

## FLEXIO: FLEXIBLE SINGLE- AND MULTI-CHANNEL SPEECH SEPARATION AND ENHANCEMENT

Masuyama, Yoshiki; Saijo, Kohei; Paissan, Francesco; Han, Jiangyu; Delcroix, Marc; Aihara, Ryo;  
Germain, François G; Wichern, Gordon; Le Roux, Jonathan

TR2026-034 March 19, 2026

### Abstract

Speech separation and enhancement (SSE) has advanced remarkably and achieved promising results in controlled settings, such as a fixed number of speakers and a fixed array configuration. Towards a universal SSE system, single-channel systems have been extended to deal with a variable number of speakers (i.e., outputs). Meanwhile, multi-channel systems accommodating various array configurations (i.e., inputs) have been developed. However, these attempts have been pursued separately. In this paper, we propose a flexible input and output SSE system, named FlexIO. It performs conditional separation using prompt vectors, one per speaker as a condition, allowing separation of an arbitrary number of speakers. Multi-channel mixtures are processed together with the prompt vectors via an array-agnostic channel communication mechanism. Our experiments demonstrate that FlexIO successfully covers diverse conditions with one to five microphones and one to three speakers. We also confirm the robustness of FlexIO on CHiME-4 real data.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)  
2026*



# FLEXIO: FLEXIBLE SINGLE- AND MULTI-CHANNEL SPEECH SEPARATION AND ENHANCEMENT

Yoshiki Masuyama<sup>1</sup>, Kohei Saijo<sup>2</sup>, Francesco Paissan<sup>1,3</sup>, Jiangyu Han<sup>4</sup>, Marc Delcroix<sup>5</sup>,  
Ryo Aihara<sup>1</sup>, François G. Germain<sup>1</sup>, Gordon Wichern<sup>1</sup>, Jonathan Le Roux<sup>1</sup>

<sup>1</sup>Mitsubishi Electric Research Laboratories (MERL), Cambridge, USA

<sup>2</sup>Waseda University, Tokyo, Japan <sup>3</sup>University of Trento, Trento, Italy

<sup>4</sup>Brno University of Technology, Brno, Czechia <sup>5</sup>NTT, Inc., Kyoto, Japan

## ABSTRACT

Speech separation and enhancement (SSE) has advanced remarkably and achieved promising results in controlled settings, such as a fixed number of speakers and a fixed array configuration. Towards a universal SSE system, single-channel systems have been extended to deal with a variable number of speakers (i.e., outputs). Meanwhile, multi-channel systems accommodating various array configurations (i.e., inputs) have been developed. However, these attempts have been pursued separately. In this paper, we propose a flexible input and output SSE system, named FlexIO. It performs conditional separation using prompt vectors, one per speaker as a condition, allowing separation of an arbitrary number of speakers. Multi-channel mixtures are processed together with the prompt vectors via an array-agnostic channel communication mechanism. Our experiments demonstrate that FlexIO successfully covers diverse conditions with one to five microphones and one to three speakers. We also confirm the robustness of FlexIO on CHiME-4 real data.

**Index Terms**— Speech enhancement, speech separation, universality, array-agnostic processing, prompting

## 1. INTRODUCTION

Speech separation and enhancement (SSE) aims to isolate individual speech signals from mixtures contaminated by background noise and reverberation [1, 2]. SSE is crucial not only for improving speech quality but also as a front-end for speech recognition [3–5]. With the advent of deep learning, SSE has shown remarkable results even under single-channel conditions [6, 7]. Dual-path modeling [8] in the short-time Fourier transform (STFT) domain, in particular, has yielded promising performance [9–11]. This approach has also been used in multi-channel SSE systems by concatenating the channel-wise STFT coefficients as input [11, 12]. These studies, however, mainly focused on matched conditions, i.e., the numbers of speakers and microphones are the same during training and evaluation. A single SSE model capable of flexibly handling various numbers of microphones (input) and speakers (output) remains underexplored.

Several attempts have been made in single-channel SSE to increase the flexibility against the number of speakers in the mixture. Recursive separation schemes extract each speaker one by one from the mixture [13, 14], and attractor-based methods internally estimate the number of speakers before separation [15–17]. More recently, task-aware unified source separation (TUSS) introduced prompt-conditional separation that can deal with an arbitrary

**Table 1.** Comparison of FlexIO and flexible SSE systems, where  $M$  and  $N$  are the number of microphones and speakers, respectively. Controllability indicates whether users can specify the number of output streams during inference.

	Flexible $M$	Flexible $N$	Controllability
TPARN [25], USES [26]	✓	✗	✗
SepEDA [15], TUSS [18]	✗	✓	✓
VarArray [27]	✓	✓	✗
DNN-IVA [28]	✓ <sup>†</sup>	✓	✓
FlexIO (Proposed)	✓	✓	✓

<sup>†</sup>DNN-IVA is not applicable to single-channel scenarios.

