



# CASA-ASR: Context-Aware Speaker-Attributed ASR

Mohan Shi<sup>1</sup>, Zhihao Du<sup>2</sup>, Qian Chen<sup>2</sup>, Fan Yu<sup>2</sup>,  
Yangze Li<sup>2</sup>, Shiliang Zhang<sup>2</sup>, Jie Zhang<sup>1</sup>, Li-Rong Dai<sup>1</sup>

<sup>1</sup>NERC-SLIP, University of Science and Technology of China (USTC), China

<sup>2</sup>Speech Lab of DAMO Academy, Alibaba Group, China

smohan@mail.ustc.edu.cn, {jzhang6; lrdai}@ustc.edu.cn, neo.dzh@alibaba-inc.com

## Abstract

Recently, speaker-attributed automatic speech recognition (SA-ASR) has attracted a wide attention, which aims at answering the question “who spoke what”. Different from modular systems, end-to-end (E2E) SA-ASR minimizes the speaker-dependent recognition errors directly and shows a promising applicability. In this paper, we propose a context-aware SA-ASR (CASA-ASR) model by enhancing the contextual modeling ability of E2E SA-ASR. Specifically, in CASA-ASR, a contextual text encoder is involved to aggregate the semantic information of the whole utterance, and a context-dependent scorer is employed to model the speaker discriminability by contrasting with speakers in the context. In addition, a two-pass decoding strategy is further proposed to fully leverage the contextual modeling ability resulting in a better recognition performance. Experimental results on AliMeeting corpus show that the proposed CASA-ASR model outperforms the original E2E SA-ASR system with a relative improvement of 11.76% in terms of speaker-dependent character error rate.

**Index Terms:** Rich transcription, speaker-attributed, multi-talker ASR, Alimeeting

## 1. Introduction

Multi-talker modeling in meeting scenarios, also known as the cocktail party problem, is one of the most challenging tasks in the field of speech signal processing [1, 2]. The purpose of speaker-attributed automatic speech recognition (SA-ASR) is not only to obtain a multi-talker transcription [3, 4, 5, 6, 7], but also to predict the speaker for each character or word in the utterance, that is, to solve the problem of “who spoke what”.

To obtain the multi-talker transcriptions, serialized output training (SOT) [7] can be used to implement multi-talker ASR, which simply concatenates the transcriptions of different speakers through a special symbol (*sc*), usually in the order of the starting time of the utterance, known as first-in first-out. Compared with permutation invariant training (PIT) [3, 4, 5], where multi-talker transcriptions are obtained through multiple decoders in the model, SOT avoids the limit of the maximum number of speakers and duplicated hypotheses, resulting in a better performance in multi-talker ASR tasks.

In literature, modular [8, 9] and end-to-end (E2E) [10, 11, 12, 13] approaches for SA-ASR have been proposed. There are three kinds of modular approaches, including frame-level diarization with SOT (FD-SOT), word-level diarization with

SOT (WD-SOT) and target-speaker separation and ASR (TS-ASR). In FD-SOT, the results of frame-level speaker diarization and the hypotheses of SOT-based ASR are aligned to obtain speaker-attributed transcriptions. In WD-SOT, speaker-attributed transcriptions are obtained by word-level speaker diarization on the hypotheses of SOT-based ASR. TS-ASR was proposed to train the joint model of target speaker separation and ASR, which achieves the best SA-ASR performance among the above three kinds of modular approaches. In contrast to modular approaches that do not directly optimize the speaker-dependent objective, E2E SA-ASR approaches [10, 11, 12, 13] mainly consists of four modules: ASR encoder (ASR-Enc), ASR decoder (ASR-Dec), speaker encoder (Spk-Enc) and speaker decoder (Spk-Dec). The ASR-Enc and Spk-Enc obtain the speech representations for multi-talker ASR and speaker identification, respectively. The Spk-Dec mainly takes the token sequence and the representation obtained by Spk-Enc as input and outputs the speaker representation for each token. The posterior probability of each speaker for each token is obtained by cosine similarity scoring between the output of Spk-Dec and speaker profiles. The speaker profile weighted by the posterior probability is fed to the ASR-Dec to help resolve multi-talker ASR transcriptions.

