



Serialized Output Training by Learned Dominance

Ying Shi¹, Lantian Li³, Shi Yin⁴, Dong Wang^{2*}, Jiqing Han^{1*}

¹School of Computer Science and Technology, Harbin Institute of Technology, China

²Center for Speech and Language Technologies, BNRist, Tsinghua University, China

³School of Artificial Intelligence, Beijing University of Posts and Telecommunications, China

⁴Huawei Technologies Co., Ltd. China

*Corresponding authors: wangdong99@mails.tsinghua.edu.cn, jqhan@hit.edu.cn

Abstract

Serialized Output Training (SOT) has showcased state-of-the-art performance in multi-talker speech recognition by sequentially decoding the speech of individual speakers. To address the challenging label-permutation issue, prior methods have relied on either the Permutation Invariant Training (PIT) or the time-based First-In-First-Out (FIFO) rule. This study presents a model-based serialization strategy that incorporates an auxiliary module into the Attention Encoder-Decoder architecture, autonomously identifying the crucial factors to order the output sequence of the speech components in multi-talker speech. Experiments conducted on the LibriSpeech and LibriMix databases reveal that our approach significantly outperforms the PIT and FIFO baselines in both 2-mix and 3-mix scenarios. Further analysis shows that the serialization module identifies dominant speech components in a mixture by factors including loudness and gender, and orders speech components based on the dominance score.

Index Terms: multi-talker speech recognition, serialized output training, dominant speech

1. Introduction

Automatic Speech Recognition (ASR) technology has experienced rapid advancements in recent years, attributed in part to the introduction of end-to-end architectures [1, 2, 3] based on state-of-the-art network structures like Transformers [4, 5] and Conformers [6, 7, 8]. For example, the Conformer model [8] has achieved a Word Error Rate (WER) below 2% on the Librispeech benchmark¹. However, it is important to note that such low WERs are attained in the context of single-talker speech. In scenarios where multiple speakers talk simultaneously, single-talker systems encounter confusion in determining which speaker to focus on. Multi-talker speech recognition sets a more ambitious objective: transcribing the speech of *all* speakers, even when their signals overlap [9, 10, 11, 12].

Early research adopted a multi-head design, where each head produces transcriptions for one speaker within the overlapped speech [13, 14, 15]. A key challenge of this architecture is the uncertainty in label assignment, where it remains unclear which speaker's text labels should be assigned to each head during loss computation. While some studies address this uncertainty by leveraging biased information such as energy [13] or starting time [14], Permutation Invariant Training (PIT) has gained widespread popularity [15, 16, 17, 18].

One significant drawback of the aforementioned multi-head architecture is the necessity to pre-define the number of

speakers, with larger speaker counts leading to increased computational demands. Serialized Output Training (SOT) [19, 20], built on the Attention Encoder-Decoder (AED) architecture [21], addresses these challenges by sequentially outputting the transcriptions of each speaker. SOT utilizes a unified decoder for transcribing all speakers' speech, eliminating the need for multiple heads. However, determining the sequence of target text labels, known as *label serialization*, remains crucial, akin to label assignment in the multi-head architecture. Two prevalent serialization strategies are First-In-First-Out (FIFO) [22] and sequential PIT [19, 23]. FIFO concatenates labels of single-talker speech speaker by speaker, ordered according to starting time, while sequential PIT orders the labels of different speakers in a manner that results in minimal loss based on the output of the current model. Neither FIFO nor PIT is flawless. Sequential PIT employs an 'indeterministic' strategy, lacking a guarantee for consistent logic in serialization — a prerequisite for effective model learning from serialized labels. In contrast, FIFO is inherently deterministic and theoretically superior to PIT, which has been supported by the findings from Kanda et al. [19]. However, when the starting time bias is ambiguous, uncertainty affects both training and inference, resulting in a significant performance drop, as demonstrated in the experimental section.

In this paper, we present a novel serialization strategy for multi-talker ASR with SOT. The fundamental concept involves training an explicit serialization module to dynamically determine the label order, departing from predefined biases or the minimum-loss order associated with PIT. We hope that the model can learn to discover factors by itself and use these factors to order the speech components, akin to human decision-making on which speech component to focus first. Specifically, the serialization module comprises a connectionist temporal classification (CTC) decoder [24] integrated with the encoder of the AED architecture, trained by minimizing the *minimum* CTC loss between the encoder output and labels of each speech component. This minimum CTC loss enforces the serialization module to identify a 'dominant component' from the mixed speech to recognize, akin to humans unconsciously paying attention to the dominant speech in a cocktail party environment [25, 26, 27]. Once the serialization module is fully trained, the CTC loss assigned to each speech component, which can be regarded as a *dominance score*, is used to order the speech components in mixed signals. We, therefore, name the new approach **dominance-based serialization**.