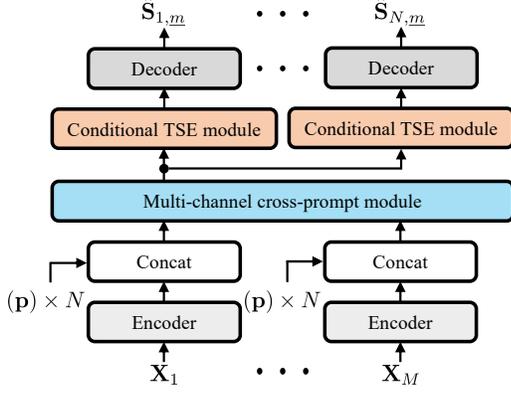
number of speakers depending on the number of given prompt vectors [18]. Prompt-conditional separation allows users to explicitly specify the number of output streams by changing the number of prompt vectors, and this controllability is beneficial if the number of speakers is known or estimated in advance.

In multi-channel SSE, pioneering work was based on array signal processing [19, 20] and has been combined with single-channel neural networks, e.g., in mask-based beamforming [21, 22]. Mask-based beamforming has been widely used due to its flexibility with regard to the number of microphones and the array geometry. Recent attempts solely adopt neural networks since the traditional array signal processing can be a performance bottleneck [11, 12, 23]. A naïve implementation is to concatenate the channel-wise STFT coefficients as the network input [11, 12], but this cannot accommodate mixtures with different numbers of channels. To overcome this limitation, array-agnostic channel communication mechanisms have been developed [24, 25], which deal with an arbitrary number of microphones and exchange the information across channels. Especially, the transform-average-concatenate (TAC) mechanism [24] has been successfully applied to universal speech enhancement [26].

Prior works have typically tackled the flexibility with regard to either the number of microphones  $M$  or the number of speakers  $N$  in isolation, as summarized in Table 1. Towards a more flexible SSE system, VarArray [27] exploits TAC and supports a variable number of speakers in a continuous speech separation scheme. However, the number of concurrent output streams is fixed, and the model does not exploit prior knowledge about the expected number of speakers. Independent vector analysis with neural source modeling (DNN-IVA) [28] offers both flexibility and controllability, but its performance is limited in underdetermined scenarios, i.e.,  $N > M$ , due to its dependence on array signal processing.

In this paper, we propose FlexIO, a flexible and versatile SSE

This work was done while F. Paissan was an intern at MERL. Jiangyu Han was partly supported by BUT IGA project no. FIT-S-23-8278.



**Fig. 1.** FlexIO takes  $N$  prompt vectors  $\mathbf{p}$  and an  $M$ -channel mixture  $(\mathbf{X}_1, \dots, \mathbf{X}_M)$  as input, and isolates each speaker at a reference channel  $\underline{m}$ . Modules of the same color share parameters.

system that can handle an arbitrary combination of  $M$  and  $N$ . FlexIO extends the single-channel TUSS framework [18] to support multi-channel input in an array-agnostic manner as depicted in Fig. 1. The multi-channel cross-prompt module processes the mixture representation and the given prompt vectors using time-frequency dual-path modeling with channel communication mechanisms. Then, we perform target speaker extraction (TSE) based on the processed mixture representation and prompt vectors at the reference channel. Our experiments on comprehensive scenarios demonstrate that FlexIO successfully copes with various input and output conditions with strong performance.

## 2. PRELIMINARIES

### 2.1. Problem settings

Let us consider the STFT of an  $M$ -channel reverberant noisy mixture of  $N$  speakers,  $(\mathbf{X}_1, \dots, \mathbf{X}_M)$ , where  $\mathbf{X}_m \in \mathbb{C}^{T \times F}$ , and  $T$  and  $F$  denote the number of time frames and frequency bins, respectively. Our goal is to estimate the anechoic source image of each speaker regardless of the number of microphones and speakers:

$$\{\hat{\mathbf{S}}_{1,\underline{m}}, \dots, \hat{\mathbf{S}}_{N,\underline{m}}\} \leftarrow \text{SSE}(\mathbf{X}_1, \dots, \mathbf{X}_M), \quad (1)$$

where  $\mathbf{S}_{n,\underline{m}} \in \mathbb{C}^{T \times F}$  denotes the desired source image of the  $n$ -th speaker at a reference channel  $\underline{m}$ . The number of speakers  $N$  can be different for each mixture, but we assume it is estimated in advance.

