

Disentangling Speakers in Multi-Talker Speech Recognition with Speaker-Aware CTC

Jiawen Kang, Lingwei Meng, Mingyu Cui, Yuejiao Wang, Xixin Wu, Xunying Liu, Helen Meng

The Chinese University of Hong Kong, Hong Kong SAR, China

Abstract—Multi-talker speech recognition (MTASR) faces unique challenges in disentangling and transcribing overlapping speech. To address these challenges, this paper investigates the role of Connectionist Temporal Classification (CTC) in speaker disentanglement when incorporated with Serialized Output Training (SOT) for MTASR. Our visualization reveals that CTC guides the encoder to represent different speakers in distinct temporal regions of acoustic embeddings. Leveraging this insight, we propose a novel Speaker-Aware CTC (SACTC) training objective, based on the Bayes risk CTC framework. SACTC is a tailored CTC variant for multi-talker scenarios, it explicitly models speaker disentanglement by constraining the encoder to represent different speakers’ tokens at specific time frames. When integrated with SOT, the SOT-SACTC model consistently outperforms standard SOT-CTC across various degrees of speech overlap. Specifically, we observe relative word error rate reductions of 10% overall and 15% on low-overlap speech. This work represents an initial exploration of CTC-based enhancements for MTASR tasks, offering a new perspective on speaker disentanglement in multi-talker speech recognition.¹

Index Terms—multi-talker speech recognition, speech recognition, Connectionist Temporal Classification, cocktail party problem, speech separation

I. INTRODUCTION

Natural human conversations always involve multiple speakers, with varying degrees of speech overlaps. Multi-talker speech recognition (MTASR) has emerged as a critical field, aiming to transcribe these natural conversational speech. While traditional automatic speech recognition (ASR) tasks typically perform monotonic speech-to-text sequence mapping, MTASR presents unique challenges: recognition models are required to disentangle speech from distinct speakers, and separately transcribe their speech.

In recent years, many approaches have been proposed to address this challenge. These approaches can be categorized into two types based on their ways of differentiating speakers: branched acoustic encoder (BAE) based models and serialized output training [1] based models. BAE models leverage structural priors to disentangle speakers: they separate mixed speech into independent branches, then use shared recognition blocks to transcribe different speakers in parallel. To align branches with respective speakers, permutation invariant training (PIT) [2], [3] is applied to calculate ASR loss. Yu et al. [4] first adopt a BAE-style model with PIT loss. Seki et al. [5] further extend this approach in a fully end-to-end manner. Subsequently, Chang et al. [6] incorporated transformer backbone into this framework. Further works [7]–[9] explored streaming ASR following the BAE framework. More recently, sidecar separator-based methods [10]–[12] were proposed to convert a single-talker ASR system into a multi-talker one through model-internal separation.

Another line of research lies on Serialized Output Training (SOT) [1]. The SOT approach serializes text from different speakers as a single stream, with a special token $\langle sc \rangle$ as a delimiter between speakers. In contrast to the structural priors of BAE models, this approach relies on attention mechanisms in attention encoder-decoder (AED) [13] to disambiguate speakers. This confers an advantage in

that it does not require pre-defining speakers and branch numbers, allowing it to handle a variable number of speakers. The superiority of SOT methods has been demonstrated in the M2Met challenges [14], [15], which provided challenging “in the wild” multi-talker meeting speech. SOT methods have been further enhanced with speaker information [16], [17], time boundary [18], [19], learnable speaker orders [20], large language models [21], and integrated with BAE structures as a hybrid branched SOT model [22].

