



Streaming Speaker-Attributed ASR with Token-Level Speaker Embeddings

Naoyuki Kanda¹, Jian Wu¹, Yu Wu², Xiong Xiao¹, Zhong Meng¹, Xiaofei Wang¹, Yashesh Gaur¹,
Zhuo Chen¹, Jinyu Li¹, Takuya Yoshioka¹

¹Microsoft Cloud+AI, USA ²Microsoft Research Asia, China

{nakanda, wujian, yuwu1, xioxiao, zhme, xiaofewa, yagaur, zhuc, jinyuli, tayoshio}@microsoft.com

Abstract

This paper presents a streaming speaker-attributed automatic speech recognition (SA-ASR) model that can recognize “who spoke what” with low latency even when multiple people are speaking simultaneously. Our model is based on token-level serialized output training (t-SOT) which was recently proposed to transcribe multi-talker speech in a streaming fashion. To further recognize speaker identities, we propose an encoder-decoder based speaker embedding extractor that can estimate a speaker representation for each recognized token not only from non-overlapping speech but also from overlapping speech. The proposed speaker embedding, named t-vector, is extracted synchronously with the t-SOT ASR model, enabling joint execution of speaker identification (SID) or speaker diarization (SD) with the multi-talker transcription with low latency. We evaluate the proposed model for a joint task of ASR and SID/SD by using LibriSpeechMix and LibriCSS corpora. The proposed model achieves substantially better accuracy than a prior streaming model and shows comparable or sometimes even superior results to the state-of-the-art offline SA-ASR model.

Index Terms: multi-talker speech recognition, serialized output training, speaker identification, speaker diarization

1. Introduction

Speaker-attributed automatic speech recognition (SA-ASR), a task to recognize “who spoke what” from audio input, has long been studied for conversation analysis [1, 2, 3, 4, 5]. An SA-ASR system typically consists of multiple modules such as speech separation for handling overlapping speech [6], speaker identification (SID) [7] or speaker diarization (SD) [8] for estimating the speaker identity, and ASR for transcribing each utterance [9]. While substantial advancement has been made for SA-ASR systems, such a modular system tends to be complicated, and difficult to optimize for the best accuracy.

To further improve the SA-ASR accuracy, various attempts have been made to jointly optimize multiple modules either with partially joint approaches (e.g. speech separation and ASR [10, 11, 12, 13, 14, 15]) or fully joint approaches [16, 17, 18]. In particular, the end-to-end (E2E) SA-ASR model of [19] which jointly performs multi-talker ASR and SID was shown to considerably outperform the modular-approach-based systems [20]. However, the E2E SA-ASR model is based on the attention encoder-decoder architecture [21, 22, 23] with serialized output training (SOT)-based ASR [24], which renders the model usable only for offline (i.e. non-streaming) inference.

Only a limited number of studies addressed the streaming SA-ASR problem through the joint optimization approach. Shafey et al. [16] proposed to insert speaker role tags (e.g., “doctor” and “patient”) in output transcriptions for the offline recurrent neural network transducer (RNNT)-based ASR. Saltau et al. [25] later showed a good performance with this idea

for medical conversation transcription with streaming RNNT models. However, their model is difficult to extend to general scenarios with arbitrary number of speakers since it requires each speaker to have a unique role. Their model also cannot deal with overlapping speech which is ubiquitous in human conversations [26]. Lu et al. [27] proposed streaming unmixing, recognition and identification transducer (SURIT) that jointly performs multi-talker ASR and SID in a streaming fashion. Their model has two output branches for ASR and two additional output branches for SID to generate ASR and SID results for up to two simultaneous speakers with limited latency. However, their word error rate (WER) and speaker error rate (SER) were significantly worse than the state-of-the-art (SOTA) results of the offline SA-ASR models [28].

To address these limitations, we propose a novel streaming SA-ASR model that works with low latency even for overlapping speech. Our model is based on token-level serialized output training (t-SOT) [29], which was recently proposed for low-latency multi-talker speech transcription. To further estimate the speaker identities, we propose an encoder-decoder based speaker embedding extractor that can estimate a speaker representation for each recognized token not only from non-overlapping speech but also from overlapping speech. The proposed speaker embedding, named t-vector, is extracted synchronously with the t-SOT ASR to jointly perform SID/SD and the multi-talker transcription for arbitrary number of speakers with low latency. In our evaluation using LibriSpeechMix [19] and LibriCSS [30] corpora, we show that the proposed model achieves substantially better accuracy than the prior streaming SA-ASR model. The model also yields on par or sometimes even better accuracy than the SOTA offline SA-ASR model.