### 2.2. Single-channel TUSS for SSE

We review the application to SSE of TUSS [18], a prompt-conditional separation system on top of which we build our proposed FlexIO. In this section, we omit the microphone index  $m$  as TUSS was developed as a single-channel system. First, we encode the real and imaginary parts of the mixture  $\mathbf{X}$  into an initial mixture representation  $\mathbf{Z} \in \mathbb{R}^{D \times T \times F}$  by applying 2D convolution and global layer normalization, where  $D$  is the feature dimension. Then, to extract  $N$  speakers, we concatenate  $N$  repeated prompt vectors  $\mathbf{p} \in \mathbb{R}^D$  to the mixture representation for subsequent prompt-aware processing, leading to the representation  $\mathbf{Z}' = [\mathbf{P}_1, \dots, \mathbf{P}_N, \mathbf{Z}] \in \mathbb{R}^{D \times T' \times F}$ , where  $\mathbf{P}_n$  is an  $F$  repeat of  $\mathbf{p}$  regardless of  $n$ , and  $T' = T + N$ . We share the same initial prompt vectors across speakers to handle an arbitrary number of speakers, but they will be updated to capture the characteristics of different speakers in the following module.

A key component of TUSS is the cross-prompt module that updates the mixture representation and prompt vectors simultaneously

to capture the relation between the mixture and each speaker. This module applies multiple TF-locoformer blocks [10], time-frequency dual-path modeling blocks with multi-head self-attention (MHSA), to the concatenated representation  $\mathbf{Z}'$ . In the temporal modeling,  $\mathbf{Z}'$  is viewed as  $F$  separate sequences of  $D$ -dimensional features with length  $T'$ . Then, MHSA and 1D convolution with a SwiGLU activation [10] are applied along the time axis:

$$\mathbf{Z}' \leftarrow \mathbf{Z}' + \text{ConvSwiGLU}(\mathbf{Z}'), \quad (2)$$

$$\mathbf{Z}' \leftarrow \mathbf{Z}' + \text{MHSA}(\text{Norm}(\mathbf{Z}')), \quad (3)$$

$$\mathbf{Z}' \leftarrow \mathbf{Z}' + \text{ConvSwiGLU}(\mathbf{Z}'), \quad (4)$$

where Norm is the root mean square group normalization [10]. Thanks to the rotary positional encoding [29] in MHSA, the identical prompt vectors are updated to different values. The frequency modeling performs similar operations along the frequency axis at each time frame.

The subsequent conditional TSE module isolates speaker-wise representations. The output  $\tilde{\mathbf{Z}}'$  of the cross-prompt module is split into components corresponding to each prompt and the mixture:  $(\tilde{\mathbf{P}}_1, \dots, \tilde{\mathbf{P}}_N, \tilde{\mathbf{Z}}) \leftarrow \tilde{\mathbf{Z}}'$ , where  $\tilde{\mathbf{P}}_n$  is expected to capture the characteristics of each speaker. The speaker-wise representation is then obtained via a Hadamard product following the conventional TSE framework [30]:  $\tilde{\mathbf{Z}}_n = \tilde{\mathbf{P}}_n \odot \tilde{\mathbf{Z}} \in \mathbb{R}^{D \times T \times F}$ , where  $\tilde{\mathbf{P}}_n$  is broadcasted along the time axis. We then separately refine  $\tilde{\mathbf{Z}}_n$  for each speaker using additional TF-locoformer blocks, where the model parameters are shared across speakers. The output of the conditional TSE module is fed to a 2D deconvolution decoder and multiplied by the STFT of the mixture  $\mathbf{X}$  as complex time-frequency masks<sup>1</sup> [31]. TUSS can deal with a variable number of speakers, including a single speaker, by adjusting the number of prompt vectors, although its input is limited to a single channel.

## 3. FLEXIO

### 3.1. Overview of FlexIO

To realize flexible multi-channel SSE, FlexIO integrates a channel communication mechanism into the cross-prompt module of TUSS, and its entire architecture is designed to generalize to diverse array configurations as depicted in Fig. 1. First, each channel component of the  $M$ -channel mixture  $(\mathbf{X}_1, \dots, \mathbf{X}_M)$  is separately encoded. The multi-channel cross-prompt module handles the  $M$ -channel features  $(\mathbf{Z}'_1, \dots, \mathbf{Z}'_M)$  separately, except for the additional channel communication mechanism: TAC [24], cross-channel attention [25], or co-attention [32]. Then, we take the processed representations at the reference channel and feed them into the conditional TSE module:  $\tilde{\mathbf{Z}}_{n,\underline{m}} = \tilde{\mathbf{P}}_{n,\underline{m}} \odot \tilde{\mathbf{Z}}_{\underline{m}}$ . We expect this to encourage FlexIO to produce estimates aligned with the ground truth at the reference channel [26, 27]. Unlike existing systems based on array signal processing [28], FlexIO can handle even single-channel mixtures due to its careful design. In addition, inheriting the prompt-conditional separation of TUSS, FlexIO can deal with a variable number of speakers by adjusting the number of prompt vectors. With one prompt vector, FlexIO focuses on suppressing background noise and reverberation, thus covering both enhancement and separation in a unified manner.

