



# Separating Long-form Speech with Group-wise Permutation Invariant Training

Wangyou Zhang<sup>1,2,\*</sup>, Zhuo Chen<sup>2</sup>, Naoyuki Kanda<sup>2</sup>, Shujie Liu<sup>2</sup>, Jinyu Li<sup>2</sup>, Sefik Emre Eskimez<sup>2</sup>, Takuya Yoshioka<sup>2</sup>, Xiong Xiao<sup>2</sup>, Zhong Meng<sup>2</sup>, Yanmin Qian<sup>1</sup>, Furu Wei<sup>2</sup>

<sup>1</sup> MoE Key Lab of AI, X-LANCE Lab, CSE Dept, Shanghai Jiao Tong University, Shanghai, China  
<sup>2</sup> Microsoft Corporation

{wyz-97, yanminqian}@sjtu.edu.cn, {zhuc, Naoyuki.Kanda, shujliu, jinyli, Sefik.Eskimez, tayoshio, Xiong.Xiao, Zhong.Meng, fuwei}@microsoft.com

## Abstract

Multi-talker conversational speech processing has drawn many interests for various applications such as meeting transcription. Speech separation is often required to handle overlapped speech that is commonly observed in conversation. Although the original utterance-level permutation invariant training-based continuous speech separation approach has proven to be effective in various conditions, it lacks the ability to leverage the long-span relationship of utterances and is computationally inefficient due to the highly overlapped sliding windows. To overcome these drawbacks, we propose a novel training scheme named Group-PIT, which allows direct training of the speech separation models on the long-form speech with a low computational cost for label assignment. Two different speech separation approaches with Group-PIT are explored, including direct long-span speech separation and short-span speech separation with long-span tracking. Experiments on simulated meeting-style data demonstrate the effectiveness of our proposed approaches, especially in dealing with a very long speech input.

**Index Terms:** Continuous speech separation, permutation invariant training, long-form speech, overlapped speech

## 1. Introduction

Speech processing for multi-talker conversational speech, such as meeting recordings, is very challenging in the real world. It differs from single-talker scenarios in two main aspects. Firstly, it naturally contains overlapped speech from multiple speakers, so a speech separation process is often required. Secondly, a conversation can be of any length without any segmentation, which poses a challenge to the long-form speech processing capability of the system. There have been increasing interests in the conversational speech processing, including automatic speech recognition (ASR) [1–3], speech separation [4–7], and speaker diarization [8, 9]. In this paper, we specifically focus on the speech separation problem for long-form speech.

Continuous speech separation (CSS) [5] is a framework to convert long-form unsegmented audio into  $N$  overlap-free audio streams. In its representative instantiation with utterance-level permutation invariant training (uPIT) [10–13], the input speech is first segmented by using a sliding window with overlaps, and speech separation is independently performed on each segment to generate  $N$  separated signals. The separated signals in adjacent segments are then aligned via a stitching algorithm. This approach has not only proven to be effective in speech separation of simulated long-form signals [14, 15], but also shown large improvement in ASR [16, 17] and speaker diarization [9] tasks in realistic conversation scenarios. However,

\*Work done during internship at Microsoft.

there are some drawbacks in such a uPIT-CSS approach: (1) It is computationally inefficient due to the large overlap between adjacent windows, which is essential for better stitching performance. (2) More importantly, the uPIT-CSS approach can only model the short-span relationship of utterances, e.g. 1.6s in [4], as it assumes at most  $N$  active speakers in each window where  $N$  is typically 2. When a long window is used for local separation, the assumption is likely to be broken as  $> N$  speakers are likely to be present within a window. Therefore, its performance is limited due to the lack of access to a long-span context.

A recent study [18, 19] proposed a novel method tackling the above problems in the CSS framework. The authors show that the label assignment in long-form speech separation can be regarded as a graph coloring problem, which leads to a generalized uPIT criterion named Graph-PIT. The computational complexity in [18] scales exponentially with the number of utterances in each segment, and was later reduced to be linear in the number of utterances via dynamic programming [19].