In contrast to these two paradigms, there is a lack of investigation on the role of connectionist temporal classification (CTC) [23] in MTASR. CTC has become a fundamental training criterion for sequence-to-sequence tasks including speech recognition. Specifically, CTC introduces a blank token to construct alignments between input and target sequences, providing a method to compute posterior probabilities by summing over all possible alignment paths between inputs and target sequences. Compared to other ASR architecture of AED and Neural Transducer [24], [25], CTC generates all tokens in the sequence simultaneously in a non-autoregressive manner, thus enabling faster decoding speeds. CTC was also used together with AED models as a joint CTC/Attention model [26], which has long been considered one of the state-of-the-art approaches for speech recognition. In the MTASR domain, although the original SOT adopted the AED architecture without including CTC loss, many studies have empirically demonstrated that the joint CTC/Attention SOT model can effectively improve recognition accuracy on overlapped speech [18], [27]–[29]. However, given CTC’s monotonicity assumption, it is counter-intuitive that CTC can non-monotonically map overlapped speech to serialized transcriptions, and there is a lack of literature investigating these results.

In this work, we investigated the effect of CTC, especially when incorporated with SOT. Our experiments with conformer encoders reveal that CTC loss enables the acoustic encoder to represent different speakers at distinct temporal within the acoustic embeddings. Distinct from existing BAE and SOT approaches, we attribute CTC’s speaker distinction capability to its non-autoregressive reordering capability [30], [31], which is potentially a novel direction for speaker disentanglement in MTASR. Building on this insight, we proposed a novel speaker-aware CTC (SACTC) as an enhanced and tailored CTC variant for multi-talker scenarios. This SACTC explicitly models speaker disentanglement by constraining the encoder to represent different speakers’ tokens at specific temporal locations. This is achieved by the Bayes risk CTC framework, where we introduced a speaker-aware risk function to penalty CTC paths with undesired token emit. In experiments, SACTC was used as an auxiliary loss for SOT-based MTASR models. Experimental results demonstrate that the SOT-SACTC model consistently outperforms the standard SOT-CTC approach across various degrees of speech overlap. Notably, we observe word error rate reductions of 10% overall and of 15% on low-overlap speech. To our knowledge, this work represents the first exploration of CTC-based enhancements for MTASR tasks.

¹Code is available at <https://github.com/kjw11/Speaker-Aware-CTC>

II. METHODS

A. Revisit CTC in speech recognition

CTC loss guides sequence-to-sequence models by maximizing the posterior probability $p(l|x)$ of the target sequence, where $x = [x_1, \dots, x_T]$ represents the input acoustic embedding, e.g., from an acoustic encoder, and $l = [l_1, \dots, l_U]$ represents the transcription label sequence. To compensate for the length discrepancy between x and l , CTC introduces blank tokens \emptyset into the label sequence l to construct alignment paths (also as known as CTC labels) $\pi = [\pi_1, \dots, \pi_T]$ between x and l . π_t denotes the output token at time step t . A collapsing function $B(\pi) = l$ maps alignment paths to text labels by collapsing repeated consecutive labels into a single label and removing all blank labels (e.g., $B(\emptyset a \emptyset a a b b) = a a b$). Subsequently, the posterior probability $p(l|x)$ of the label sequence can be calculated by summing up the posteriors of all possible alignments:

$$P(l|x) = \sum_{\pi \in B^{-1}(l)} p(\pi|x) \quad (1)$$

where $\pi \in B^{-1}(l)$ if $B(\pi) = l$. $p(\pi|x)$ denotes the posterior probability of path π , calculated by the product of posterior probabilities of π_t cross T time steps:

$$p(\pi|x) = \prod_{t=1}^T p(\pi_t|x_t) = \prod_{t=1}^T y_{\pi_t}^t \quad (2)$$

Here $p(\pi_t|x_t)$ typically consists of linear projection and softmax function to generate frame-wise posterior of π_t at the t -th frame. We denote the output as $y_{\pi_t}^t$.