### 3.2. TAC and cross-channel attention

We require a channel communication mechanism to exchange information across channels while keeping the entire architecture array-

<sup>1</sup>We use 2D deconvolution to predict time-frequency masks instead of band-split processing used in [18], since we focus only on 16 kHz speech.

agnostic. TAC is a well-developed channel communication mechanism [24], and we adopt it after each TF-locoformer block. In detail, we separately encode each TF-bin to a hidden dimension  $E$ :

$$\mathbf{W}_m \leftarrow \text{FC}_{\text{in}}(\mathbf{Z}'_m) \in \mathbb{R}^{E \times T' \times F}, \quad (5)$$

where FC denotes a fully-connected layer with PReLU activation. The projected representation is averaged over channels and fed into another fully-connected layer:

$$\underline{\mathbf{W}} \leftarrow \text{FC}_{\text{avg}}\left(\frac{1}{M} \sum_{m=1}^M \mathbf{W}_m\right) \in \mathbb{R}^{E \times T' \times F}. \quad (6)$$

The global representation  $\underline{\mathbf{W}}$  is concatenated with the channel-wise representation  $\mathbf{W}_m$  and projected back to the original dimension with a skip connection:

$$\mathbf{Z}'_m \leftarrow \mathbf{Z}'_m + \text{Norm}(\text{FC}_{\text{cat}}([\mathbf{W}_m; \underline{\mathbf{W}}])). \quad (7)$$

This mechanism can handle an arbitrary number of microphones due to the nature of average pooling. In addition, it falls back to standard fully-connected layers in single-channel scenarios.

Another channel communication mechanism is cross-channel attention [25, 33]. It applies MHSA across channels without positional encoding, which enables generalization to a different number of microphones. The cross-channel attention realizes a dynamic weighted average, while TAC aggregates the representation with the same weight for each channel regardless of the given mixture.

### 3.3. Channel communication via co-attention

We also explore the co-attention mechanism, originally proposed for diarization [32]. In contrast to TAC and cross-channel attention, co-attention does not introduce additional layers but tweaks existing channel-wise MHSA. Specifically, the original channel-wise MHSA for temporal modeling in (3) can be realized by

$$\mathbf{A}_{m,f}^{(h)} \leftarrow \text{softmax}\left(\frac{\mathbf{Q}_{m,f}^{(h)} \mathbf{K}_{m,f}^{(h)\top}}{\sqrt{D_{\text{att}}}}\right) \in \mathbb{R}^{T' \times T'}, \quad (8)$$

$$\mathbf{U}_{m,f} \leftarrow \text{FC}_{\text{att}}\left([\mathbf{A}_{m,f}^{(1)} \mathbf{V}_{m,f}^{(1)}; \dots; \mathbf{A}_{m,f}^{(H)} \mathbf{V}_{m,f}^{(H)}]\right) \in \mathbb{R}^{D \times T'}, \quad (9)$$

where  $H$  is the number of attention heads, and  $\mathbf{Q}_{m,f}^{(h)}$ ,  $\mathbf{K}_{m,f}^{(h)}$  and  $\mathbf{V}_{m,f}^{(h)}$  respectively are the query, key, and value for the  $h$ -th head at the  $f$ -th frequency. They are calculated from  $\mathbf{Z}'_m$  after the normalization using fully-connected layers. In (9),  $\text{FC}_{\text{att}}$  aggregates the outputs of multiple heads, and the output  $\mathbf{U}_{m,f}$  is stacked over all frequencies. The co-attention mechanism replaces the channel-wise weight in (8) with the following channel-invariant weight [32]:

$$\mathbf{A}_f^{(h)} \leftarrow \text{softmax}\left(\frac{\sum_{m=1}^M \mathbf{Q}_{m,f}^{(h)} \mathbf{K}_{m,f}^{(h)\top}}{\sqrt{D_{\text{att}} M}}\right) \in \mathbb{R}^{T' \times T'}, \quad (10)$$

and  $\mathbf{A}_f^{(h)}$  is used at all channels. The additional factor  $\sqrt{M}$  arises from interpreting the numerator of (10) as the concatenation of  $M$  keys and queries. While co-attention does not explicitly exchange the representation across channels, it can improve multi-channel SSE performance as shown in the following experiments.

## 4. EXPERIMENTS

### 4.1. Datasets

To validate the flexibility of our proposed FlexIO, we combined several SSE datasets, as summarized in Table 2. The CHiME-4 dataset

**Table 2.** Datasets used in our SSE experiments. ‘‘A’’ and ‘‘R’’ in the condition denote anechoic and reverberant settings. The 3- and 5-channel data are excluded from the training and validation sets.