Although the E2E SA-ASR paradigm shows a promising applicability in multi-talker scenarios, the contextual information (that might be beneficial for understanding the acoustic scene of “who spoke what”) was not sufficiently taken into account in existing E2E approaches. This can be interpreted from two aspects. For instance, in [13] the Spk-Dec takes the token sequence as the query for source-target attention and the speech representation as the key and value to generate the speaker representation for each token. The problem lies in that the raw token sequence lacks context information, which may cause an inaccurate attention range of each speaker and a non-informative speaker representation. Besides, the involved cosine similarity is a kind of context-independent scorer (CI-Scorer) when calculating the posterior probability of each speaker for each token, which is another source of context insufficiency and also lacks nonlinear modeling ability.

In this work, we therefore propose a context-aware speaker-attributed ASR (CASA-ASR) to enhance the context modeling ability accordingly. We first introduce a contextual text encoder (Context-Enc) to aggregate the semantic information of the whole utterance to obtain a higher quality speaker representation. Then, a context-dependent scorer (CD-Scorer) is employed to model the local speaker discriminability by contrasting with speakers in the context. In addition, we introduce a two-pass decoding to enable Context-Enc and CD-Scorer to obtain more complete context information for better speaker identification performance. Finally, a skip connection is introduced

This work was done while the first author was an intern at Speech Lab of DAMO Academy, Alibaba Group, and supported by the National Natural Science Foundation of China (62101523), Hefei Municipal Natural Science Foundation (2022012) and USTC Research Funds of the Double First-Class Initiative (YD2100002008). (Correspondence: jzhang6@ustc.edu.cn)

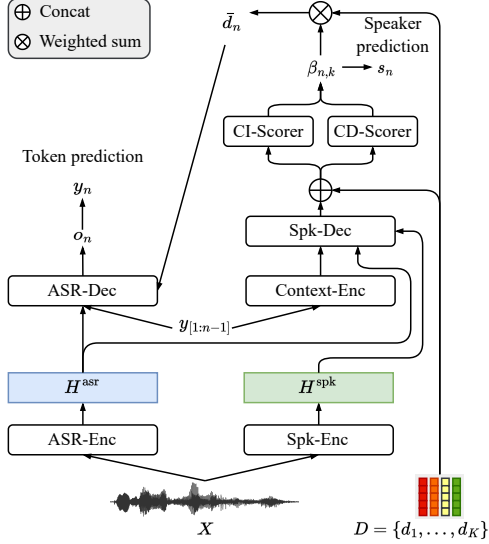


Figure 1: The overall structure of CASA-ASR.

between the raw speaker representation and the output of Spk-Dec to make the speaker information easier to flow into deep representations. Experimental results in the monaural setting of the Alimeeting corpus [14, 15] show the superiority of the proposed CASA-ASR method.

## 2. Context-aware Speaker-attributed ASR

### 2.1. SOT-based E2E SA-ASR

In the framework of SOT, the transcriptions of multiple speakers are connected by a special ‘‘speaker change’’ symbol  $\langle sc \rangle$  according to the emission time. Let the acoustic features be denoted by  $X \in \mathbb{R}^{f^a \times l^a}$  and the set of speaker profiles by  $D = \{d_k \in \mathbb{R}^{f^d} | k = 1, \dots, K\}$ , where  $f^a$  and  $l^a$  are the dimension and length of  $X$ , respectively,  $K$  is the total number of speakers in the profile,  $d_k$  is the profile of the  $k$ -th speaker, and  $f^d$  is the dimension of a profile. The goal of E2E SA-ASR is to estimate the SOT-style transcription  $Y = (y_n \in \{1, \dots, |\mathcal{V}|\} | n = 1, \dots, N)$  and the corresponding speaker of each token  $S = (s_n \in \{1, \dots, K\} | n = 1, \dots, N)$  with the inputs of  $X$  and  $D$ , where  $|\mathcal{V}|$  is the size of vocabulary  $\mathcal{V}$  containing the  $\langle sc \rangle$  symbol,  $y_n$  and  $s_n$  are the word index and speaker index for the  $n$ -th token, respectively.

### 2.2. CASA-ASR

To obtain speaker-attributed transcriptions, E2E SA-ASR mainly includes *ASR branch* and *speaker branch*. The proposed CASA-ASR is obtained by enhancing the context-aware modeling ability of E2E SA-ASR. The overall structure is shown in Figure 1.