Our contribution is twofold: (1) We present a novel dominance-based serialization strategy for multi-talker ASR with SOT. Experimental results demonstrated that this strategy outperforms the prevalent PIT approach, and beats the FIFO strategy that relies on a strong bias on starting time. (2) We conducted a deep analysis of the serialization module and com-

¹<https://paperswithcode.com/sota/speech-recognition-on-librispeech-test-clean>

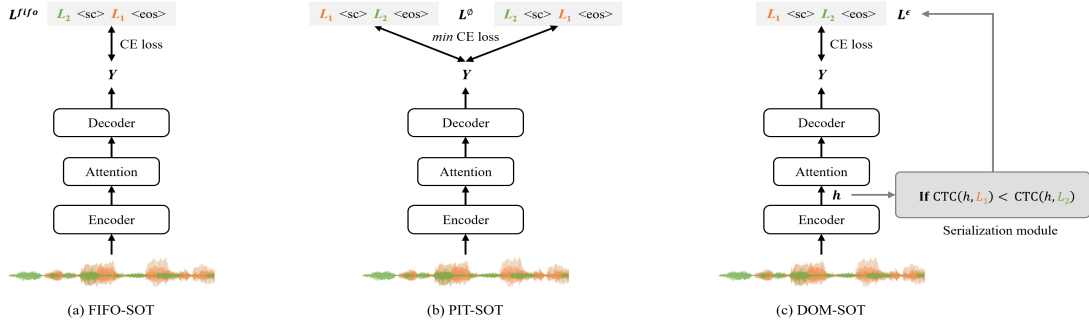


Figure 1: Diagram of FIFO-SOT, PIT-SOT, and our DOM-SOT

pared it with PIT and FIFO. This comparison effectively illustrates how the new approach functions.

2. Methods

2.1. FIFO-SOT & PIT-SOT

SOT constructs the target labels for multi-talker speech by concatenating each speaker’s labels according to specific serialization rules. Suppose a multi-speaker speech signal X involves speech from two different speakers, and these speakers appear sequentially in time, then the FIFO serialization strategy generates the target label sequence as follows:

$$L^{fif} = [L_1 \langle sc \rangle L_2 \langle sc \rangle \langle eos \rangle], \quad (1)$$

where L_1 , L_2 represent the content labels of the speech components of the two speakers, respectively. $\langle sc \rangle$ represents ‘speaker change’, and $\langle eos \rangle$ represents ‘end of sentence’. Figure 1(a) illustrates the FIFO strategy. The loss function for the FIFO-SOT with the target label L^{fif} is therefore defined as:

$$Loss^{fif} = \sum_n CE(y_n, L_n^{fif}),$$

where CE represents the cross-entropy loss and L_n^{fif} is the n^{th} token in L^{fif} , and y_n is the n^{th} output of the model.

For PIT, the order of $\{L_i\}$ in Eq. (1) can be permuted arbitrarily and the minimum CE loss resultant from these permutations is used as the loss for model training, formulated as:

$$Loss^{pit} = \min_{\phi \in \Phi} \sum_n CE(y_n, L_n^\phi),$$

where Φ represents the set of all possible permutations on $\{L_i\}$. Figure 1(b) illustrates the computation of the PIT-SOT loss.

2.2. DOM-SOT

The proposed dominance-based SOT, abbreviated by DOM-SOT, is illustrated in Figure 1(c). Overall, it is an AED architecture augmented by a serialization module (shown by the gray block); this module determines the order of the speakers when constructing the target label for model training, and the serialization module and the AED model are trained jointly. Note that once the model has been trained, the serialization module is discarded during inference.

More specifically, given the multi-talker speech X , the feature sequence h is derived by the AED encoder. h is then forwarded to the serialization module, where the CTC losses are computed between h and the labels of all speakers, i.e., $\{L_i\}$. According to these CTC losses, the speakers are arranged in

ascending order, denoted by ϵ , which is the output of the serialization module, and the target label L^ϵ is constructed according to ϵ . Finally, the entire model is trained by:

$$Loss^{dom} = \alpha * \min_{i=1, \dots, N} \{CTC(h, L_i)\} + (1 - \alpha) * CE(y, L^\epsilon),$$

where N is the number of speakers in multi-talker speech. α is a hyper-parameter to tune the behavior of the encoder. Our experiments show that a small α is preferred, probably because it avoids the encoder concentrating on a single speaker and ignoring others.