	#Ch $M$	#Spks $N$	Acoustic condition
CHiME-4 [34]	{1, 2, 4}, {3, 5}	1	Noisy A
WSJ0-mix [6]	1	{2, 3}	Clean
WHAM! [35]	1	{1, 2}	Noisy A
WHAMR! [36]	{1, 2}	{1, 2}	Noisy A/R
WSJ1-CHiME [28]	{2, 4}, 3	{2, 3}	Noisy R

provides not only simulated but also real 6-channel noisy recordings of a single speaker [34]. WSJ0-mix is a standard speech separation dataset [6], and we further incorporated its noisy (WHAM! [35]) and noisy reverberant (WHAMR! [36]) extensions to cover diverse acoustic conditions. In addition, our training set includes WSJ1-CHiME developed in [28], which covers various numbers of microphones and speakers<sup>2</sup>. Finally, our training dataset covers 1 to 3 speakers with 1 to 4 channels, where the sampling rate was set to 16 kHz for all recordings. The 3- and 5-channel recordings were excluded from the training and validation sets to assess the generalization capability under unseen conditions.

### 4.2. Model and training details

The multi-channel cross-prompt module and the conditional TSE module consisted of 2 and 4 TF-locoformer blocks [10]. In our medium models, the configuration of each TF-locoformer block mainly followed TUSS with medium size settings [18], i.e.,  $D = 64$ ,  $H = 4$ , and the kernel size and stride of the 1D convolution were 4 and 1, respectively, except for the temporal convolution in the multi-channel cross-prompt module. In the multi-channel cross-prompt module, we omitted ConvSwiGLU before MHSA in (2) to improve the efficiency of multi-channel processing following [37]. For large models, we set  $D$  to 96 and did not omit ConvSwiGLU in (2). Regarding the channel communication mechanism, the hidden dimension  $E$  for TAC was 128, and the cross-channel attention leveraged MHSA with 4 heads and a dimension 16 for each head.

During training, we first sampled a pair of  $N$  and  $M$  for each batch. Then, the corresponding mixtures were randomly selected and truncated at 4 seconds. The medium and large FlexIOs were trained up to 100 and 150 epochs, respectively, with batch sizes of 16 and 8. Each epoch consisted of 2.5k steps. The AdamW optimizer was used with a weight decay factor 0.01. The learning rate was warmed up to 0.001 within 30k steps. The learning rate was halved if the validation loss did not improve over 5 successive epochs, and we terminated the training when the best validation loss was not updated over 10 epochs. The negative signal-to-noise ratio was used as the loss function with permutation invariant training [6, 7]. For evaluation, we used the signal-to-distortion ratio (SDR) [38], signal-to-interference ratio (SIR) [38], wide-band PESQ [39], and STOI [40].

### 4.3. Results

Tables 3 and 4 respectively show enhancement and separation results under diverse conditions, where FlexIO and TUSS share the same models in both tables. ‘‘TUSS [18]’’ corresponds to FlexIO without the channel communication mechanism, and it performs single-channel processing only on the reference channel. In the enhancement experiment, USES [26] is an array-agnostic system trained on multiple datasets. To compute scores not reported in the

<sup>2</sup>[https://github.com/fakufaku/create\\_ws\\_j1\\_2345\\_db](https://github.com/fakufaku/create_ws_j1_2345_db)

**Table 3.** Speech enhancement performance under diverse conditions. The numbers in parentheses denote the numbers of speakers and microphones, i.e.,  $(N-M)$ . For the WHAMR! dataset, “A” and “R” indicate the anechoic and reverberant conditions, respectively. “Comm.” shows the channel communication mechanism, where “1ch” indicates that speech enhancement is performed solely on the reference channel. The best and second best results are highlighted in blue and orange, respectively

	Comm.	#Params ( $10^6$ )	WHAMR! (1-1)			WHAMR! A (1-2)			WHAMR! R (1-2)			CHiME-4 (1-4)			CHiME-4 (1-5)		
			SDR	STOI	PESQ	SDR	STOI	PESQ	SDR	STOI	PESQ	SDR	STOI	PESQ	SDR	STOI	PESQ
USES [26]	TAC	3.05	10.2	85.7	1.65	15.8	96.4	2.55	13.8	96.0	2.51	18.3	96.6	2.46	18.3	97.8	2.95
TUSS [18]	1ch	3.42	13.6	94.2	2.29	13.6	94.1	2.28	12.4	93.7	2.24	17.1	95.7	2.31	17.1	95.7	2.31
FlexIO	TAC (M)	3.59	13.5	94.0	2.25	15.4	95.6	2.55	14.2	95.3	2.50	19.3	97.3	2.54	19.6	97.5	2.60
	ChAtt (M)	3.49	13.6	94.1	2.25	15.6	95.9	2.57	14.8	95.7	2.56	19.5	97.4	2.55	20.2	97.6	2.61
	CoAtt (M)	3.42	13.5	94.0	2.26	15.5	95.7	2.56	14.4	95.4	2.52	20.8	97.8	2.72	21.7	98.1	2.81
	CoAtt (L)	7.35	13.8	94.4	2.31	15.8	96.0	2.61	14.9	95.8	2.58	21.3	98.0	2.86	22.3	98.3	2.96

**Table 4.** Speech separation performance in various settings. FlexIO and TUSS are the same models as those used for speech enhancement in Table 3. While DNN-IVA uses the same model for all conditions, TF-GridNet<sup>†</sup> is separately trained for the 1- and 2-channel conditions.