#### 2.2.1. ASR encoder and speaker encoder

We use Conformer [16] as ASR encoder (ASR-Enc). Given the acoustic input  $X$ , the output of ASR-Enc is denoted as  $H^{asr} \in \mathbb{R}^{f^h \times l^h}$ , where  $f^h$  and  $l^h$  are the embedding dimension and the length of the sequence, respectively. We use Res2Net [17] as speaker encoder (Spk-Enc) to transform acoustic features  $X$  into speaker embedding  $H^{spk} \in \mathbb{R}^{f^h \times l^h}$ .

#### 2.2.2. ASR decoder

The speaker-aware transformer decoder is employed to predict the multi-talker transcriptions, where the weighted speaker pro-

file  $\bar{d}_n \in \mathbb{R}^{f^d}$  obtained from the weighted summation of  $D$  is added into its first layer and  $f^d$  represents the dimension of the speaker profile. The calculation procedure of ASR decoder (ASR-Dec) can be summarized as follows:

$$z_{[1:n-1],1}^{asr} = \text{PosEnc}(\text{Embed}(y_{[1:n-1]})), \quad (1)$$

$$\bar{z}_{n-1,l}^{asr} = z_{n-1,l}^{asr} + \text{MHA}_l^{asr\text{-self}}(z_{n-1,l}^{asr}, z_{[1:n-1],l}^{asr}, z_{[1:n-1],l}^{asr}), \quad (2)$$

$$\bar{z}_{n-1,l}^{asr} = \bar{z}_{n-1,l}^{asr} + \text{MHA}_l^{asr\text{-src}}(\bar{z}_{n-1,l}^{asr}, H^{asr}, H^{asr}), \quad (3)$$

$$z_{n-1,l+1}^{asr} = \begin{cases} \bar{z}_{n-1,l}^{asr} + \text{FF}_l^{asr}(\bar{z}_{n-1,l}^{asr} + W^{spk} \cdot \bar{d}_n) & (l = 1) \\ \bar{z}_{n-1,l}^{asr} + \text{FF}_l^{asr}(\bar{z}_{n-1,l}^{asr}) & (l > 1) \end{cases} \quad (4)$$

$$o_n = \text{Softmax}(W^o \cdot z_{n-1,L}^{asr} + b^o). \quad (5)$$

The token sequence is first processed to obtain  $z_{[1:n-1],1}^{asr} \in \mathbb{R}^{f^d \times n-1}$  by embedding and positional encoding. For each layer  $l$  in ASR-Dec, the MHA operation and the source-target-(src-) attention operation are applied at  $z_{n-1,l}^{asr}$  to obtain  $\bar{z}_{n-1,l}^{asr}$ . Then, the position-wise feed forward layer is applied to obtain the input of the next layer  $z_{n-1,l+1}^{asr}$ . When  $l = 1$ ,  $\bar{d}_n$  is projected through a linear transformation by multiplying with  $W^{spk} \in \mathbb{R}^{f^h \times f^d}$ . Then, it is added with  $\bar{z}_{n-1,l}^{asr}$  and sent to the feed-forward layer. Finally, the outputs of the last decoder layer is passed through a linear transformation and softmax activation function to obtain  $o_n$ . The posterior probability of token  $i$  (i.e., the  $i$ -th token in  $\mathcal{V}$ ) at the  $n$ -th decoder step is represented as

$$Pr(y_n = i | y_{[1:n-1]}, s_{[1:n]}, X, D) = o_{n,i}, \quad (6)$$

where  $o_{n,i}$  represents the  $i$ -th element of  $o_n$ .

#### 2.2.3. Context-aware speaker decoder

Considering that the semantic information of  $\bar{z}_{[1:n-1],1}^{asr}$  is relatively weak, in order to obtain a better context-aware speaker representation, we apply a contextual text encoder (Context-Enc) to  $\bar{z}_{[1:n-1],1}^{asr}$  to obtain  $\bar{z}_{[1:n-1]}^{\text{context}}$ . The first layer of Spk-Dec with Context-Enc can be represented as follows:

$$\bar{z}_{n-1}^{\text{context}} = \text{Context-Enc}(\bar{z}_{[1:n-1],1}^{asr}), \quad (7)$$

$$\bar{z}_{n-1,1}^{\text{spk}} = \text{MHA}_{n-1,1}^{\text{spk}\text{-src}}(\bar{z}_{n-1}^{\text{context}}, H^{asr}, H^{spk}), \quad (8)$$

$$z_{n-1,2}^{\text{spk}} = \bar{z}_{n-1,1}^{\text{spk}} + \text{FF}_1^{\text{spk}}(\bar{z}_{n-1,1}^{\text{spk}}). \quad (9)$$