3. Experiment Settings

3.1. Training Data

The official Train_960h set from the LibriSpeech database [28] was used to construct the training data. Each training sample could be clean speech, 2-talker speech, or 3-talker speech. These three types of speech samples are equal in quantity in the final training set. When constructing 2-talker or 3-talker samples, the clean utterances were randomly sampled from the Train_960h set and then mixed using random and normalized weights.

To train the FIFO-SOT model, a 0.25s - 4s offset is imposed when mixing any two utterances. Without such offset, the model is hard to train as the bias on starting time is ambiguous. For the PIT-SOT and DOM-SOT models, the same offset occurs in 40% of the training samples, while the remaining samples are mixed without any offset.

3.2. Test Data

We tested the various SOT models on both clean speech and multi-talker speech. The clean test used the *test-clean* dataset from LibriSpeech. The 2-talker and 3-talker test samples were constructed by mixing the utterances listed in *2mix-test-clean* and *3mix-test-clean* from LibriMix [29], respectively. These utterances are mixed with an offset of 0/1/2/3 seconds, to test performance in different scenarios.

3.3. Model Configuration

All models in the experiments used 40-dimensional Mel-filter Fbank features as the input and adopted the same AED model structure where the encoder is a Conformer and the decoder is a Transformer. Before the encoder, two CNN layers perform sub-sampling. The kernel sizes and strides of both CNN layers were set to 3 and 2. The first and second CNN layer’s input/output channels were 1/256 and 256/256, respectively.

The encoder consisted of 9 Conformer blocks, each Conformer block consisted of a 4-head self-attention layer with 256 hidden units and two 2048-dimensional feed-forward layers (macaron style [30]). The decoder comprised 4 transformer blocks, each containing a 4-head 256-dimensional self-attention layer, a 256-dimensional cross-attention module, and a feed-forward linear layer with 2048 hidden units. For the DOM-SOT model, there is a serialization module. In our implementation, it was a simple linear layer that projects the encoder output to the word space, where the CTC loss was computed. The number of parameters of all the models is about 33M.

The output labels consisted of 5k wordpieces generated by the SentencePiece [31] toolkit¹ executed in the mode of the unigram language model. A blank label, an <sc> token, and <eos> token were also included in the label repository.

3.4. Training Strategy

The Adam optimizer was utilized for model training, employing a learning rate of 1e-3. All models were trained using a batch size of 32 over 60 epochs, with the initial 10 epochs designated as the warm-up phase. Upon completion of training, the average of the checkpoints from the last 10 epochs was selected as the final model. The hyper-parameter α in $Loss^{dom}$ was set to 0.1.

3.5. Speaker-Aware WER

The widely used naive WER is speaker-blind: it permutes the labels of all the speakers in the mixed speech and chooses the permutation that leads to the minimum WER by comparing the model output and the label sequence derived from the selected permutation. This naive WER tends to yield over-optimal results. For instance, if the placement of the <sc> token in the output is inaccurately predicted or not predicted at all, the WER does not significantly change, but the performance could be substantially worse if we measure how the speech of *each speaker* is recognized in average.

This motivated the design of a speaker-aware WER. Specifically, the output of the model is firstly segmented into n single-speaker segments (hypothesis) based on the <sc> token, where n represents the number of speakers predicted by the model. It is crucial to note that n does not always correspond to the actual number of speakers, m , present in the mixed speech. A greedy approach is adopted to align the n hypothesis segments and the m references. For each reference, the best-matched hypothesis is discovered in terms of WER, and the matched reference and hypothesis are removed from the reference set and the hypothesis set respectively. The match continues until either the reference set or the hypothesis set is empty. Ultimately, each unpaired reference is assigned an empty hypothesis, and each unpaired segment is assigned an empty reference. After this alignment, insertions, deletions, and substitutions are calculated with each reference/hypothesis pair, and the speaker-aware WER is computed simply by computing the ratio of these errors among all the tokens in the target label.

$$WER = \frac{\sum_{j=1}^{\max(n,m)} (S_j + I_j + D_j)}{\sum_{j=1}^{\max(n,m)} C_j},$$

where j indexes the hypothesis/reference pair, S_j , I_j , D_j are the number of substitutions, insertions, and deletions computed from the j -th pair, and C_j is the number of tokens of the j -th reference segment.