	Comm.	#Params ( $10^6$ )	WHAMR! R (2-1)			WHAMR! R (2-2)			WSJ1-CHiME (2-2)			WSJ1-CHiME (2-4)			WSJ1-CHiME (3-3)		
			SDR	SIR	PESQ	SDR	SIR	PESQ	SDR	SIR	PESQ	SDR	SIR	PESQ	SDR	SIR	PESQ
DNN-IVA [28]	-	5.13	N/A	N/A	N/A	-	-	-	10.7	24.1	-	-	-	-	7.7	20.1	-
TF-GridNet <sup>†</sup> [5]	-	8.38	9.3	27.0	1.79	11.7	29.6	2.20	-	-	-	-	-	-	-	-	-
TUSS [18]	1ch	3.42	9.5	25.3	1.82	9.5	25.3	1.82	14.9	29.1	2.73	15.1	29.4	2.75	11.4	23.1	2.15
FlexIO	TAC (M)	3.59	8.9	24.0	1.77	11.8	28.5	2.14	18.5	33.9	3.22	19.5	25.0	3.37	15.7	28.4	2.81
	ChAtt (M)	3.49	9.0	23.7	1.75	12.4	29.5	2.22	19.5	34.9	3.34	20.6	36.2	3.49	16.9	30.1	2.97
	CoAtt (M)	3.42	9.1	24.4	1.78	12.1	29.0	2.18	18.9	34.7	3.27	20.6	36.4	3.49	16.7	30.0	2.96
	CoAtt (L)	7.35	9.7	25.5	1.84	12.5	29.6	2.22	19.6	35.0	3.36	21.6	37.3	3.60	17.3	30.4	3.03

original USES paper, we used its pre-trained model<sup>3</sup>. When the number of microphones increased to more than one, FlexIO consistently outperforms its single-channel version (“TUSS”) regardless of the choice of the channel communication mechanism. These results confirm that FlexIO successfully exploits spatial information in an array-agnostic manner. On the single-channel WHAMR! dataset, the models with different channel communication mechanisms result in comparable performance, slightly lagging behind the single-channel specific TUSS. On the 2-channel WHAMR!, the cross-channel attention works marginally better than TAC and co-attention. Meanwhile, FlexIO with co-attention performs best on the CHiME-4 dataset with more channels. Although the model was not trained on 5-channel recordings, it successfully improves the performance when the number of microphones increases from 4 to 5. This result confirms the generalization capability of FlexIO to unseen numbers of microphones. Finally, the large FlexIO with co-attention achieves comparable or slightly better performance compared with the enhancement-specific universal model, USES [26].

Table 4 summarizes the results on speech separation with various numbers of speakers. DNN-IVA [28] uses the same model across different conditions, while TF-GridNet [11] is specific to 2-speaker cases and separately trained for 1- and 2-channel conditions following [5]. FlexIOs with different channel communication mechanisms show the same tendency as in the enhancement experiment. We emphasize that FlexIO works well even on the 3-channel mixtures, even though the models were not trained on 3-channel recordings. On WSJ1-CHiME, our FlexIO significantly outperformed DNN-IVA [28]. Furthermore, the large FlexIO achieves substantially better performance than task-specific TF-GridNet under both 1- and 2-channel conditions. These results confirm that FlexIO successfully covers diverse conditions with arbitrary numbers of microphones and speakers.

<sup>3</sup>[https://huggingface.co/espnet/Wangyou\\_Zhang\\_universal\\_train\\_enh\\_uses\\_refch0\\_2mem\\_raw](https://huggingface.co/espnet/Wangyou_Zhang_universal_train_enh_uses_refch0_2mem_raw)

**Table 5.** DNSMOS and recognition performance on CHiME-4 recordings. FlexIO adopts co-attention for channel communication.

	Simulated				Real			
	DNSMOS		WER (%)		DNSMOS		WER (%)	
	1ch	5ch	1ch	5ch	1ch	5ch	1ch	5ch
Noisy/MVDR	2.08	2.57	5.8	4.0	1.46	1.95	6.7	4.5
USES [26]	3.03	3.22	15.2	4.4	3.07	1.58	7.4	85.9
U2-C [41]	-	-	-	-	-	3.08	-	10.3
FlexIO (M)	3.11	3.17	6.0	3.9	2.91	3.05	6.8	4.5

To validate the generalization capability of FlexIO, we assessed DNSMOS (OVRL) [42] and word error rates (WERs) on the CHiME-4 simulated and real recordings. We used the Whisper Large v2 model for recognition following [26, 41]. As shown in Table 5, USES faces severe performance degradation on the 5-channel real data, which may be due to inconsistent SDR between channels in real recordings [26]. U2-C substantially mitigates this issue by carefully combining TAC and cross-channel attention [41]. Its WER, however, is still worse than mask-based MVDR beamforming used in [43] and even observed signals without processing. On the other hand, our FlexIO simultaneously achieves promising DNSMOS and WER not only on simulated but also on real recordings.