In this paper, we aim to solve the long-span speech processing approach without changing the PIT objective function. We propose Group-PIT (gPIT in short), a simple training data construction strategy to address this problem, to allow the separation network to directly process long-form speech in both training and inference stages. We show that by carefully designing the data simulation procedure and arranging the long-form reference signal into utterance groups, the number of possible permutations in each long-form audio (e.g. 60s) can be constrained to  $N!$  regardless of the number of active speakers and utterances. This allows training of speech separation models directly on long-form speech with the same training objective as uPIT, except that it is used for utterance groups rather than individual utterances. We also explore different long-form speech processing approaches with Group-PIT. Firstly, we show that the straightforward extension of CSS to gPIT-CSS with long-span separation can better process the long-form speech, which benefits from the direct long-span modeling. Secondly, we explore a two-stage gPIT-CSS approach with short-span separation and long-span tracking. It combines the properties of the local- and long-span processing, which is suitable for conditions where the long-form training data is difficult to obtain or simulate, e.g., realistic long-form conversation speech with spontaneous speaker interactions. Our proposed methods are validated on simulated meetings based on the WSJ corpus [20, 21].

## 2. Stitching-based uPIT-CSS

We suppose the long-form input speech mixture  $\mathbf{Y}$  consists of  $U$  utterances and in total  $C$  speakers. In the CSS framework, it is assumed that at most  $N$  speakers are active at the same time




Figure 1: The ideal CSS output channels ( $N = 2$ ). Blocks in different colors represent utterances from different speakers.

so that all utterances in  $\mathbf{Y}$  can be separated and placed into  $N$  channels. Each output channel has the same length as the input, and only contains overlap-free utterances, as shown in Fig. 1.

A typical CSS pipeline with uPIT-based speech separation [4] is composed of three stages: segmentation, separation, and stitching. As illustrated in Fig. 2, the segmentation stage divides the long-form audio into several overlapped segments using a fixed-length sliding window. Each sliding window consists of three parts with  $T_h$ ,  $T_c$ , and  $T_f$  frames that represent the history, current, and future frames, respectively. The overlap length between adjacent segments is  $T_h + T_f$ . The speech separation is then performed on each segment independently to generate  $N$  overlap-free signals. Finally, the separation outputs in all segments are merged via a stitching algorithm to obtain the meeting-level separation result. This is done by first finding the best permutation of output channels with the highest overall similarity in each pair of adjacent segments and permuting them accordingly. Then, an overlap-and-average operation is performed along each channel across all segments.

Note that the uPIT-based CSS assumes the window length is small enough to only contain at most  $N$  speakers so that uPIT-based speech separation models can be trained. It makes the uPIT-CSS difficult to use a long window where more than  $N$  speakers will be likely to appear and potentially limits the modeling capacity due to the lack of access to the long-span context.

### 3. Group-PIT for long-span modeling

In this section, we introduce the proposed Group-PIT approach for long-form speech separation. First, we define the data arrangement in our proposed approach. Later, we propose two different approaches based on Group-PIT.

#### 3.1. Group-PIT and the corresponding data arrangement

As illustrated in Fig. 1, in the original CSS pipeline, the placement of some separated utterances in different output channels may not be unique. For example, the last two utterances (orange and yellow) in Ch 1 can be swapped with the last utterance (blue) in Ch 2, while still satisfying the CSS constraint introduced in Section 2. This phenomenon is common and can easily happen when a relatively long silence exists in the midst of the input speech mixture. As a result, the number of possible permutations of  $U$  separated utterances in  $N$  output channels can be up to  $N^U$ . More specifically, if we define an *utterance group* as a consecutive segment in which all utterances in  $N$  channels only have  $N!$  possible permutations, then the number of possible permutations in the CSS problem is up to  $(N!)^G$ , where  $G$  is the number of utterance groups in the CSS output. In Fig. 1, we can see that  $G = 2$  and  $N = 2$ , so there are  $(2!)^2 = 4$  permutations in total. Since this number increases exponentially with the number of utterance groups, it would be computationally expensive to extend CSS by using a longer window directly.

To remedy this issue, we propose to proactively arrange the training data so that  $G = 1$  is guaranteed for every long-form sample. This makes the number of possible permutations significantly small. In the following discussion, we adopt  $N = 2$  as three-fold overlaps are rarely observed in real meetings [22].




Figure 2: Segment-wise processing in stitching-based CSS. Separation is performed in each segment independently with a sliding window. The window length is  $T_w = T_h + T_c + T_f$ .

