13.52	12.11
1b	–	IRIS Chang et al. (2022)	Conv-TasNet	✓	1	–	3.16	2.03	6.12	3.92
2a	7c of Table 3	SuperM2M	TFGridNetv2	✗	1	–	3.39	1.84	6.57	3.04
2b	7c of Table 3	SuperM2M	TFGridNetv2	✗	1	10	<b>2.40</b>	<b>1.64</b>	<b>4.54</b>	<b>2.40</b>
3a	–	MultiIRIS Masuyama, Chang, Cornell, Watanabe, and Ono (2023)	Neural WPD	✗	2	–	2.28	2.06	2.30	3.63
3b	–	MultiIRIS Masuyama, Chang, Cornell, et al. (2023)	Neural WPD	✓	2	–	2.04	1.66	2.04	2.65
4a	7c of Table 2	SuperM2M	TFGridNetv2	✗	2	–	1.50	1.40	2.08	1.94
4b	7c of Table 2	SuperM2M	TFGridNetv2	✗	2	10	<b>1.28</b>	<b>1.33</b>	<b>1.88</b>	<b>1.84</b>
5a	–	MultiIRIS Masuyama, Chang, Cornell, et al. (2023)	Neural WPD	✗	6	–	1.19	1.32	1.29	1.85
5b	–	MultiIRIS Masuyama, Chang, Cornell, et al. (2023)	Neural WPD	✓	6	–	1.22	1.33	<b>1.24</b>	1.77
6a	7c of Table 2	SuperM2M	TFGridNetv2	✗	6	–	<b>0.83</b>	1.26	1.37	1.61
6b	7c of Table 2	SuperM2M	TFGridNetv2	✗	6	10	<b>0.83</b>	<b>1.23</b>	1.37	<b>1.58</b>
7	–	Close-talk Mixture	–	–	–	–	–	1.14	–	1.49

Notes: The best scores are highlighted in bold in the 1-, 2- and 6-channel cases separately.

Table 6

ASR results on CHiME-4 obtained by using Whisper Large v2 vs. ESPnet ASR model (Chang et al., 2022) under same Whisper text normalization.

System	Cross reference	Approach	ASR model (same Whisper text norm.)	Input #mics	Official CHiME-4 Test Utterances			
					Val. WER (%)↓		Test WER (%)↓	
					SIMU	REAL	SIMU	REAL
1a	–	Mixture	Whisper Large v2	1	7.43	4.96	10.97	7.69
1b	–	Mixture	ESPnet ASR Model Chang et al. (2022)	1	<b>6.94</b>	<b>4.94</b>	<b>9.60</b>	<b>5.42</b>
2a	7c of Table 3	SuperM2M	Whisper Large v2	1	7.05	4.78	11.34	6.02
2b	7c of Table 3	SuperM2M	ESPnet ASR Model Chang et al. (2022)	1	<b>3.96</b>	<b>2.30</b>	<b>7.84</b>	<b>3.65</b>
3a	7c of Table 2	SuperM2M	Whisper Large v2	6	3.38	3.81	3.77	4.04
3b	7c of Table 2	SuperM2M	ESPnet ASR Model Chang et al. (2022)	6	<b>1.19</b>	<b>1.63</b>	<b>1.45</b>	<b>1.90</b>

## 9. Conclusion

We have proposed to adapt UNSSOR and M2M training for neural speech enhancement, where the models can be trained on real-recorded far-field mixtures in an unsupervised way, and on real-recorded close-talk and far-field mixture pairs in a weakly-supervised way. To improve UNSSOR and M2M training, we have proposed SuperM2M, a novel co-learning algorithm that trains neural speech enhancement models by alternating between supervised training on simulated data and UNSSOR/M2M training on real data. Evaluation results on the challenging CHiME-4 benchmark show the effectiveness of SuperM2M for speech enhancement and robust ASR. Future research will modify and evaluate SuperM2M on conversational speech separation and recognition.