The Context-Enc consists of several layers of transformer encoder containing self-attention and feed forward layers. By treating  $\bar{z}_{n-1}^{\text{context}}$  as the query,  $H^{asr}$  as the key and  $H^{spk}$  as the value, a context-dependent speaker representation  $\bar{z}_{n-1,1}^{\text{spk}}$  are obtained for each token through the src-attention. Since  $\bar{z}_{[1:n-1]}^{\text{context}}$  contains rich contextual information, it can aggregate more speaker embeddings  $H^{spk}$  by interacting with  $H^{asr}$ , which leads to a higher-quality speaker representation for each token. We employ a standard transformer-based decoder consisting of self-attention, source-attention and point-wise feed forward layers as the remaining layer of Spk-Dec to predict the speaker representation for each token, which is computed as:

$$\bar{z}_{n-1,l}^{\text{spk}} = z_{n-1,l}^{\text{spk}} + \text{MHA}_l^{\text{spk}\text{-self}}(z_{n-1,l}^{\text{spk}}, z_{[1:n-1],l}^{\text{spk}}, z_{[1:n-1],l}^{\text{spk}}), \quad (10)$$

$$\bar{z}_{n-1,l}^{\text{spk}} = \bar{z}_{n-1,l}^{\text{spk}} + \text{MHA}_l^{\text{spk}\text{-src}}(\bar{z}_{n-1,l}^{\text{spk}}, H^{spk}, H^{spk}), \quad (11)$$

$$z_{n-1,l+1}^{\text{spk}} = \bar{z}_{n-1,l}^{\text{spk}} + \text{FF}_l^{\text{spk}}(\bar{z}_{n-1,l}^{\text{spk}}), \quad (12)$$

$$q_n = W^q \cdot (z_{n-1,L}^{\text{spk}} + \bar{z}_{n-1,1}^{\text{spk}}). \quad (13)$$

Since the Spk-Enc is pre-trained, the initial speaker representation from the first layer of Spk-Dec is also meaningful. Therefore, we add a skip connection from the first layer to the output layer to make the speaker information easier to flow into deep representations, i.e., the speaker representation for each token  $\bar{z}_{n-1,1}^{\text{spk}}$  is added to the  $L^{\text{spk}}$ -th layer output. Finally, the output  $q_n$  is obtained by linear transformation.

#### 2.2.4. CI-CD Scorer

In E2E SA-ASR [13], the output of Spk-Dec  $q_n$  and the  $k$ -th speaker vector  $d_k$  are computed by cosine similarity to obtain the weight  $\beta_{n,k}$  of  $k$ -th speaker for  $n$ -th token. However, the cosine similarity is a context-independent scorer (CI-Scorer), which lacks context information and nonlinear modeling ability. Inspired by [18], we employ an extra context-dependent scorer (CD-Scorer) to enhance the context-aware and nonlinear modeling ability, resulting in the proposed CI-CD Scorer:

$$\text{score}_{n,k}^{\text{CI}} = \cos(q_n, d_k), \quad (14)$$

$$\text{score}_{n,k}^{\text{CD}} = \tanh(\text{CD-Scorer}(q_{[1:n]}, d_k)), \quad (15)$$

$$\beta_{n,k} = \frac{\exp(\text{score}_{n,k}^{\text{CI}} + \text{score}_{n,k}^{\text{CD}})}{\sum_j^K \exp(\text{score}_{n,j}^{\text{CI}} + \text{score}_{n,j}^{\text{CD}})}, \quad (16)$$

$$\bar{d}_n = \sum_{k=1}^K \beta_{n,k} d_k. \quad (17)$$

where  $\cos$  means the cosine similarity. The concatenation of  $q_{[1:n]} \in \mathbb{R}^{f^d \times n}$  with  $d_k \in \mathbb{R}^{f^d}$  is used as the input of CD-Scorer, which consists of several transformer encoder layers. The output of CD-Scorer is passed through a tanh function to obtain  $\text{score}_{n,k}^{\text{CD}}$ , which matches the range of  $[-1, 1]$  as  $\text{score}_{n,k}^{\text{CI}}$ . Then,  $\text{score}_{n,k}^{\text{CI}}$  and  $\text{score}_{n,k}^{\text{CD}}$  are added and normalized over  $K$  speakers to get  $\beta_{n,k} \in \mathbb{R}$  as in (16). Finally,  $\bar{d}_n$  is obtained by the weighted sum of  $d_k$  and  $\beta_{n,k}$ . The posterior probability of person  $k$  speaking the  $n$ -th token is given by