Considering the combinational explosion of permutating all alignment paths, forward-backward algorithm [32] is commonly used to calculate $P(l|x)$ effectively. First, the original label sequence l is extended by inserting \emptyset symbol between any two non-blank tokens: $l' = [\emptyset, l_1, \emptyset, \dots, \emptyset, l_U, \emptyset]$. Then recursively compute the forward-backward variables $\alpha(t, v)$ and $\beta(t, v)$:

$$\alpha(t, v) = \sum_{\substack{\pi: B(\pi_{1:t})=B(l'_{1:v}) \\ \pi_t=l'_v}} \prod_{t'=1}^t y_{\pi_{t'}}^{t'} \quad (3)$$

$$\beta(t, v) = \sum_{\substack{\pi: B(\pi_{t:T})=B(l'_{v:2U+1}) \\ \pi_t=l'_v}} \prod_{t'=t}^T y_{\pi_{t'}}^{t'} \quad (4)$$

with (t, v) is a node in CTC lattice and $1 \leq t \leq T$, $1 \leq v \leq 2U + 1$. These two variables summarized posterior of all paths passing through node (t, v) , with exact $1 : v$ prefix and $v : 2U + 1$ suffix alignment. Subsequently, for any consent time step t , enumerating all possible tokens v in l' will consider all possible paths. Therefore the CTC posteriors can be calculated by:

$$P(l|x) = \sum_{\pi \in B^{-1}(l)} p(\pi|x) = \sum_{v=1}^{2U+1} \frac{\alpha(t, v) \cdot \beta(t, v)}{y_{l'_v}^t} \quad (5)$$

B. Speaker-aware CTC based on minimizing Bayes risk

Consider a two-talker scenario with serialized target label $l = [l_1^a, \dots, l_M^a, \langle sc \rangle, l_1^b, \dots, l_N^b]$, where a and b stand for 2 speakers and $\langle sc \rangle$ token separates them. We denote the alignment path as $\pi = [\pi_1^a, \dots, \pi_K^a, \pi_{K+1}^b, \dots, \pi_T^b]$, where π_K^a represent the last $\langle sc \rangle$ token. During CTC training, we consider the case where multi-talker overlapped speech was encoded by an acoustic encoder, resulting in embedding x . Eq. 2 suggests that the information carried by x is

inherently encouraged to align with π . I.e., $[x_1, \dots, x_{K-1}]$ encodes speaker a and $[x_{K+1}, \dots, x_T]$ encodes speaker b . We note that the "speaker boundary" K varies across alignment paths, potentially confusing the encoder on how different speakers are represented. Furthermore, embedding x with a nondeterministic speaker boundary may complicate subsequent processing, e.g., hindering a cascaded ASR decoder from recognizing different speakers.

Addressing this issue, we propose a speaker-aware CTC training objective as an enhanced and tailored loss function for MTASR. The core idea is to constrain the encoder model to represent different speakers' tokens at specific time frames, which *explicitly models speaker disentanglement*. To control CTC prediction, the Bayes risk CTC (BRCTC) framework [33] was used to introduce preference over alignment paths. Specifically, BRCTC defined Bayes risk function $r(\pi)$ over posteriors of alignment paths, and the training objective is:

$$\mathcal{J}_{br}(l, x) = \sum_{\pi \in B^{-1}(l)} r(\pi) \cdot p(\pi|x) \quad (6)$$

As paths with the same concerned property could share the same risk value, we group paths by the ending point (frame) of a certain non-blank token and use group-wise risk functions, to control the encoding frames of specific speakers. Given a constant non-blank token $l_u = l'_{2u}$, the ending point of l_u is exclusive over time frames t , thus CTC posterior can be alternatively calculated by enumerating all possible frames, and the training objective can be reformulated as:

$$\mathcal{J}_{brctc}(l, x) = \sum_{t=1}^T r_g(t) \cdot \frac{\alpha(t, 2u) \cdot \hat{\beta}(t, 2u)}{y_{\pi_t}^t} \quad (7)$$

in which $\hat{\beta}(t, 2u)$ is a revised backward variable, summarizing posteriors of the paths where l_u ends at t , i.e., $t = \operatorname{argmax}_{\tau} \pi_{\tau} = l'_{2u}$. This can be achieved by eliminating non-ending paths such that $\pi_{t+1} = l'_{2u}$ for $t < T$. Accordingly:

$$\hat{\beta}(t, 2u) = \begin{cases} \beta(t, 2u) - \beta(t+1, 2u) \cdot y_{\pi_t}^t & \text{if } t < T \\ \beta(t, 2u), & \text{Otherwise} \end{cases} \quad (8)$$

The above grouping strategy inherits the original derivation in [33].