Our study, we think, provides illuminating findings towards improving the generalizability of modern neural speech enhancement models to real-recorded data, since it, for the first time since the introduction of the challenging and representative CHiME-4 benchmark nearly a decade ago, shows that, on the real mixtures of CHiME-4, feeding in the immediate outputs of neural speech enhancement models for ASR decoding can produce remarkable improvement over feeding in unprocessed mixtures and neural beamforming results, breaking out to new highs in ASR performance even though joint frontend-backend training is not employed and even if the ASR backend, which leverages strong self-supervised learning representations, is a very strong one. This success is realized by SuperM2M, which trains enhancement models not only on simulated data but also on real data, and through our accumulative efforts on complex spectral mapping (Tan et al., 2022; Wang et al., 2020; Wang, Wang, & Wang, 2021), loss functions dealing with implicit magnitude-phase compensation (Wang, Wichern, & Le Roux, 2021c), FCP (Wang, Wichern, & Le Roux, 2021a, 2021b), TF-GridNet (Wang et al., 2023), UNSSOR (Wang & Watanabe, 2023), USDnet (Wang, 2024b), and M2M (Wang, 2024a), which have firmly built up the foundation of SuperM2M.

We point out that nearly all the current supervised neural speech enhancement algorithms can be seamlessly integrated with SuperM2M to improve their generalization abilities, by including real-recorded close-talk and far-field mixture pairs, or far-field mixtures alone if close-talk mixtures are unavailable, for M2M training. This indicates that SuperM2M can ride on the development of large-scale supervised neural speech enhancement models trained on simulated data, and has strong potential to grow into a representative algorithm for training speech enhancement models directly on real-recorded data.

In closing, we would like to highlight the learning-based methodology for solving blind deconvolution problems, which has been developed along our line of research on FCP (Wang, Wichern, & Le Roux, 2021a, 2021b), UNSSOR (Wang & Watanabe, 2023), M2M (Wang, 2024a), and SuperM2M. By training DNNs in an un-, weakly- or semi-supervised way to estimate sources, filter estimation becomes differentiable so that the DNNs can be trained to optimize mixture-constraint losses to realize separation. Based on the challenging real-recorded mixtures in CHiME-4 and through the integration with supervised learning, this paper, for the first time, has demonstrated that this methodology is effective for real-recorded data, and is also effective at neural speech enhancement. Considering that blind deconvolution broadly exists in many application domains, we expect the methodology to be also effective in similar applications and generate broader impact beyond speech enhancement.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

We would like to thank Dr. Xuankai Chang at Carnegie Mellon University for constructive discussions on robust ASR.

## Data availability

Data will be made available on request.