$$\beta_{n,k} = P_r(s_n = k | y_{[1:n-1]}, s_{[1:n-1]}, X, D). \quad (18)$$

The entire model is trained as a joint task using the loss function:

$$\text{Loss}^{\text{joint}} = \lambda \text{Loss}^{\text{spk}}(\hat{s}, s) + (1 - \lambda) \text{Loss}^{\text{asr}}(\hat{y}, y) \quad (19)$$

where  $s$  and  $y$  are the true speaker and transcription, and  $\hat{s}$  and  $\hat{y}$  are the corresponding hypotheses, respectively. The cross-entropy function is used as the speaker loss while joint attention-CTC loss [19, 20] is used to train the ASR model.

#### 2.2.5. Two-Pass Decoding

Context-Enc and CD-Scorer need the whole sequence to achieve a better context modeling. However, at each step of decoding, the input of the Context-Enc and CD-Scorer is the partial sequence obtained from the previous steps, i.e.,  $\bar{z}_{[1:n-1],1}^{\text{asr}}$ ,  $q_{[1:n]}$ , which leads to insufficient global context information. Therefore, after the best hypothesis of ASR  $y_{[1:N]}$  is obtained in the first-pass decoding, we take  $y_{[1:N]}$  as the input for the second-pass decoding, which is only used to predict speakers. In the second-pass decoding, the input of Context-Enc and CD-Scorer has global context, i.e.,  $\bar{z}_{[1:N],1}^{\text{asr}}$ ,  $q_{[1:N]}$ , which can make the speaker identification more accurate. Note that as in the second-pass decoding the input of Context-Enc and CD-Scorer is complete, the speaker prediction can therefore be obtained in one feed-forward without any auto-regression. With the two-pass decoding, the posterior probability of person  $k$  speaking the  $n$ -th token is thus given by:

$$\beta_{n,k}^{\text{two-pass}} = P_r(s_n^{\text{two-pass}} = k | y_{[1:N]}, s_{[1:N]}, X, D). \quad (20)$$

## 3. Experiments

### 3.1. Dataset and evaluation metrics

In this work, Alimeeting corpus is used to evaluate the performance, which is an open source Mandarin corpus of meeting scenario. This dataset contains 104.75 hours for training (Train), 4 hours for evaluation (Eval) and 10 hours for test (Test). Among them, the Train and Eval sets contain not only the 8-channel far-field audios recorded by a microphone array (*Ali-far*), but also the single-channel near-field audios recorded by the headset microphone of the participants (*Ali-near*), while the Test set only contains the far-field audios. Because a monaural setup is used for training and evaluation in this work, we apply CDDMA beamforming [21, 22] to *Ali-far* to obtain a single-channel data set *Ali-far-bf*. The prefixes *Train-*, *Eval-* and *Test-* stand for different sets. For the speaker profile, we use a 256-dim d-vector (i.e.,  $f^d = 256$ ) which is extracted by Res2Net pre-trained on CN-Celeb corpus [23]. Since there are at most four participants in each scenario, in case of meetings with less participants, the profiles are padded with speakers from other meetings.

We use speaker independent- (SI-) and speaker dependent- (SD-) character error rate (CER) [24] as evaluation metrics. SI-CER is used to evaluate the multi-talker ASR task in the SOT framework, which is computed in the same way as the normal CER, ignoring the speaker labels. While SD-CER is calculated by matching the ASR hypothesis to the corresponding speaker reference transcription, which is a rigorous evaluation metric used to evaluate SA-ASR in meeting scenarios.