With Eq. 7 enables computing the training objective by summing over time steps, we define the following speaker-aware risk function to constrain the emitting time of tokens with consideration of their belonging speakers:

$$r_{sa}(s, t) = \begin{cases} -\frac{1}{1+e^{(\lambda(\frac{t}{T}-b))}} & \text{if } s = 1 \\ -\frac{1}{1+e^{-(\lambda(\frac{t}{T}-b))}} & \text{Otherwise} \end{cases}, \quad b = \frac{M}{M+N} \quad (9)$$

in which s represents speaker index where speakers were ordered chronologically following the first-in-first-out serialization strategy. λ is an adjustable Bayes risk factor controlling the sharpness of risks, $\lambda = 0$ leads to uniform risks for all paths. And b is a ratio of speaker utterance lengths, used to determine a speaker boundary K invariant to alignment paths. The $r_{sa}(s, t)$ overall is a conditional Sigmoid function, assigning high or low risks according to the established speaker boundary. Subsequently, the training objective for a certain l_u is:

$$\mathcal{J}'_{sa}(l, x, s, u) = \sum_{t=1}^T r_{sa}(s, t) \cdot \frac{\alpha(t, 2u) \cdot \hat{\beta}(t, 2u)}{y_{\pi_t}^t} \quad (10)$$

Note that $r_{sa}(s, t)$ is consistently < 0 , so here we minimize expected Bayes risk \mathcal{J}'_{sa} , which contrasts with maximizing a posterior as in




Fig. 1. A simplified illustration of the proposed speaker-aware risk function with CTC lattice. Red area indicates high risk and green for low risk. Tokens 1 and 2,3,4 are from different speakers. Two encouraged alignments are shown as examples.

Eq. 1. Furthermore, to constrain every token $l_u \in l$ from all speakers, the final training objective is to minimize the following:

$$\mathcal{J}_{sa}(l, x) = \frac{1}{S} \cdot \sum_{s=1}^S \left[\frac{1}{U} \cdot \sum_u \mathbb{1}(s, u) \cdot \log \mathcal{J}'_{sa}(l, x, s, u) \right] \quad (11)$$

where S is the total speaker numbers and indicator $\mathbb{1}(s, u) = 1$ if token l_u belongs to speaker s , otherwise 0. Note that if $r_{sa}(s, t)$ provides uniform risks across all paths, \mathcal{J}_{sa} degenerates to vanilla CTC, as it treats all paths equally. From this perspective, this training objective can be understood as adding a path penalty upon the vanilla CTC objective, where the penalty corresponds to the risk function.

Fig. 1 presents a simplified illustration of how the proposed training objective works. It requires the front-end encoder to disentangle separate speakers onto specific frames.

III. EXPERIMENTAL SETUP

Dataset Our experiments employed LibriSpeechMix (LSM) [1] as a benchmark dataset. This dataset is derived from the LibriSpeech (LS) [34] corpus, simulated both 2-speaker (LSM-2mix) and 3-speaker (LSM-3mix) mixed speech. As LSM only provides development and test sets, we generated 2-speaker mixed speeches for training using the similar protocol as in [1], [22]. Specifically, for each sample in the LS 960-hour training set, we randomly sample another sample with a random offset to mix with. We expect a practical MTASR model can simultaneously handle single- and multi-talker scenarios. Thus the generated mixed data was combined with the single-talker LS training set, resulting in our training set containing around 560k utterances with 1.7k hours of speech. To probe model performance on varying degrees of overlapped speech, we further divided the LSM test set into three subsets, representing low, medium, and high overlap conditions. The corresponding overlap ratios are (0, 0.2], (0.2, 0.5], and (0.5, 1.0] respectively. The overlap ratio here is defined as the duration of overlaps divided by the total duration of mixed speech. Besides, we concatenate transcriptions from separate speakers as text labels, using the first-in-first-out serialization strategy.