## 5. CONCLUSION

We presented FlexIO, a versatile SSE system that flexibly deals with arbitrary numbers of microphones (inputs) and speakers (outputs). In FlexIO, the multi-channel cross-prompt module takes a sequence of prompt vectors, controlling the number of speakers to be separated. In addition, it exchanges information across channels in an array-agnostic manner. Our experiments show the efficacy of FlexIO on diverse conditions. Its extension to general audio source separation, similar to the original TUSS [18], and to joint source counting are directions of our future work.

## 6. REFERENCES

- [1] S. Gannot, E. Vincent, S. Markovich-Golan, and A. Ozerov, "A consolidated perspective on multi-microphone speech enhancement and source separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 25, pp. 692–730, 2017.
- [2] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [3] T. von Neumann, C. Boeddeker, L. Drude, K. Kinoshita, M. Delcroix, T. Nakatani, and R. Haeb-Umbach, "Multi-talker ASR for an unknown number of sources: Joint training of source counting, separation and ASR," in *Proc. Interspeech*, 2020, pp. 3097–3101.
- [4] D. Raj, P. Denisov, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo, N. Kanda, J. Li, S. Wisdom, and J. R. Hershey, "Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis," in *Proc. SLT*, 2021, pp. 897–904.
- [5] Y. Masuyama, X. Chang, W. Zhang, S. Cornell, Z.-Q. Wang, N. Ono, Y. Qian, and S. Watanabe, "An end-to-end integration of speech separation and recognition with self-supervised learning representation," *Comput. Speech Lang.*, vol. 95, p. 101813, 2026.
- [6] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. ICASSP*, 2016, pp. 31–35.
- [7] M. Kolbæk, D. Yu, Z. H. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 25, no. 10, pp. 1901–1913, Jul. 2017.
- [8] Y. Luo, Z. Chen, and T. Yoshioka, "Dual-path RNN: efficient long sequence modeling for time-domain single-channel speech separation," in *Proc. ICASSP*, 2020, pp. 46–50.
- [9] L. Yang, W. Liu, and W. Wang, "TFPSNet: Time-frequency domain path scanning network for speech separation," in *Proc. ICASSP*, 2022, pp. 6842–6846.
- [10] K. Saijo, G. Wichern, F. G. Germain, Z. Pan, and J. Le Roux, "TF-LoCoformer: Transformer with local modeling by convolution for speech separation and enhancement," in *Proc. IWAENC*, 2024.
- [11] Z.-Q. Wang, S. Cornell, S. Choi, Y. Lee, B.-Y. Kim, and S. Watanabe, "TF-GridNet: Integrating full- and sub-band modeling for speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 31, pp. 3221–3236, 2023.
- [12] V. A. Kalkhorani and D. Wang, "TF-CrossNet: Leveraging global, cross-band, narrow-band, and positional encoding for single- and multi-channel speaker separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 32, pp. 4999–5009, 2024.
- [13] K. Kinoshita, L. Drude, M. Delcroix, and T. Nakatani, "Listening to each speaker one by one with recurrent selective hearing networks," in *Proc. ICASSP*, 2018, pp. 5064–5068.
- [14] N. Takahashi, S. Parthasarathy, N. Goswami, and Y. Mitsufuji, "Recursive speech separation for unknown number of speakers," in *Proc. Interspeech*, 2019, pp. 1348–1352.
- [15] S. R. Chetupalli and E. A. P. Habets, "Speaker counting and separation from single-channel noisy mixtures," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 31, pp. 1681–1692, 2023.
- [16] K. Saijo, W. Zhang, Z.-Q. Wang, S. Watanabe, T. Kobayashi, and T. Ogawa, "A single speech enhancement model unifying dereverberation, denoising, speaker counting, separation, and extraction," in *Proc. ASRU*, 2023.
- [17] S. Dang, L. Li, S. Seki, and H. Kudo, "First analyze then enhance: A task-aware system for speech separation, denoising, and dereverberation," in *Proc. Interspeech*, 2025, pp. 3848–3852.
- [18] K. Saijo, J. Ebberts, F. G. Germain, G. Wichern, and J. Le Roux, "Task-aware unified source separation," in *Proc. ICASSP*, 2025.
- [19] A. Hiroe, "Solution of permutation problem in frequency domain ICA using multivariate probability density functions," in *Proc. ICA*, 2006, pp. 601–608.