### 3.2. Model configuration

The ASR-Enc contains 12 layers of conformer with 4-head multi-head attention (MHA) and convolutional kernel size of 15, where the dimensions of MHA and feed-forward network (FFN) are set to be 256 (i.e.,  $f^h = 256$ ) and 2048, respectively. The ASR-Dec contains 6 layers of transformer decoder. The Spk-Enc has the same architecture as the d-vector extractor except that an additional linear layer is added to map the  $f^d$ -dim output into  $f^h$ -dim embedding. Spk-Enc is initialized with the pretrained d-vector extractor. The Spk-Dec has 3 layers, i.e.,  $L^{\text{spk}} = 3$ . Both Context-Enc and CD-Scorer consist of 4 transformer encoder layers with 4-head MHA. The settings of MHA and FFN modules are similar to ASR-Enc. The weight of the additional CTC loss in [20] is set to be 0.3.

### 3.3. Training details

In this work, we use the 80-dimensional log Mel filterbank (Fbank) as the input feature. The window size is 32 ms and the window shift is 8 ms. The vocabulary used for ASR consists of 4950 common Chinese characters ( $|\mathcal{V}| = 4950$ ). We use the ESPnet toolkit [25] to train the models.

First, we train a SOT-based multi-talker ASR model under the above configurations using *Train-Ali-far-bf* and *Train-Ali-near*. Then, we use the pre-trained parameters of the ASR model and d-vector extractor to initialize ASR-Enc, ASR-Dec, and Spk-Enc. Finally, we jointly fine-tune the model using *Train-Ali-far-bf* and *Train-Ali-near*, where the weight  $\lambda$  for speaker loss is set to 0.5. The total amount of data used in the experiment is consistent with that in [8] to ensure fair comparison. We used the Adam [26] optimizer to jointly train the model for 60 epochs with a warmup for 2000 steps, and the maximum learning rate is set to  $5e-4$ . The speaker with the highest averaged  $\beta_{n,k}$  of tokens between two  $\langle sc \rangle$  is selected as the predicted speaker of that utterance.

Table 2: Ablation study on different modules of the proposed CASA-ASR in terms of SI-CER (%) and SD-CER (%).

Approach	Parameters	SI-CER			SD-CER		
		Eval	Test	Average	Eval	Test	Average
E2E SA-ASR	60.07 M	26.4	28.1	27.6	31.8	34.7	34.0
+ skip connection	60.07 M	26.5	28.1	27.7	30.4	34.7	33.6
+ CD-Scorer	65.46 M	26.5	28.1	27.7	29.7	32.5	31.8
+ two-pass decoding	65.46 M	26.5	28.1	27.7	28.7	31.4	30.7
+ Context-Enc (CASA-ASR)	70.79 M	<b>26.3</b>	<b>28.0</b>	<b>27.6</b>	<b>27.9</b>	<b>30.8</b>	<b>30.0</b>

Table 3: Comparison of recognition performance on different data settings in terms of SI-CER (%) and SD-CER (%).

Data Selection	Duration	SI-CER			SD-CER		
		Eval	Test	Average	Eval	Test	Average
<i>Train-Ali-far-bf</i>	105 Hours	27.1	28.6	28.2	29.9	32.2	31.6
+ <i>Train-Ali-near</i>	210 Hours	26.4	28.2	27.7	29.0	32.2	31.4
+ interference for <i>Train-Ali-near</i>	210 Hours	26.4	28.2	27.8	28.5	30.8	30.2
+ interference for <i>Train-Ali-far-bf</i>	210 Hours	<b>26.3</b>	<b>28.0</b>	<b>27.6</b>	<b>27.9</b>	<b>30.8</b>	<b>30.0</b>

Table 1: Comparison of various approaches on AliMeeting sets in terms of SD-CER (%).

Approach	SD-CER		
	Eval	Test	Average
FD-SOT [8]	41.0	41.2	41.2
WD-SOT [8]	36.0	37.1	36.8
TS-ASR [8]	32.5	35.1	34.4
E2E SA-ASR [13]	31.8	34.7	34.0
CASA-ASR (ours)	<b>27.9</b>	<b>30.8</b>	<b>30.0</b>

\*: This models is re-implemented by ourselves.

### 3.4. Experimental Results

#### 3.4.1. Comparison of various SA-ASR approaches

Table 1 shows the comparison of the proposed CASA-ASR, three modular approaches (i.e., FD-SOT, WD-SOT, TS-ASR) and the E2E SA-ASR. As expected, the E2E SA-ASR slightly outperforms the best modular approach, TS-ASR, in terms of the average SD-CER on Eval and Test sets (from 34.4% to 34.0%). The proposed method CASA-ASR achieves 11.8% (from 34.0% to 30.0%) relative SD-CER reduction compared to E2E SA-ASR.