When simulating a long-form speech sample, we first generate two overlap-free reference signals corresponding to two output channels by iteratively appending  $U$  utterances to either of the channels based on the following rules: (i) the first utterance is appended to Ch 1, (ii) the  $u$ -th utterance is appended to the channel  $n_u \in \{1, 2\}$  where the end time,  $e_{n_u}$ , of the lastly-appended utterance is earlier than the end time,  $e_{\bar{n}_u}$ , of the lastly-appended utterance on the other channel. When the  $u$ -th utterance is appended, its onset is randomly sampled from the uniform distribution  $\mathcal{U}(e_{n_u}, e_{\bar{n}_u})$ . After generating two overlap-free reference signals, they are mixed to form the long-form audio mixture for training. This constraint guarantees that only one utterance group exists in each long-form speech sample. Given the small number of possible permutations ( $N! = 2$ ), we can apply the conventional uPIT criterion except that it is applied for utterance groups rather than individual utterances. We call this method Group-PIT.

Compared to our proposed method, Graph-PIT [18, 19] is a more generalized approach that directly extends uPIT for long-span modeling. On the other hand, our proposed method simplifies the permutation problem, thus reducing the computational cost during training. The proposed method can be viewed as a special solution, which leverages the prior knowledge in the CSS problem, of the Graph-PIT. Such prior knowledge can result in different behaviors for the separation network when  $C > 2$ , while we focus on the case of  $C = 2$  in this paper and leave such cases for future work. It should be noted that with additional constraints, it is also possible for Graph-PIT to converge to the same solution as the proposed method.

Note that modeling long-form samples containing multiple utterance groups ( $G > 1$ ) during training is potentially inefficient. From the practical perspective, it is relatively easy to detect long silence regions by applying voice activity detection (VAD) as a preprocessing. The input mixture can then be divided into chunks without such silence [19], and the separation output for each chunk can still be regarded as a single utterance group. Note that in this work, utterances with short silence in between are considered to belong to the same utterance group.

#### 3.2. gPIT-CSS with long-span separation

One straightforward way to extend the uPIT-CSS with  $N = 2$  by Group-PIT is using a longer sliding window that covers more than two utterances. In the training stage, we assume that the reference signal  $\mathbf{X}_n$  ( $n = 1, 2$ ) only contains one utterance group. The training objective is then given by:

$$\mathcal{L}^{(\text{Group-PIT})} = \min_{\pi \in \mathcal{P}_2} \sum_{n=1}^2 \mathcal{L}(\mathbf{X}_n, \hat{\mathbf{X}}_{\pi(n)}), \quad (1)$$

where  $\hat{\mathbf{X}}_{\pi(n)}$  is the  $\pi(n)$ -th output signal from the speech separation model,  $\pi \in \mathcal{P}_2$  enumerates all possible permutations for  $N = 2$  channels, and  $\pi(n)$  denotes the permuted index for the  $n$ -th channel.  $\mathcal{L}$  is the loss function in time/frequency domain.

During inference, the same stitching-based process as in Section 2 is used for processing the entire meeting, except that a much longer window can be used. It is thus possible to directly utilize the long-span audio context for better speech separation.

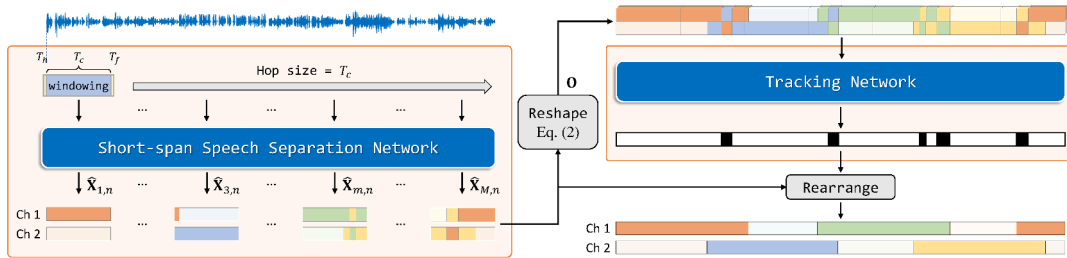


Figure 3: Proposed gPIT-CSS approach with short-span separation and long-span tracking.