## References

- Aralikatti, R., Boeddeker, C., Wichern, G., Subramanian, A. S., & Le Roux, J. (2023). Reverberation as supervision for speech separation. In *ICASSP*.
- Bando, Y., Sekiguchi, K., Masuyama, Y., Nugraha, A. A., Fontaine, M., & Yoshii, K. (2021). Neural full-rank spatial covariance analysis for blind source separation. *IEEE Signal Processing Letters*, 28, 1670–1674.
- Barker, J., Watanabe, S., Vincent, E., & Trmal, J. (2018). The fifth ‘CHiME’ speech separation and recognition challenge: Dataset, task and baselines. In *Interspeech* (pp. 1561–1565).
- Boeddeker, C., Heitkaemper, J., Schmalenstroer, J., Drude, L., Heymann, J., & Haeb-Umbach, R. (2018). Front-end processing for the chime-5 dinner party scenario. In *CHiME* (pp. 35–40).
- Carletta, J., et al. (2006). The AMI meeting corpus: A pre-announcement. In *Machine learning for multimodal interact.: vol. 3869*, (pp. 28–39).
- Chang, X., Maekaku, T., Fujita, Y., & Watanabe, S. (2022). End-to-end integration of speech recognition, speech enhancement, and self-supervised learning representation. In *Interspeech* (pp. 3819–3823).
- Chang, X., Zhang, W., Qian, Y., Le Roux, J., & Watanabe, S. (2019). MIMO-SPEECH: End-to-end multi-channel multi-speaker speech recognition. In *ASRU* (pp. 237–244).
- Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., et al. (2022). WavLM: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6), 1505–1518.
- Chen, J., Wang, Y., Yoho, S. E., Wang, D., & Healy, E. W. (2016). Large-scale training to increase speech intelligibility for hearing-impaired listeners in novel noises. *Journal of the Acoustical Society of America*, 139(5), 2604–2612.
- Chen, Z., Yoshioka, T., Lu, L., Zhou, T., Meng, Z., Luo, Y., et al. (2020). Continuous speech separation: Dataset and analysis. In *ICASSP* (pp. 7284–7288).
- Chetupalli, S. R., & Habets, E. A. (2023). Speaker counting and separation from single-channel noisy mixtures. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 1681–1692.
- Cornell, S., Wiesner, M., Watanabe, S., Raj, D., Chang, X., Garcia, P., et al. (2023). The chime-7 DASR challenge: Distant meeting transcription with multiple devices in diverse scenarios. In *CHiME*.
- Cox, T. J., Barker, J., Bailey, W., Graetzer, S., Akeroyd, M. A., Culling, J. F., et al. (2023). Overview of the 2023 ICASSP SP clarity challenge: Speech enhancement for hearing aids. In *ICASSP*.
- Du, Y.-Q., Zhang, J., Fang, X., Wu, M.-H., & Yang, Z.-W. (2023). A semi-supervised complementary joint training approach for low-resource speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 3908–3921.
- Ephrat, A., Mosseri, I., Lang, O., Dekel, T., Wilson, K., Hassidim, A., et al. (2018). Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *ACM Transactions on Graphics*, 37(4).
- Fujimura, T., Koizumi, Y., Yatabe, K., & Miyazaki, R. (2021). Noisy-target training: A training strategy for DNN-based speech enhancement without clean speech. In *EUSIPCO* (pp. 436–440).