#### 3.4.2. The ablation study of CASA-ASR

We show the ablation study on different modules of the proposed CASA-ASR in Tabel 2. Skip connection is applied to assist speaker information flow into deep representations, leading to a slight SD-CER reduction (from 34.0% to 33.6%) on Eval and Test sets. Involving the CD-Scorer leads to a large relative SD-CER reduction (from 33.6% to 31.8%), which indicates that CD-Scorer can improve the performance of SA-ASR by contrasting with other speakers in the context. Due to the properties of autogressive decoding, the context of each step is incomplete when decoding. The two-pass decoding is involved to solve this problem, which further reduces the average SD-CER (from 31.8% to 30.7%). The Context-Enc is involved to obtain the final CASA-ASR, which achieves the best performance in terms of both SD-CER (30.0%) and SI-CER (27.6%). As shown in Fig. 2, by involving the Context-Enc, the tokens (query for src-attention) between two speaker change symbols <sc> have a more focused attention weights for the speech representation. This indicates that the Context-Enc can aggregate the context information of the whole sentence to obtain a higher-quality speaker representation through source-target attention.

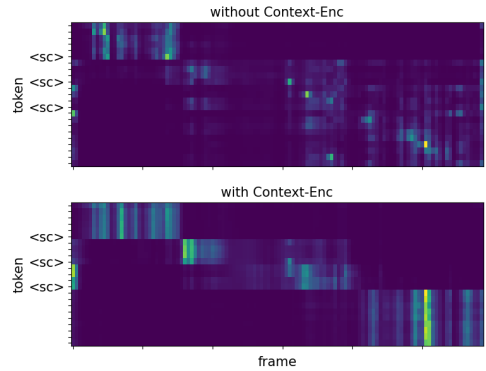


Figure 2: Visualization of source-target attention with and without Context-Enc.

#### 3.4.3. Comparison of performance on different data settings

We also investigate the SA-ASR performance under different data selection and the results are shown in Tabel 3. In case of only training with *Train-Ali-far-bf*, we achieve an average SI-CER of 28.2% and SD-CER of 31.6% on Eval and Test sets. Simply adding *Train-Ali-near* to the training set can improve the performance in terms of SI-CER (from 28.2% to 27.7%). However, it can not take a substantial improvement for SD-CER because there is only one speaker in each audio in *Train-Ali-near*. After adding interference speakers to *Train-Ali-near*, the average SD-CER is significantly reduced (from 31.4% to 30.2%). Additionally adding the interference speaker to the *Train-Ali-far-bf* achieves the lowest SD-CER (30.0%).

## 4. Conclusion

In this paper, we proposed a context-aware speaker-attributed ASR model by enhancing the contextual modeling capability of E2E speaker-attributed ASR. It was shown that adding skip connection can make the speaker information easier to flow into deep representations. Then Context-Enc enhances the context-aware modeling and obtains a higher-quality speaker representation for each token. With the CD-Scorer, stronger context awareness and nonlinear modeling ability are involved into the model. The proposed two-pass decoding strategy makes the context of Context-Enc and CD-Scorer more complete, which further improves performance. The proposed CASA-ASR approach achieves the best performance on the open-source Alimeeting corpus.