### 3.3. gPIT-CSS with short-span separation and long-span tracking

The separation approach in Section 3.2 solves the long-span separation problem in one shot. However, it usually requires matched training data to maximize its advantage in long-form modeling, which is not always available in practice. For example, in the multi-channel scenario, it is challenging to simulate all the varieties in realistic long-form conversation speech that includes speaker movements. Without matched data, the long-span modeling could be potentially sub-optimal. Thus, we explore another approach to apply Group-PIT in the CSS pipeline, where the separation of one long-segment is decomposed to short-span separation and long-span tracking procedures.

The overview to process one long audio segment (such as 24s) is depicted in Fig. 3. In this approach, the long audio segment is further segmented by short sliding windows (such as 4s) with almost no overlap where  $T_h = T_f = 1$  frame. For each short window indexed by  $m$ , a short-span separation model trained with the conventional uPIT objective function is applied to generate two overlap-free signals  $\hat{\mathbf{X}}_{m,n}$  where  $n \in \{1, 2\}$ . The long-span tracking network is then applied on the separated signals from all short windows to predict the frame-wise permutation for output channels as following:

$$\mathbf{O} = \text{Concat}_{m=1}^M (\text{Splice}([\hat{\mathbf{X}}_{m,1}^T; \hat{\mathbf{X}}_{m,2}^T]^T)), \quad (2)$$

$$\mathbf{P} = \text{TrackNet}(\mathbf{O}). \quad (3)$$

Here,  $\text{Splice}(\cdot)$  denotes the operation that stacks each frame and its adjacent frames along the feature dimension.  $\text{Concat}_{m=1}^M(\cdot)$  denotes the operation to concatenate all features from  $M$  short windows along the frame dimension. The term  $\mathbf{P} \in \mathbb{R}^{1 \times T}$  represents the frame-wise permutation indicator, where “1” indicates swapping the two separation results from current frame in Ch 1 and Ch 2, while “0” means no change.  $T$  refers to the length of the input sequence. According to the permutation indicator, the short-span separation result  $\hat{\mathbf{X}}_{m,n}$  is rearranged to form the final long-span output signal for each channel. Note that the above procedure is the explanation for the speech separation of one (relatively long) audio segment, which is still shorter than the duration of an entire recording. The entire recording is processed with the stitching algorithm as used in the gPIT-CSS with long-span separation in Section 3.2.

The tracking network is trained with the cross-entropy loss:

$$\mathcal{L}_{\text{track}} = \text{CrossEntropy}(\mathbf{P}, \mathbf{R}), \quad (4)$$

where  $\mathbf{R} \in \mathbb{R}^{1 \times T}$  is the oracle frame-wise permutation. The oracle permutation label is formed by comparing the frame-wise separation result with the two-channel reference spectrum. The data arrangement from Group-PIT is applied to construct the reference signals, i.e., after tracking alignment, the final separation result should follow the CSS arrangement as shown in Fig. 1. We freeze the short-span separation model when training the tracking network.

Note that the idea of combining short-span separation and long-span tracking was already investigated in the literature

[23], but only for utterance-level mixtures. Therefore, it is still unclear how well this approach works in the CSS framework. Our proposed extension with Group-PIT naturally fills this gap.

## 4. Experimental Setup

### 4.1. Data description

We experimented with simulated multi-talker recordings based on the WSJ corpus [20, 21]. The training and development sets were simulated based on the WSJ1 [21] training set, with 283 speakers in total. The evaluation set was simulated based on the `si_dt_05` and `si_et_05` subsets from WSJ0 [20], with 18 speakers in total. The sampling rate of the audio was 16kHz. The simulation of all datasets follows the description in Section 3.1. The number of speakers ranges from 2 to 5, while the meeting length is fixed to 80s. For all datasets, we simulated two types of mixtures, i.e. partially overlapped mixtures (`partial`) and sequential mixtures (`seq.`). For training, development and evaluation, we used 27000, 2992, 2999 overlapped samples and 8000, 1500, 3000 sequential samples, respectively. For overlapped mixtures, the overlap ratio ranges from 20% to 60%. For sequential mixtures, note that we considered sequential utterances with a short pause ( $< 0.5s$ ) belong to the same utterance group, and constrain them to be assigned to different channels even they are not overlapped. This property of separation networks has been shown to be important to handle quick speaker turns in real conversation [17].