Model settings We implemented CTC and AED ASR models with the ESPnet2 toolkit [35]. For the CTC model, we use a conformer encoder with 12 conformer blocks. Each block has 4-head self-attention with 256 hidden units and two 1024-dimensional feed-forward layers (macaron style). The AED model has an additional transformer decoder, comprising 8 transformer blocks with also 4 heads self-attention and 256 hidden units, but a 2048-dimensional feed-forward layer. As a result, CTC model has 22.14M parameters and AED model has 34.18M parameters.

Training settings The CTC models were trained with vanilla CTC or proposed speaker-aware CTC (SACTC) objectives. And AED model was trained with sole AED loss (w/o CTC) or joint-CTC/attention

loss, where CTC weight was set as 0.3. During training, Adam optimizer was used with learning rate of $5e-4$, warm-up steps of 25k, and batch bins of 35M. Our preliminary study shows CTC converges slower than AED in MTASR, thus we trained CTC model for 80 epochs, while AED models for 50 epochs. After training, the best 10 epochs on the dev set were fused as the final models.

Metrics For single-talker condition, we used the standard word error rate (WER) as the evaluation metric. For multi-talker condition, we deployed permutation-invariant WER, a common metric for SOT approaches. This metric compares all possible permutations of speaker orders and picks up the lowest WER. Additionally, we also implemented overlap-aware WER (OA-WER) [22]. OA-WER averages WERs across various overlap ratios, aiming to balance the impact of different degrees of overlapped speech.

TABLE I
WER (%) OF VANILLA CTC AND SACTC IN MTASR. C1 AND E1 ARE THE MAIN EXPERIMENTS. "DEC." STANDS FOR DECODING.

ID	System	Librispeech		LibrispeechMix-2mix					
		dev	test	Dev	Test (Overall)	Test (Conditional)			
						low	mid	high	OA-WER
A1	SOT	4.1	4.6	7.9	9.2	9.0	8.0	12.8	9.9
B1	CTC	5.0	5.4	11.7	11.1	7.5	12.4	18.2	12.7
C1	SOT+CTC	4.3	4.5	8.4	8.8	7.1	9.0	13.1	9.7
C2	\hookrightarrow AED only dec.	4.8	5.4	11.5	12.9	11.7	12.5	17.3	13.8
C3	\hookrightarrow CTC only dec.	5.5	5.6	12.7	12.0	7.8	13.8	19.7	13.8
D1	SACTC	5.4	5.6	13.5	12.3	8.0	13.9	20.8	14.2
E1	SOT+SACTC	3.9	4.1	8.2	8.0	6.0	8.4	12.8	9.1
E2	\hookrightarrow AED only dec.	4.0	4.5	8.4	8.8	8.2	8.0	12.3	9.5
E3	\hookrightarrow CTC only dec.	5.5	5.8	12.5	11.9	8.1	13.1	19.9	13.7

IV. RESULTS AND DISCUSSIONS

In this section, we first analyze the effect of vanilla CTC in SOT-based MTASR. We then present and discuss the experimental results of the proposed SACTC approach, comparing it to vanilla CTC.

A. Analysis of vanilla CTC

Previous research has demonstrated that integrating CTC with SOT improves MTASR performance. We reproduced these experiments, with results presented in Table I, systems A to C. Comparing CTC with SOT, we observe that while CTC generally performed worse, it achieved better WER on the low-overlap subset (7.5 vs. 9.0). Moreover, incorporating CTC with SOT (C1) didn't enhance single-talker performance but improved multi-talker recognition, particularly on the low-overlap subset (9.0 \rightarrow 7.5). However, *the addition of CTC led to decreased performance on mid- and high-overlap speech*. These results validate that CTC could assist in recognizing low-overlap speech, while it degrades performance when encountering more severe overlaps.