## 5. References

- [1] J. G. Fiscus, J. Ajot, M. Michel, and J. S. Garofolo, “The rich transcription 2006 spring meeting recognition evaluation,” in *MLMI*, ser. Lecture Notes in Computer Science, vol. 4299. Springer, 2006, pp. 309–322.
- [2] J. G. Fiscus, J. Ajot, and J. S. Garofolo, “The rich transcription 2007 meeting recognition evaluation,” in *CLEAR*, ser. Lecture Notes in Computer Science, vol. 4625. Springer, 2007, pp. 373–389.
- [3] D. Yu, X. Chang, and Y. Qian, “Recognizing multi-talker speech with permutation invariant training,” in *INTERSPEECH*. ISCA, 2017, pp. 2456–2460.
- [4] X. Chang, W. Zhang, Y. Qian, J. L. Roux, and S. Watanabe, “Mimo-speech: End-to-end multi-channel multi-speaker speech recognition,” in *ASRU*. IEEE, 2019, pp. 237–244.
- [5] W. Zhang, X. Chang, Y. Qian, and S. Watanabe, “Improving end-to-end single-channel multi-talker speech recognition,” *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 28, pp. 1385–1394, 2020.
- [6] Z. Chen, J. Droppo, J. Li, and W. Xiong, “Progressive joint modeling in unsupervised single-channel overlapped speech recognition,” *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 26, no. 1, pp. 184–196, 2018.
- [7] N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, “Serialized output training for end-to-end overlapped speech recognition,” in *INTERSPEECH*. ISCA, 2020, pp. 2797–2801.
- [8] F. Yu, Z. Du, S. Zhang, Y. Lin, and L. Xie, “A comparative study on speaker-attributed automatic speech recognition in multi-party meetings,” in *INTERSPEECH*. ISCA, 2022, pp. 560–564.
- [9] M. Shi, J. Zhang, Z. Du, F. Yu, S. Zhang, and L.-R. Dai, “A comparative study on multichannel speaker-attributed automatic speech recognition in multi-party meetings,” *arXiv preprint arXiv:2211.00511*, 2022.
- [10] 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,” in *INTERSPEECH*. ISCA, 2020, pp. 36–40.
- [11] N. Kanda, X. Chang, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, “Investigation of end-to-end speaker-attributed ASR for continuous multi-talker recordings,” in *SLT*. IEEE, 2021, pp. 809–816.
- [12] 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 *ICASSP*. IEEE, 2021, pp. 6763–6767.
- [13] N. Kanda, G. Ye, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, “End-to-end speaker-attributed ASR with transformer,” in *Interspeech*. ISCA, 2021, pp. 4413–4417.
- [14] F. Yu, S. Zhang, Y. Fu, L. Xie, S. Zheng, Z. Du *et al.*, “M2Met: The ICASSP 2022 multi-channel multi-party meeting transcription challenge,” in *ICASSP*. IEEE, 2022, pp. 6167–6171.
- [15] F. Yu, S. Zhang, P. Guo, Y. Fu, Z. Du *et al.*, “Summary on the ICASSP 2022 multi-channel multi-party meeting transcription grand challenge,” in *ICASSP*. IEEE, 2022, pp. 9156–9160.
- [16] A. Gulati, J. Qin, C. Chiu, N. Parmar *et al.*, “Conformer: Convolution-augmented transformer for speech recognition,” in *INTERSPEECH*. ISCA, 2020, pp. 5036–5040.
- [17] T. Zhou, Y. Zhao, and J. Wu, “Resnext and res2net structures for speaker verification,” in *SLT*. IEEE, 2021, pp. 301–307.
- [18] Z. Du, S. Zhang, S. Zheng, and Z. Yan, “Speaker overlap-aware neural diarization for multi-party meeting analysis,” in *EMNLP*. Association for Computational Linguistics, 2022, pp. 7458–7469.
- [19] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *ICML*, ser. ACM International Conference Proceeding Series, vol. 148. ACM, 2006, pp. 369–376.
- [20] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, “Hybrid ctc/attention architecture for end-to-end speech recognition,” *IEEE J. Sel. Top. Signal Process.*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [21] W. Huang and J. Feng, “Differential beamforming for uniform circular array with directional microphones,” in *INTERSPEECH*. ISCA, 2020, pp. 71–75.
- [22] S. Zheng, W. Huang, X. Wang, H. Suo, J. Feng, and Z. Yan, “A real-time speaker diarization system based on spatial spectrum,” in *ICASSP*. IEEE, 2021, pp. 7208–7212.
- [23] Y. Fan, J. W. Kang, L. T. Li, K. C. Li, H. L. Chen, S. T. Cheng, P. Y. Zhang, Z. Y. Zhou, Y. Q. Cai, and D. Wang, “Cn-celeb: A challenging chinese speaker recognition dataset,” in *ICASSP*. IEEE, 2020, pp. 7604–7608.
- [24] Y. Fu, L. Cheng, S. Lv, Y. Jv, Y. Kong *et al.*, “AISHELL-4: an open source dataset for speech enhancement, separation, recognition and speaker diarization in conference scenario,” in *INTERSPEECH*. ISCA, 2021, pp. 3665–3669.
- [25] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba *et al.*, “Espnet: End-to-end speech processing toolkit,” in *INTERSPEECH*. ISCA, 2018, pp. 2207–2211.
- [26] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR (Poster)*, 2015.