¹<https://github.com/google/sentencepiece.git>

4. Experimental Result

We compare the three SOT systems constructed with different serialization strategies: FIFO-SOT, PIT-SOT, and DOM-SOT. The test includes three groups: clean test, 2-mix talker test, and 3-mix talker test. In each of the mix-talkers tests, 4 test conditions are designed, according to the offset among speakers. Note that the models are the same for all these 3 groups and 4 conditions. We report both speaker-blind WERs and speaker-aware WERs.

4.1. FIFO-SOT vs PIT-SOT

The comparison between FIFO-SOT and PIT-SOT aligns closely with our expectations. The performance difference between the two models on clean data is insignificant. Moreover, in the 2-mix and 3-mix groups, FIFO-SOT consistently outperforms PIT-SOT when the offset is not 0s. This is not surprising as the bias on starting time is clear in these test conditions, and the bias has been explicitly exposed to the model when training the FIFO-SOT model. This also aligns with the results reported in [14, 19]. However, when the offset is 0s, the performance of FIFO-SOT drastically declines. It is noteworthy that the true offset of the speakers in the 0s-offset condition is not exactly 0 due to the leading silence within each utterance. However, this slight offset is not sufficient to form a clear bias, resulting in poor performance. This is an illuminating result and indicates that if the bias is ambiguous in the test, SOT models relying on that bias may completely fail. This is perhaps a general rule and may apply to any bias.

In comparison, PIT-SOT remains relatively stable in all the test conditions: it is slightly worse than FIFO-SOT when the bias on time is clear, but remains a rather good performance when the bias disappears. This underscores its effectiveness in dealing with multi-talkers ASR tasks.

4.2. DOM-SOT vs PIT-SOT and FIFO-SOT

First of all, DOM-SOT outperforms PIT-SOT under all testing conditions. This is a highly promising result, as the two models were trained with identical resources and nearly identical structures (the extra parameters from the additional projection layer of the serialization module in DOM-SOT can be largely ignored). This indicates that the performance gain obtained with DOM-SOT is of no cost. We therefore conclude that DOM-SOT, like PIT-SOT, is a general architecture for dealing with multi-talkers speech, and it is more powerful than PIT-SOT.

When comparing with FIFO-SOT, we can find that in the 0s-offset group, DOM-SOT shows an overwhelming advantage. Even on the conditions with a longer offset for which the bias is very clear and the training and test conditions match perfectly for FIFO-SOT, DOM-SOT can achieve comparable or even better performance compared to FIFO-SOT. We highlight that in these conditions, PIT-SOT never beats FIFO-SOT. This provides strong evidence that DOM-SOT is a more powerful architecture than PIT-SOT in tackling multi-talkers speech.

4.3. Speaker Aware vs Speaker Blind

Table 1 presents both speaker-blind WER and speaker-aware WER. It is evident that the speaker-aware WER is a more conservative metric than the speaker-blind WER, as discussed in Section 3.5. Moreover, the comparison between the two types of WERs reveals extra information about the model behavior. For instance, DOM-SOT outperforms in all the conditions in terms of speaker-blind WER but is slightly worse than FIFO-

Table 1: Performance of SOT models with different serialization strategies.

Offset	Clean	2-Mix				3-Mix			
	-	3s	2s	1s	0s	3s	2s	1s	0s
Speaker-Blind WER (%)									
FIFO-SOT	5.36	5.94	6.41	7.00	14.87	9.96	11.65	16.53	42.66
PIT-SOT	5.30	6.52	6.74	7.09	8.95	11.82	12.76	16.89	22.19
DOM-SOT	5.17	5.56	5.82	6.06	7.11	9.96	11.44	15.02	19.49
Speaker-Aware WER (%)									
FIFO-SOT	5.69	6.26	6.58	7.22	21.53	14.96	13.06	17.58	53.68
PIT-SOT	5.48	8.05	6.95	7.29	9.23	21.17	17.38	17.99	23.26
DOM-SOT	5.25	6.93	6.08	6.28	7.31	15.74	14.78	16.01	20.51

SOT in the 3s-offset conditions if the metric is speaker-aware WER. This implies that DOM-SOT produces better target labels overall but the location of the <sc> token is not predicted as well as FIFO-SOT.