### 4.2. Network Architectures

We adopt the time-frequency masking [24] based speech separation method to examine the effectiveness of the proposed approaches. The window size and hop size for short-time Fourier transform (STFT) are 512 and 256, respectively. The loss function  $\mathcal{L}$  in Eq. (1) is the L2 loss between estimated and reference magnitude spectra. For the gPIT-CSS approach with long-span separation, we adopt the dual-path transformer (DP-transformer) [25, 26] architecture for its capability and efficiency in long sequence modeling. It consists of 16 encoder layers with 4 attention heads, and each layer has 128 attention dimensions and 1024 FF dimensions. For the gPIT-CSS approach with separation and tracking, we adopt the transformer model with 16 layers and a similar amount of parameters to the DP-transformer for short-span separation. For the tracking network, we adopt the DP-transformer architecture for long-span modeling, which consists of 16 encoder layers with 128 attention dimensions and 1024 FF dimensions. The chunk size and hop size in the inter- and intra-chunk processing in all DP-transformer models are 150 and 75, respectively. The batch size is 96 for training Group-PIT models with a 4s sliding window on 8 GPUs. For other window lengths, we adjust the batch size accordingly to fit approximately the same amount of data into each batch as long as the memory can hold. The AdamW optimizer is used for training.

Table 1: Average utterance-level SI-SNR (dB) of gPIT-CSS based long-span separation models with different window sizes.

Model	$T_{tr}$ (s)	Sliding window size (s)			
		4	16	32	60
Original partial mixture	-	-2.84			
Baseline uPIT-CSS	4	7.72	2.21	1.90	-
+ Oracle permutation	4	13.71	6.60	3.85	-
gPIT-CSS with long-span separation	4	9.02	2.70	2.40	2.31
	16	0.94	13.92	7.08	6.46
	32	1.04	13.47	14.10	12.26
	60	1.84	11.61	13.01	<b>14.58</b>
+ Oracle permutation	4	11.93	8.59	5.59	3.66
	16	8.55	15.64	11.18	8.19
	32	8.84	15.87	<b>16.12</b>	13.67
	60	8.42	15.00	15.38	15.25

Table 2: Average frame-wise accuracy (%) of gPIT-CSS based long-span separation models with different window sizes.

Model	$T_{tr}$ (s)	Sliding window size (s)			
		4	16	32	60
Baseline uPIT-CSS	4	87.54	66.44	63.64	-
	4	77.87	63.12	62.71	61.99
gPIT-CSS with long-span separation	16	56.65	94.22	77.65	77.00
	32	53.46	94.21	94.66	91.92
	60	53.53	90.43	93.40	<b>97.45</b>

## 5. Experimental Results

### 5.1. gPIT-CSS with long-span separation

As mentioned in Section 3.2, the proposed Group-PIT allows training of speech separation models on much longer segments than uPIT. Therefore, we first compared the performance of direct long-span separation models trained with different window lengths  $T_{tr}$ . The best permutation of the meeting-level separation output channels is first determined, and the oracle utterance boundaries in each channel are then used to calculate the utterance-level scale-invariant signal-to-noise ratio (SI-SNR) [27]. The overlap between adjacent windows is set to 2s by default, i.e.  $T_h = T_f = 1s$  for all models.

Table 1 shows the separation performance on the overlapped evaluation data (partial). The performance of the uPIT-CSS baseline ( $T_{tr} = 4s$ ) is presented in the 2nd and 3rd rows<sup>1</sup>, where we retrained the model from [14] by removing RIRs and noise in data simulation. Albeit not fully comparable due to different training data, the baseline performance can still serve as an indicator of the relative performance of our proposed method. It is shown that when evaluated with different sliding window lengths, models trained with a longer window tend to have better performance. This verifies our conjecture that a longer context can benefit the separation of long-form audios. In all conditions, the best performance is achieved when the same window length is used for both training and evaluation. In addition, we can observe that models trained with longer windows tend to reach the performance with oracle permutations<sup>2</sup>, which further demonstrates the effectiveness of the proposed approach.

For the sequential evaluation data (seq.), since no overlap exists, the SI-SNRs of the separation outputs tend to be very large ( $> 30dB$ ), which is inappropriate to compare due to the nonlinear scale in SI-SNR. Instead, we compare the frame-wise accuracy of speaker assignment in each output channel in Table 2. This is obtained by calculating the percentage of

<sup>1</sup>We cannot obtain the result for 60s-window because of the huge memory cost with the baseline model.