- Gannot, S., Vincent, E., Markovich-Golan, S., & Ozerov, A. (2017). A consolidated perspective on multi-microphone speech enhancement and source separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25, 692–730.
- H. Taal, C., C. Hendriks, R., Heusdens, R., & Jensen, J. (2011). An algorithm for intelligibility prediction of time-frequency weighted noisy speech. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(7), 2125–2136.
- Haeb-Umbach, R., Heymann, J., Drude, L., Watanabe, S., Delcroix, M., & Nakatani, T. (2020). Far-field automatic speech recognition. *Proceedings of the IEEE*.
- Haeb-Umbach, R., Watanabe, S., Nakatani, T., Bacchiani, M., Hoffmeister, B., Seltzer, M. L., et al. (2019). Speech processing for digital home assistants: Combining signal processing with deep-learning techniques. *IEEE Signal Processing Magazine*, 36(6), 111–124.
- Han, C., Wilson, K., Wisdom, S., & R. Hershey, J. (2024). Unsupervised multi-channel separation and adaptation. In *ICASSP* (pp. 721–725).
- Hao, X., Xu, C., & Xie, L. (2023). Neural speech enhancement with unsupervised pre-training and mixture training. *Neural Networks*, 158, 216–227.
- Heymann, J., Drude, L., Boeddeker, C., Hanebrink, P., & Haeb-Umbach, R. (2017). BEAMNET: End-to-end training of a beamformer-supported multi-channel ASR system. In *ICASSP* (pp. 5325–5329).
- Heymann, J., Drude, L., Chinaev, A., & Haeb-Umbach, R. (2015). BLSTM supported GEV beamformer front-end for the 3rd CHiME challenge. In *ASRU* (pp. 444–451).
- Kavalerov, I., Wisdom, S., Erdogan, H., Patton, B., Wilson, K., Le Roux, J., et al. (2019). Universal sound separation. In *WASPAA* (pp. 175–179).
- Le Roux, J., Wisdom, S., Erdogan, H., & Hershey, J. R. (2019). SDR - half-baked or well done? In *ICASSP* (pp. 626–630).
- Leglaive, S., Borne, L., Tzinis, E., Sadeghi, M., Fraticelli, M., Wisdom, S., et al. (2023). The CHiME-7 UDASE task: Unsupervised domain adaptation for conversational speech enhancement. In *CHiME*.
- Levin, A., Weiss, Y., Durand, F., & T. Freeman, W. (2011). Understanding blind deconvolution algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12), 2354–2367.
- Lu, Y.-J., Cornell, S., Chang, X., Zhang, W., Li, C., Ni, Z., et al. (2022). Towards low-distortion multi-channel speech enhancement: The ESPNet-SE submission to the L3DAS22 challenge. In *ICASSP* (pp. 9201–9205).
- Luo, Y., Chen, Z., & Yoshioka, T. (2020). Dual-path RNN: Efficient long sequence modeling for time-domain single-channel speech separation. In *ICASSP* (pp. 46–50).
- Luo, Y., & Mesgarani, N. (2019). Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27, 1256–1266.
- Masuyama, Y., Chang, X., Cornell, S., Watanabe, S., & Ono, N. (2023). End-to-end integration of speech recognition, dereverberation, beamforming, and self-supervised learning representation. In *SLT* (pp. 260–265).
- Masuyama, Y., Chang, X., Zhang, W., Cornell, S., Wang, Z.-Q., Ono, N., et al. (2023). Exploring the integration of speech separation and recognition with self-supervised learning representation. In *WASPAA* (pp. 1–5).
- Nachmani, E., Adi, Y., & Wolf, L. (2020). Voice separation with an unknown number of multiple speakers. In *ICML* (pp. 7121–7132).
- Narayanan, A., & Wang, D. (2015). Improving robustness of deep neural network acoustic models via speech separation and joint adaptive training. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(1), 92–101.
- Pandey, A., & Wang, D. (2019). A new framework for CNN-based speech enhancement in the time domain. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 27, 1179–1188.
- Pandey, A., & Wang, D. (2020). On cross-corpus generalization of deep learning based speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28, 2489–2499.
- Pishdadian, F., Wichern, G., & Le Roux, J. (2020). Finding strength in weakness: Learning to separate sounds with weak supervision. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28, 2386–2399.
- Pons, J., Liu, X., Pascual, S., & Serra, J. (2024). GASS: Generalizing audio source separation with large-scale data. In *ICASSP*.
- Quan, C., & Li, X. (2024). SpatialNet: Extensively learning spatial information for multichannel joint speech separation, denoising and dereverberation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32, 1310–1323.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., & Sutskever, I. (2023). Robust speech recognition via large-scale weak supervision. *Proceedings Machine Learning Research*, 202, 28492–28518.
- Rix, A., Beerends, J., Hollier, M., & Hekstra, A. (2001). Perceptual evaluation of speech quality (PESQ)-A new method for speech quality assessment of telephone networks and codecs. In *ICASSP: vol. 2*, (pp. 749–752).
- Saijo, K., Zhang, W., Wang, Z.-Q., Watanabe, S., Kobayashi, T., et al. (2023). A single speech enhancement model unifying dereverberation, denoising, speaker counting, separation, and extraction. In *ASRU*.
- Sawada, H., Araki, S., & Makino, S. (2007). A two-stage frequency-domain blind source separation method for underdetermined convolutive mixtures. In *WASPAA* (pp. 139–142).
- Sawada, H., Ono, N., Kameoka, H., Kitamura, D., & Saruwatari, H. (2019). A review of blind source separation methods: Two converging routes to ILRMA originating from ICA and NMF. *APSIPA Transactions on Signal and Information Processing*, 8, 1–14.