To better understand the interaction of CTC and multi-talker speech, we examined the attention patterns in the conformer encoder for different speakers. In detail, we visualized the top 50 attended frames in self-attention for each CTC-emitted token, then accumulated attentions for tokens from two speakers in distinct colors. Fig. 2 illustrates an interesting pattern: two speakers generally attended all frames in shallower blocks, while from layer 10 onwards, two speakers began to focus on distinct regions. This aligns with our derivation in Section II-B. Notably, we did not observe this phenomenon in the sole SOT system. Fig. 3 further visualizes the attention matrices in the last conformer block. Compared to the sole SOT system, the use of CTC leads to information re-ordering: certain portions of the input embedding are attended to distinct regions of the output embedding (illustrated in Fig. 3(b) Head 4). Moreover, these repeatedly attended regions show a direct correlation to the




Fig. 2. Visualization of top-50 attended frames for two speakers (red and blue colors). Purple colors represent two speakers attending simultaneously.

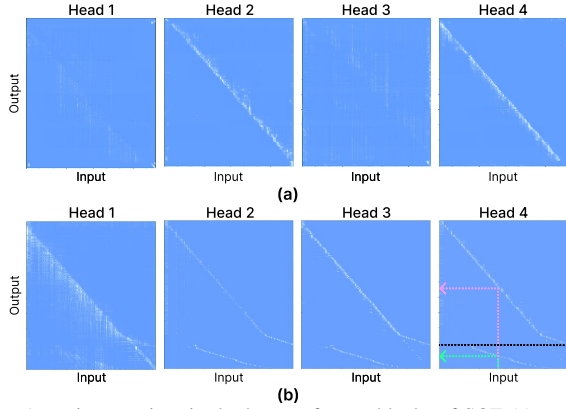


Fig. 3. Attention matrices in the last conformer blocks of SOT (a) and SOT-CTC (b) models. In (b), the overlapped area was encoded into separate output frames.

overlapped area in the input speech. These visualizations suggest that with CTC guidance, the self-attention modules disentangle different speakers along the time dimension to align with concatenated labels. Together with WER results, we hypothesize that this capacity is limited for handling severely overlapped speech.

B. Performance of SACTC

We evaluated the proposed SACTC approach using a default risk factor parameter of 15. Initially, we trained a model with the SACTC objective alone. As shown in Table I, SACTC by itself did not outperform the vanilla CTC model (B1 vs. D1). However, when combined with SOT (D1), the model showed significant improvements over vanilla CTC: overall LSM-2mix WER improved from 8.8 to 8.0, and mid-overlap WER from 12.4 to 8.4. This outcome is understandable, as SACTC is designed to enhance MTASR embedding with deterministic speaker disentangling, thus not necessarily improving token-level recognition². When integrated with SOT, SACTC enhanced low-overlap recognition similar to vanilla CTC, while mitigating performance degradation on more severe overlaps. Experiment E2 also supports this interpretation. For the model trained with SOT+SACTC, AED-only decoding led to performance gains compared to SOT+CTC, particularly improving recognition in high-overlap conditions (13.1→12.3). This suggests that SACTC resulted in embeddings with enhanced speaker discriminability. Based

²For encoder-only models, there might exist potential trade-offs between these two aspects.

TABLE II
WER (%) OF SOT+SACTC WITH DIFFERENT RISK FACTORS, WHERE 15 IS THE DEFAULT SETTING.

ID	Risk factor	Librispeech		LibrispeechMix-2mix					
		dev	test	Dev	Test (Overall)	Test (Conditional)			
						low	mid	high	OA-WER
C1	SOT+CTC	4.3	4.5	8.4	8.8	7.1	9.0	13.1	9.7
F1	5	4.3	4.6	8.4	8.8	7.3	8.9	12.7	9.6
F2	10	4.0	4.4	8.1	8.3	6.5	8.6	12.4	9.2
F3	15	3.9	4.1	8.2	8.0	6.0	8.4	12.8	9.1
F4	20	4.1	4.3	8.1	8.3	6.4	8.6	12.7	9.2

on these findings, we propose that AED-only decoding should be preferred for heavily overlapped scenarios.