<sup>2</sup>“oracle permutations” means using the reference signal to decide the permutation of each window to stitch adjacent separated segments.

Table 3: Performance comparison of two different gPIT-CSS approaches. “stitching” denotes the gPIT-CSS approach with long-span separation, while “tracking” denotes the gPIT-CSS approach with short-span separation and long-span tracking.

Model	Approach	Tracking acc.	SI-SNR (dB)
Original partial mixture	no processing	-	2.84
	+ stitching	-	4.36
gPIT-CSS ( $T_{tr} = 2s$ )	+ tracking	91.81%	9.21
	+ oracle tracking	100%	17.16
	+ stitching	-	7.50
gPIT-CSS ( $T_{tr} = 4s$ )	+ tracking	90.81%	8.22
	+ oracle tracking	100%	17.20

speaker turns in the best frame-wise permutation based on the final meeting-level separation output. It can be seen that models trained with longer windows also show higher frame-wise accuracies on the sequential mixture, which further shows the benefit of the proposed approach.

### 5.2. gPIT-CSS with short-span separation and long-span tracking

In this section, we evaluate the gPIT-CSS approach with short-span separation and long-span tracking. In contrast with the relatively large window length and overlap length used in Section 3.2, we only use a short sliding window (2s and 4s) with a 2-frame overlap for short-span separation. The tracking network is trained and evaluated using a 24s sliding window. The overlap between adjacent tracking windows is 12s and 2s for training and evaluation, respectively. Table 3 shows the performance of tracking-based models trained with different window lengths. Although the frame-wise tracking accuracy is not low, the overall SI-SNR performance is not as good as the direct long-span separation approaches with the best stitching window configuration in Table 1, which suggests this approach is more sensitive to frame-wise tracking errors. However, such comparison is unfair because much longer overlap sizes are used to achieve good performance in Table 1, leading to higher computational overhead. If we reduce the overlap size to only 2 frames, as shown in Table 3, the performance of the long-span speech separation (denoted as “stitching” in Table 3) is severely degraded. On the other hand, the tracking-based approach can significantly improve the final separation result while enjoying a much lower computational cost<sup>3</sup>. It is especially helpful when a shorter separation window is used, as more improvement is achieved with  $T_{tr} = 2s$  over  $T_{tr} = 4s$ .

## 6. Conclusion

In this paper, we explored the long-span speech separation approaches in the meeting scenario. A novel training scheme called Group-PIT was proposed to cope with the permutation problem in long-form speech. We showed that Group-PIT-based speech separation models can be trained directly on the arranged long-form speech with the same computational complexity as in uPIT. Moreover, we explored two different Group-PIT-based speech separation approaches for long-span speech processing, and their effectiveness was validated on the simulated data based on the WSJ corpus.

## 7. Acknowledgements

Wangyou Zhang and Yanmin Qian were supported in part by China NSFC projects under Grants 62122050 and 62071288, and in part by Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102.

<sup>3</sup>The computational cost for our tracking network is roughly 1/3 of that required for uPIT-CSS with 2s overlap, and the total computational cost becomes lower even with the overhead for the tracking network.