- [20] M. Souden, J. Benesty, and S. Affes, "On optimal frequency-domain multichannel linear filtering for noise reduction," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 2, pp. 260–276, 2010.
- [21] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *Proc. ICASSP*, 2016, pp. 196–200.
- [22] H. Erdogan, J. R. Hershey, S. Watanabe, M. I. Mandel, and J. Le Roux, "Improved MVDR beamforming using single-channel mask prediction networks," in *Proc. Interspeech*, 2016, pp. 1981–1985.
- [23] Z.-Q. Wang, P. Wang, and D. Wang, "Complex spectral mapping for single- and multi-channel speech enhancement and robust ASR," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 28, pp. 1778–1787, 2020.
- [24] Y. Luo, Z. Chen, N. Mesgarani, and T. Yoshioka, "End-to-end microphone permutation and number invariant multi-channel speech separation," in *Proc. ICASSP*, 2020, pp. 6394–6398.
- [25] A. Pandey, B. Xu, A. Kumar, J. Donley, P. Calamia, and D. Wang, "TPARN: Triple-path attentive recurrent network for time-domain multichannel speech enhancement," in *Proc. ICASSP*, 2022, pp. 6497–6501.
- [26] W. Zhang, K. Saijo, Z.-Q. Wang, S. Watanabe, and Y. Qian, "Toward universal speech enhancement for diverse input conditions," in *Proc. ASRU*, 2023.
- [27] T. Yoshioka, X. Wang, D. Wang, M. Tang, Z. Zhu, Z. Chen, and N. Kanda, "Vararray: Array-geometry-agnostic continuous speech separation," in *Proc. ICASSP*, 2022, pp. 6027–6031.
- [28] R. Scheibler and M. Togami, "Surrogate source model learning for determined source separation," in *Proc. ICASSP*, 2021, pp. 176–180.
- [29] J. Su, M. Ahmed, Y. Lu, S. Pan, W. Bo, and Y. Liu, "Roformer: Enhanced transformer with rotary position embedding," *Neurocomputing*, vol. 568, p. 127063, 2024.
- [30] K. Žmolíková, M. Delcroix, K. Kinoshita, T. Ochiai, T. Nakatani, L. Burget, and J. Černocký, "Speaker aware neural network for target speaker extraction in speech mixtures," *IEEE J. Sel. Top. Signal Process.*, vol. 13, no. 4, pp. 800–814, 2019.
- [31] D. S. Williamson, Y. Wang, and D. Wang, "Complex ratio masking for monaural speech separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 3, pp. 483–492, 2016.
- [32] S. Horiguchi, Y. Takashima, P. Garcia, S. Watanabe, and Y. Kawaguchi, "Multi-channel end-to-end neural diarization with distributed microphones," in *Proc. ICASSP*, 2022, pp. 7332–7336.
- [33] W. Wang, X. Qin, and M. Li, "Cross-channel attention-based target speaker voice activity detection: Experimental results for the m2met challenge," in *Proc. ICASSP*, 2022, pp. 9171–9175.
- [34] E. Vincent, S. Watanabe, A. A. Nugraha, J. Barker, and R. Marxer, "An analysis of environment, microphone and data simulation mismatches in robust speech recognition," *Comput. Speech Lang.*, vol. 46, pp. 535–557, 2017.
- [35] G. Wichern, J. Antognini, M. Flynn, L. R. Zhu, E. McQuinn, D. Crow, E. Manilow, and J. Le Roux, "WHAM!: Extending speech separation to noisy environments," in *Proc. Interspeech*, 2019, pp. 1368–1372.
- [36] M. Maciejewski, G. Wichern, E. McQuinn, and J. Le Roux, "WHAMR!: Noisy and reverberant single-channel speech separation," in *Proc. ICASSP*, 2020, pp. 696–700.
- [37] F. Paissan, G. Wichern, Y. Masuyama, R. Aihara, F. G. Germain, K. Saijo, and J. Le Roux, "FastTUS: Faster task-aware unified source separation," *arXiv preprint arXiv:2507.11435*, 2025.
- [38] E. Vincent, R. Gribonval, and C. Fevotte, "Performance measurement in blind audio source separation," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 14, no. 4, pp. 1462–1469, 2006.
- [39] ITU-T P.862.2, "Wideband extension to recommendation P. 862 for the assessment of wideband telephone networks and speech codecs," 2007.
- [40] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [41] W. Zhang, J.-w. Jung, and Y. Qian, "Improving design of input condition invariant speech enhancement," in *Proc. ICASSP*, 2024, pp. 10 696–10 700.
- [42] C. K. A. Reddy, V. Gopal, and R. Cutler, "DNSMOS P.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *Proc. ICASSP*, 2022, pp. 886–890.
- [43] Y. Masuyama, X. Chang, S. Cornell, S. Watanabe, and N. Ono, "End-to-end integration of speech recognition, dereverberation, beamforming, and self-supervised learning representation," in *Proc. SLT*, 2023, pp. 260–265.