Impact of hyperparameter We also investigated the impact of different risk factors (RFs) in SOT-SACTC, as shown in Table II. It shows that all tested RFs yielded improvements over the baseline C1. With a small RF of 5, the performance was close to the baseline, while the best result was achieved using RF=15.

Generalize to more speakers A key advantage of SOT-based models is their ability to generalize to a greater number of speakers than present in the training data. To test this, we evaluated our models on the LSM-3mix test set, despite all models being trained on 1 and 2 speaker scenarios. As shown in Table III, the trends observed were similar to those in the 2-speaker scenario. The SOT-CTC model significantly outperformed the SOT model on low-overlap speech (17.7% vs. 23.6%), but showed degraded performance on high-overlap speech (30.1% vs. 29.5%). In contrast, the SOT-SACTC model achieved the best performance across all conditions.

TABLE III
WER (%) OF MTASR MODELS IN 3-SPEAKER TEST SET.

System	Dev	Test (Overall)	Test (Conditional)			
			low	mid	high	OA-WER
SOT	22.9	25.3	23.6	24.3	29.5	25.8
SOT-CTC	23.5	23.6	17.7	23.3	30.1	23.7
SOT-SACTC	22.6	22.6	15.9	22.7	29.1	22.6

V. CONCLUSIONS

In this work, we investigated the effect of CTC in multi-talker speech recognition (MTASR) based on serialized output training (SOT). Our findings reveal that the CTC training objective guides the ASR encoder to encode different speakers into distinct temporal regions within acoustic embeddings. Building upon this insight, we leveraged the Bayes risk CTC framework and proposed a speaker-aware CTC (SACTC), an enhanced CTC variant tailored for MTASR. The core idea of SACTC is to constrain the encoder model to represent different speakers' tokens at specific time frames, explicitly modeling speaker disentanglement. SACTC was used as an auxiliary loss for SOT-based MTASR models in our experiments. Experimental results demonstrate that the SOT-SACTC model consistently outperforms the standard SOT-CTC approach across various degrees of speech overlap. Notably, we observe relative WER reductions of 10% overall and of 15% on low-overlap speech. To our knowledge, this work represents the first exploration of CTC-based enhancements for MTASR tasks. Future research directions may include extending SACTC to streaming seniors and exploring its potential in non-autoregressive speech recognition.

VI. ACKNOWLEDGEMENTS

This work is supported by the HKSARG Research Grants Council's Theme-based Research Grant Scheme (Project No. T45-407/19N) and the CUHK Stanley Ho Big Data Decision Research Centre.