5. Analysis & Discussion

In this section, we aim to discover how DOM-SOT works and outperforms PIT-SOT and FIFO-SOT. To achieve this, we conducted an analysis of the dominance scores of speech components within mixed speech samples, i.e., the CTC loss computed by the serialization module with the transcript of each speech component. The mixed samples are from the *2-mix* group, under the *0s-offset* condition and the *3s-offset* condition.

First of all, we found that nearly all samples adhered to the serialization rule based on dominance scores. More specifically, for the 2-mix data with 0s/3s offsets, 99.7% and 98.6% of the test data, respectively, were transcribed by putting the speech component with a lower CTC loss ahead. This observation indicates that **the DOM-SOT model has effectively learned to utilize CTC loss as a criterion for organizing its output sequence.**




Figure 2: Proportion of factors contributing to the dominance of dominant speech under the condition 2-mix with 0s/3s offset.

Next, we attempt to identify commonalities in dominant speech. We investigated five potential factors: Loudness, Gender, Content length, Length being overlapped, and Start time. These factors were selected based on human perception of dominant speech [25, 26, 27]. By comparing dominant and non-dominant speech across these factors — for instance, if the dominant speech is louder than the non-dominant speech —

one can identify potential biases used to determine dominant speech. To ensure the results are meaningful, the test sets were filtered to maintain balance across all factors. We also conducted a similar analysis for FIFO-SOT and PIT-SOT. Since the concept of ‘dominant speech’ is irrelevant in these methods, we define the first transcribed speech as dominant. This is acceptable as the comparison between the two speech components is what matters.

Figure 2 displays the results. It is evident that for DOM-SOT, there is a distinct bias on loudness and gender. This suggests that when two speeches are mixed, the encoder tends to prioritize recognizing louder male speech as dominant. In contrast, FIFO-SOT uses starting time as the bias, while PIT-SOT uses gender as the major bias. Such findings indicate that DOM-SOT employs a multi-bias approach, making it more robust than single-bias approaches. This clarifies why it outperforms FIFO-SOT and PIT-SOT. Importantly, the biases in DOM-SOT are identified by the model itself rather than being predefined by human intuition. Note that the biases shown with DOM-SOT and PIT-SOT are stable in both the 0s-offset and 3s-offset conditions, indicating that they are truly used by the model rather than from statistical randomness.

6. Conclusion & Future Work

We propose a model-based serialization strategy for the SOT framework, named DOM-SOT. This strategy augments the standard AED model with a serialization module that computes CTC loss for each speech component and orders the labels according to this loss. Surprisingly, this approach demonstrated a remarkable advantage over the prevalent PIT approach in our multi-talker ASR experiments, and yielded comparable or superior results to methods based on human-defined biases, like FIFO, even in test conditions where the bias is strong. Deep analysis showed that the serialization module has discovered multiple factors to order the speech components, especially loudness, and gender. This multi-factor nature explains its superior performance.

We also advocate the use of speaker-aware WER as the metric for measuring multi-talker ASR performance, as it highlights the importance of the speaker-change token. Analysis has shown that the performance of all SOT systems decreases when measured by speaker-aware WER, indicating potential errors in the speaker-change token, in particular with the proposed strategy. Future work will investigate this problem, and conduct extensive experiments on other datasets to confirm the superiority of the dominance-based approach over the PIT approach.