- Sivaraman, A., Wisdom, S., Erdogan, H., & R. Hershey, J. (2022). Adapting speech separation to real-world meetings using mixture invariant training. In *ICASSP* (pp. 686–690).
- Stoller, D., Ewert, S., & Dixon, S. (2018). Adversarial semi-supervised audio source separation applied to singing voice extraction. In *ICASSP* (pp. 2391–2395).
- Talmon, R., Cohen, I., & Gannot, S. (2009). Relative transfer function identification using convolutive transfer function approximation. *IEEE Transactions on Audio, Speech, and Language Processing*, 17(4), 546–555.
- Tan, K., Wang, Z.-Q., & Wang, D. (2022). Neural spectrospatial filtering. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30, 605–621.
- Tesch, K., & Gerkmann, T. (2021). Nonlinear spatial filtering in multichannel speech enhancement. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 1795–1805.
- Tzinis, E., Adi, Y., Ithapu, V. K., Xu, B., Smaragdis, P., & Kumar, A. (2022). RemixIT: Continual self-training of speech enhancement models via bootstrapped remixing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6), 1329–1341.
- Tzinis, E., Wisdom, S., Remez, T., & R. Hershey, J. (2022). AudioScopeV2: Audio-visual attention architectures for calibrated open-domain on-screen sound separation. In *ECCV* (pp. 368–385).
- Vincent, E., Gribonval, R., & Févotte, C. (2006). Performance measurement in blind audio source separation. *IEEE Transactions on Audio, Speech, and Language Processing*, 14(4), 1462–1469.
- Vincent, E., Watanabe, S., Nugraha, A. A., Barker, J., & Marxer, R. (2017). An analysis of environment, microphone and data simulation mismatches in robust speech recognition. *Computer Speech and Language*, 46, 535–557.
- Vu, D. H. T., & Haeb-Umbach, R. (2010). Blind speech separation employing directional statistics in an expectation maximization framework. In *ICASSP* (pp. 241–244).
- Wang, Z.-Q. (2024a). Mixture to mixture: Leveraging close-talk mixtures as weak-supervision for speech separation. *IEEE Signal Processing Letters*, 31, 1715–1719.
- Wang, Z.-Q. (2024b). USDnet: Unsupervised speech dereverberation via neural forward filtering. *arXiv preprint arXiv:2402.00820*.
- Wang, D., & Chen, J. (2018). Supervised speech separation based on deep learning: An overview. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(10), 1702–1726.
- Wang, Z.-Q., Cornell, S., Choi, S., Lee, Y., et al. (2023). TF-GridNet: Integrating full- and sub-band modeling for speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 3221–3236.
- Wang, Z.-Q., & Wang, D. (2016). A joint training framework for robust automatic speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(4), 796–806.
- Wang, Z.-Q., & Wang, D. (2020). Deep learning based target cancellation for speech dereverberation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28, 941–950.
- Wang, Z.-Q., Wang, P., & Wang, D. (2020). Complex spectral mapping for single- and multi-channel speech enhancement and robust ASR. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28, 1778–1787.
- Wang, Z.-Q., Wang, P., & Wang, D. (2021). Multi-microphone complex spectral mapping for utterance-wise and continuous speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 2001–2014.
- Wang, Z.-Q., & Watanabe, S. (2023). UNSSOR: Unsupervised neural speech separation by leveraging over-determined training mixtures. In *NeurIPS* (pp. 34021–34042).
- Wang, Z.-Q., Wichern, G., & Le Roux, J. (2021a). Convolutional prediction for monaural speech dereverberation and noisy-reverberant speaker separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 3476–3490.
- Wang, Z.-Q., Wichern, G., & Le Roux, J. (2021b). Convolutional prediction for reverberant speech separation. In *WASPAA* (pp. 56–60).
- Wang, Z.-Q., Wichern, G., & Le Roux, J. (2021c). On the compensation between magnitude and phase in speech separation. *IEEE Signal Processing Letters*, 28, 2018–2022.
- Watanabe, S., Mandel, M., Barker, J., Vincent, E., et al. (2020). Chime-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings. *arXiv preprint arXiv:2004.09249*.
- Wisdom, S., R. Hershey, J., Wilson, K., Thorpe, J., Chinen, M., Patton, B., et al. (2019). Differentiable consistency constraints for improved deep speech enhancement. In *ICASSP* (pp. 900–904).
- Wisdom, S., Tzinis, E., Erdogan, H., Weiss, R. J., Wilson, K., & R. Hershey, J. (2020). Unsupervised sound separation using mixture invariant training. In *NeurIPS*.
- Wu, S., Wang, C., Chen, H., Dai, Y., Zhang, C., Wang, R., et al. (2024). The multimodal information based speech processing (MISP) 2023 challenge: Audio-visual target speaker extraction. In *ICASSP*.
- Xu, C., Rao, W., Chng, E. S., & Li, H. (2020). SpEx: Multi-scale time domain speaker extraction network. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28, 1370–1384.
- Yang, Y., Pandey, A., & Wang, D. (2024). Towards decoupling frontend enhancement and backend recognition in monaural robust ASR. *arXiv preprint arXiv:2403.06387*.
- Yu, F., et al. (2022). M2MeT: The ICASSP 2022 multi-channel multi-party meeting transcription challenge. In *ICASSP* (pp. 6167–6171).
- Zeghidour, N., & Grangier, D. (2021). Wavesplit: End-to-end speech separation by speaker clustering. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 2840–2849.