REFERENCES

- [1] N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, "Serialized output training for end-to-end overlapped speech recognition," *arXiv preprint arXiv:2003.12687*, 2020.
- [2] 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.
- [3] 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.
- [4] D. Yu, X. Chang, and Y. Qian, "Recognizing multi-talker speech with permutation invariant training," *arXiv preprint arXiv:1704.01985*, 2017.
- [5] H. Seki, T. Hori, S. Watanabe, J. L. Roux, and J. R. Hershey, "A purely end-to-end system for multi-speaker speech recognition," *arXiv preprint arXiv:1805.05826*, 2018.
- [6] X. Chang, W. Zhang, Y. Qian, J. Le Roux, and S. Watanabe, "End-to-end multi-speaker speech recognition with transformer," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6134–6138.
- [7] 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.
- [8] D. Raj, D. Povey, and S. Khudanpur, "Surt 2.0: Advances in transducer-based multi-talker speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [9] 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.
- [10] L. Meng, J. Kang, M. Cui, Y. Wang, X. Wu, and H. Meng, "A sidetar 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.
- [11] L. Meng, J. Kang, M. Cui, H. Wu, X. Wu, and H. Meng, "Unified modeling of multi-talker overlapped speech recognition and diarization with a sidetar separator," in *Proceedings of Interspeech*, 2023, pp. 3467–3471.
- [12] L. Meng, J. Kang, Y. Wang, Z. Jin, X. Wu, X. Liu, and H. Meng, "Empowering whisper as a joint multi-talker and target-talker speech recognition system," *arXiv preprint arXiv:2407.09817*, 2024.
- [13] 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.
- [14] F. Yu, S. Zhang, Y. Fu, L. Xie, S. Zheng, Z. Du, W. Huang, P. Guo, Z. Yan, B. Ma *et al.*, "M2met: The icassp 2022 multi-channel multi-party meeting transcription challenge," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6167–6171.
- [15] Y. Liang, M. Shi, F. Yu, Y. Li, S. Zhang, Z. Du, Q. Chen, L. Xie, Y. Qian, J. Wu *et al.*, "The second multi-channel multi-party meeting transcription challenge (m2met 2.0): A benchmark for speaker-attributed asr," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [16] 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," *arXiv preprint arXiv:2006.10930*, 2020.
- [17] Z. Fan, L. Dong, J. Zhang, L. Lu, and Z. Ma, "Sa-sot: Speaker-aware serialized output training for multi-talker asr," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 9986–9990.
- [18] Y. Liang, F. Yu, Y. Li, P. Guo, S. Zhang, Q. Chen, and L. Xie, "Ba-sot: Boundary-aware serialized output training for multi-talker asr," *arXiv preprint arXiv:2305.13716*, 2023.
- [19] N. Kanda, J. Wu, Y. Wu, X. Xiao, Z. Meng, X. Wang, Y. Gaur, Z. Chen, J. Li, and T. Yoshioka, "Streaming multi-talker asr with token-level serialized output training," *arXiv preprint arXiv:2202.00842*, 2022.
- [20] Y. Shi, L. Li, S. Yin, D. Wang, and J. Han, "Serialized output training by learned dominance," *arXiv preprint arXiv:2407.03966*, 2024.
- [21] L. Meng, S. Hu, J. Kang, Z. Li, Y. Wang, W. Wu, X. Wu, X. Liu, and H. Meng, "Large language model can transcribe speech in multi-talker scenarios with versatile instructions," *arXiv preprint arXiv:2409.08596*, 2024.
- [22] J. Kang, L. Meng, M. Cui, H. Guo, X. Wu, X. Liu, and H. Meng, "Cross-speaker encoding network for multi-talker speech recognition," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 986–11 990.
- [23] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [24] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [25] K. Rao, H. Sak, and R. Prabhavalkar, "Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 193–199.
- [26] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [27] F. Yu, S. Zhang, P. Guo, Y. Fu, Z. Du, S. Zheng, W. Huang, L. Xie, Z.-H. Tan, D. Wang *et al.*, "Summary on the icassp 2022 multi-channel multi-party meeting transcription grand challenge," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9156–9160.
- [28] C. Shen, Y. Liu, W. Fan, B. Wang, S. Wen, Y. Tian, J. Zhang, J. Yang, and Z. Ma, "The volcspeech system for the icassp 2022 multi-channel multi-party meeting transcription challenge," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9181–9180.
- [29] S. Ye, P. Wang, S. Chen, X. Hu, and X. Xu, "The royalflush system of speech recognition for m2met challenge," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 9176–9185.
- [30] S.-P. Chuang, Y.-S. Chuang, C.-C. Chang, and H.-y. Lee, "Investigating the reordering capability in ctc-based non-autoregressive end-to-end speech translation," *arXiv preprint arXiv:2105.04840*, 2021.
- [31] Y. Shi, D. Wang, L. Li, and J. Han, "A glance is enough: Extract target sentence by looking at a keyword," *arXiv preprint arXiv:2310.05352*, 2023.
- [32] L. Rabiner and B. Juang, "An introduction to hidden markov models," *IEEE ASSP Magazine*, vol. 3, no. 1, pp. 4–16, 1986.
- [33] J. Tian, B. Yan, J. Yu, C. Weng, D. Yu, and S. Watanabe, "Bayes risk ctc: Controllable ctc alignment in sequence-to-sequence tasks," *arXiv preprint arXiv:2210.07499*, 2022.
- [34] 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.
- [35] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proceedings of Interspeech*, 2018, pp. 2207–2211. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1456>