7. References

- [1] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," *Advances in neural information processing systems*, vol. 28, 2015.
- [2] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, "Listen, attend and spell: A neural network for large vocabulary conversational speech recognition," in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 4960–4964.
- [3] A. Zeyer, P. Bahar, K. Irie, R. Schlüter, and H. Ney, "A comparison of transformer and lstm encoder decoder models for asr," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 8–15.
- [4] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang *et al.*, "A comparative study on transformer vs rnn in speech applications," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 449–456.
- [5] Y. Wang, A. Mohamed, D. Le, C. Liu, A. Xiao, J. Mahadeokar, H. Huang, A. Tjandra, X. Zhang, F. Zhang *et al.*, "Transformer-based acoustic modeling for hybrid speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6874–6878.
- [6] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *Interspeech 2020*, 2020.
- [7] M. Zeineldeen, J. Xu, C. Lüscher, W. Michel, A. Gerstenberger, R. Schlüter, and H. Ney, "Conformer-based hybrid asr system for switchboard dataset," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7437–7441.
- [8] Y. Zhang, J. Qin, D. S. Park, W. Han, C.-C. Chiu, R. Pang, Q. V. Le, and Y. Wu, "Pushing the limits of semi-supervised learning for automatic speech recognition," *arXiv preprint arXiv:2010.10504*, 2020.
- [9] J. Wu, Z. Chen, S. Chen, Y. Wu, T. Yoshioka, N. Kanda, S. Liu, and J. Li, "Investigation of practical aspects of single channel speech separation for asr," *arXiv preprint arXiv:2107.01922*, 2021.
- [10] S. Settle, J. Le Roux, T. Hori, S. Watanabe, and J. R. Hershey, "End-to-end multi-speaker speech recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 4819–4823.
- [11] R. Masumura, N. Makishima, T. Yamane, Y. Yamazaki, S. Mizuno, M. Ihori, M. Uchida, K. Suzuki, H. Sato, T. Tanaka *et al.*, "End-to-end joint target and non-target speakers asr," *arXiv preprint arXiv:2306.02273*, 2023.
- [12] J. Kang, L. Meng, M. Cui, H. Guo, X. Wu, X. Liu, and H. Meng, "Cross-speaker encoding network for multi-talker speech recognition," *arXiv preprint arXiv:2401.04152*, 2024.
- [13] C. Weng, D. Yu, M. L. Seltzer, and J. Droppo, "Deep neural networks for single-channel multi-talker speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 10, pp. 1670–1679, 2015.
- [14] L. Lu, N. Kanda, J. Li, and Y. Gong, "Streaming end-to-end multi-talker speech recognition," *IEEE Signal Processing Letters*, vol. 28, pp. 803–807, 2021.
- [15] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *The 42nd IEEE International Conference on Acoustics, Speech and Signal Processing: The Internet of Signals*. IEEE, 2017, pp. 241–245.
- [16] Y. Qian, X. Chang, and D. Yu, "Single-channel multi-talker speech recognition with permutation invariant training," *Speech Communication*, vol. 104, pp. 1–11, 2018.
- [17] L. Meng, J. Kang, M. Cui, Y. Wang, X. Wu, and H. Meng, "A sidecar separator can convert a single-talker speech recognition system to a multi-talker one," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [18] W. Zhang, X. Chang, Y. Qian, and S. Watanabe, "Improving end-to-end single-channel multi-talker speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1385–1394, 2020.
- [19] N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, "Serialized output training for end-to-end overlapped speech recognition," *Interspeech 2020*, 2020.
- [20] N. Kanda, G. Ye, Y. Wu, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, "Large-scale pre-training of end-to-end multi-talker asr for meeting transcription with single distant microphone," *arXiv preprint arXiv:2103.16776*, 2021.
- [21] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, "End-to-end continuous speech recognition using attention-based recurrent nn: First results," in *NIPS 2014 Workshop on Deep Learning, December 2014*, 2014.
- [22] A. Tripathi, H. Lu, and H. Sak, "End-to-end multi-talker overlapping speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6129–6133.
- [23] N. Kanda, Y. Gaur, X. Wang, Z. Meng, Z. Chen, T. Zhou, and T. Yoshioka, "Joint speaker counting, speech recognition, and speaker identification for overlapped speech of any number of speakers," *Interspeech 2020*, 2020.
- [24] A. Graves and A. Graves, "Connectionist temporal classification," *Supervised sequence labelling with recurrent neural networks*, pp. 61–93, 2012.
- [25] N. Y. S. Kawata, T. Hashimoto, and R. Kawashima, "Neural mechanisms underlying concurrent listening of simultaneous speech," *Brain research*, vol. 1738, p. 146821, 2020.
- [26] M. S. Mast, "Dominance as expressed and inferred through speaking time: A meta-analysis," *Human Communication Research*, vol. 28, no. 3, pp. 420–450, 2002.
- [27] C. Jie and P. Peng, "Recognize the most dominant person in multi-party meetings using nontraditional features," in *2010 IEEE International Conference on Intelligent Computing and Intelligent Systems*, vol. 1. IEEE, 2010, pp. 312–316.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [29] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "Librimix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [30] Y. Lu, Z. Li, D. He, Z. Sun, B. Dong, T. Qin, L. Wang, and T.-y. Liu, "Understanding and improving transformer from a multi-particle dynamic system point of view," in *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020.
- [31] T. Kudo and J. Richardson, "Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 66–71.