- Zhang, W., Jung, J.-w., Watanabe, S., & Qian, Y. (2024). Improving design of input condition invariant speech enhancement. In *ICASSP*.
- Zhang, W., Saijo, K., Wang, Z.-Q., Watanabe, S., et al. (2023). Toward universal speech enhancement for diverse input conditions. In *ASRU*.
- Zhang, W., Shi, J., Li, C., Watanabe, S., & Qian, Y. (2021). Closing the gap between time-domain multi-channel speech enhancement on real and simulation conditions. In *WASPAA* (pp. 146–150).
- Zhang, X., Wang, Z.-Q., & Wang, D. (2017). A speech enhancement algorithm by iterating single- and multi-microphone processing and its application to robust ASR. In *ICASSP* (pp. 276–280).
- Zhang, Z., Xu, Y., Yu, M., Zhang, S. X., Chen, L., S. Williamson, D., et al. (2021). Multi-channel multi-frame ADL-MVDR for target speech separation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29, 3526–3540.
- Zhang, N., Yan, J., & Zhou, Y. (2018). Weakly supervised audio source separation via spectrum energy preserved Wasserstein learning. In *IJCAI* (pp. 4574–4580).
- Zhang, J., Zorilă, C., Doddipatla, R., & Barker, J. (2022). On monaural speech enhancement for automatic recognition of real noisy speech using mixture invariant training. In *Interspeech* (pp. 1056–1060).
- Zheng, C., Zhang, H., Liu, W., Luo, X., Li, A., et al. (2023). Sixty years of frequency-domain monaural speech enhancement: From traditional to deep learning methods. *Trends in Hearing*, 27.
- Zhu, Q.-S., Zhang, J., Zhang, Z.-Q., & Dai, L.-R. (2023). A joint speech enhancement and self-supervised representation learning framework for noise-robust speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31, 1927–1939.
- Zhu, Q.-S., Zhang, J., Zhang, Z.-Q., Wu, M.-H., Fang, X., & Dai, L.-R. (2022). A noise-robust self-supervised pre-training model based speech representation learning for automatic speech recognition. In *ICASSP* (pp. 3174–3178).
- Zmolikova, K., Delcroix, M., Ochiai, T., Kinoshita, K., Cernocky, J., & Yu, D. (2023). Neural target speech extraction: An overview. *IEEE Signal Processing Magazine*, 40(3), 8–29.
- Zorilă, C., & Doddipatla, R. (2022). Speaker reinforcement using target source extraction for robust automatic speech recognition. In *ICASSP* (pp. 6297–6301).