## 8. References

- [1] S. Watanabe, M. Mandel, J. Barker, E. Vincent, A. Arora, X. Chang, S. Khudanpur, V. Manohar, D. Povey, D. Raj, D. Snyder, A. S. Subramanian, J. Trmal, B. B. Yair, C. Boeddeker, Z. Ni, Y. Fujita, S. Horiguchi, N. Kanda, T. Yoshioka, and N. Ryant, “CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings,” in *Proc. 6th International Workshop on Speech Processing in Everyday Environments (CHiME 2020)*, 2020, pp. 1–7.
- [2] X. Chang, N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, “Hypothesis stitcher for end-to-end speaker-attributed ASR on long-form multi-talker recordings,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6763–6767.
- [3] N. Kanda, X. Xiao, J. Wu, T. Zhou, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, “A comparative study of modular and joint approaches for speaker-attributed ASR on monaural long-form audio,” *arXiv preprint arXiv:2107.02852*, 2021.
- [4] T. Yoshioka, H. Erdogan, Z. Chen, X. Xiao, and F. Alleva, “Recognizing overlapped speech in meetings: A multichannel separation approach using neural networks,” in *Proc. Interspeech 2018*, 2018, pp. 3038–3042.
- [5] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, X. Xiao, and J. Li, “Continuous speech separation: Dataset and analysis,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7284–7288.
- [6] C. Li, Y. Luo, C. Han, J. Li, T. Yoshioka, T. Zhou, M. Delcroix, K. Kinoshita, C. Boeddeker, Y. Qian *et al.*, “Dual-path RNN for long recording speech separation,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 865–872.
- [7] Z.-Q. Wang and D. Wang, “Localization based sequential grouping for continuous speech separation,” *arXiv preprint arXiv:2107.06853*, 2021.
- [8] T. von Neumann, K. Kinoshita, M. Delcroix, S. Araki, T. Nakatani, and R. Haeb-Umbach, “All-neural online source separation, counting, and diarization for meeting analysis,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 91–95.
- [9] X. Xiao, N. Kanda, Z. Chen, T. Zhou, T. Yoshioka, S. Chen, Y. Zhao, G. Liu, Y. Wu, J. Wu *et al.*, “Microsoft speaker diarization system for the VoxCeleb speaker recognition challenge 2020,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5824–5828.
- [10] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [11] Y. Isik, J. Le Roux, Z. Chen, S. Watanabe, and J. R. Hershey, “Single-channel multi-speaker separation using deep clustering,” in *Proc. Interspeech 2016*, 2016, pp. 545–549.
- [12] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, “Permutation invariant training of deep models for speaker-independent multi-talker speech separation,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 241–245.
- [13] 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 Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [14] S. Chen, Y. Wu, Z. Chen, J. Wu, J. Li, T. Yoshioka, C. Wang, S. Liu, and M. Zhou, “Continuous speech separation with conformer,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 5749–5753.
- [15] S. Chen, Y. Wu, Z. Chen, T. Yoshioka, S. Liu, J. Li, and X. Yu, “Don’t shoot butterfly with rifles: Multi-channel continuous speech separation with early exit transformer,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6139–6143.
- [16] D. Raj, P. Denison, Z. Chen, H. Erdogan, Z. Huang, M. He, S. Watanabe, J. Du, T. Yoshioka, Y. Luo *et al.*, “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 897–904.
- [17] T. Yoshioka, I. Abramovski, C. Aksoylar, Z. Chen, M. David, D. Dimitriadis, Y. Gong, I. Gurvich, X. Huang, Y. Huang *et al.*, “Advances in online audio-visual meeting transcription,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 276–283.
- [18] T. von Neumann, K. Kinoshita, C. Boeddeker, M. Delcroix, and R. Haeb-Umbach, “Graph-PIT: Generalized permutation invariant training for continuous separation of arbitrary numbers of speakers,” in *Proc. Interspeech 2021*, 2021, pp. 3490–3494.
- [19] T. von Neumann, C. Boeddeker, K. Kinoshita, M. Delcroix, and R. Haeb-Umbach, “Speeding up permutation invariant training for source separation,” in *Speech Communication; 14th ITG Conference*. VDE, 2021, pp. 1–5.
- [20] J. S. Garofolo, D. Graff, D. Paul, and P. David, *LDC Catalog: CSR-I (WSJ0) Complete LDC93S6A*. Philadelphia: Linguistic Data Consortium, 1993.
- [21] Linguistic Data Consortium and NIST Multimodal Information Group, *LDC Catalog: CSR-II (WSJ1) Complete LDC94S13A*. Philadelphia: Linguistic Data Consortium, 1994.
- [22] Ö. Çetin and E. Shriberg, “Analysis of overlaps in meetings by dialog factors, hot spots, speakers, and collection site: Insights for automatic speech recognition,” in *Proc. Interspeech 2006*, 2006, pp. 293–296.
- [23] Y. Liu and D. Wang, “Divide and conquer: A deep CASA approach to talker-independent monaural speaker separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 12, pp. 2092–2102, 2019.
- [24] D. Wang and J. Chen, “Supervised speech separation based on deep learning: An overview,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [25] Y. Luo, Z. Chen, and T. Yoshioka, “Dual-path RNN: efficient long sequence modeling for time-domain single-channel speech separation,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 46–50.
- [26] C. Subakan, M. Ravanelli, S. Cornell, M. Bronzi, and J. Zhong, “Attention is all you need in speech separation,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 21–25.
- [27] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, “SDR-half-baked or well done?” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 626–630.