2. Streaming Multi-Talker ASR with t-SOT

2.1. t-SOT

The t-SOT framework was recently proposed to recognize multi-talker conversations with low latency [29]. In t-SOT, we assume up to M utterances are overlapping at the same time in the input audio. Here, we explain the t-SOT for $M = 2$. Refer [29] for a case with $M > 2$. In t-SOT, the transcriptions for multiple speakers are serialized into a single sequence of recognition tokens (e.g., words, subwords) by sorting the tokens in a chronological order. A special token $\langle cc \rangle$, which indicates a change of “virtual” output channels, is inserted between two adjacent words if they are spoken by two different speakers. An example of such a serialized transcription is shown at the middle of Fig. 1. A streaming end-to-end ASR model is trained to generate such serialized transcriptions from the corresponding audio samples. During inference, a serialized transcription including $\langle cc \rangle$ is produced by the ASR model in a streaming fashion, which is then deserialized to separate transcriptions by switching the virtual output channel when $\langle cc \rangle$ is encountered.

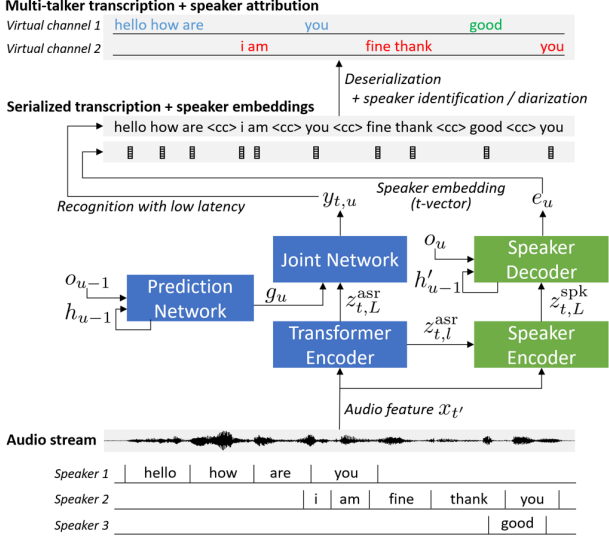


Figure 1: The overview of t -SOT-based TT with t -vector extraction. ASR-related blocks are colored by blue while speaker-related blocks are colored by green.

The deseriaization process is illustrated at the top of Fig. 1. The t -SOT model was found to significantly outperform prior multi-talker ASR models in both the recognition accuracy and latency while keeping the model architecture as simple as conventional single-talker ASR models [29].

2.2. Transformer Transducer

We use transformer transducer (TT) [31] as a backbone ASR model in this work. The components of TT are shown by blue blocks in Fig. 1. TT is an RNN model using Transformer [32] as the encoder. The TT encoder is represented as follows.

$$z_{[1:T],0}^{\text{asr}} = \text{CNN}(x_{[1:T']}) \quad (1)$$

$$\tilde{z}_{t,l}^{\text{asr}} = z_{t,l-1}^{\text{asr}} + \text{MHA}_l^{\text{asr}}(z_{t,l-1}^{\text{asr}}, z_{m(t),l-1}^{\text{asr}}, z_{m(t),l-1}^{\text{asr}}), \quad (2)$$

$$z_{t,l}^{\text{asr}} = \tilde{z}_{t,l}^{\text{asr}} + \text{FF}_l^{\text{asr}}(\tilde{z}_{t,l}^{\text{asr}}) \quad (3)$$

Here, $x_{[1:T']}$ is a sequence of input audio features with length T' . $\text{CNN}()$ represents a stack of convolution layers that subsamples the sequence length to T . $\text{MHA}_l^*(Q, K, V)$ represents a multi-head attention (MHA) of the l -th layer [32] with query Q , key K , and value V matrices. The streaming mask function $m(t)$ generates limited time frames $[t_s : t_e] \in [1 : T]$ to compute the attention at time frame t by using an attention mask. With an appropriate attention mask, TT can be executed with limited algorithmic latency as proposed in [33]. Finally, $\text{FF}_l^*(\cdot)$ represents a position-wise feed forward network at the l -th layer. The decoder of TT consists of the prediction network and joint network as follows.

$$g_u, h_u = \text{Prediction}(o_{u-1}, h_{u-1}) \quad (4)$$

$$y_{t,u} = \text{Joint}(z_{t,L}^{\text{asr}}, g_u) \quad (5)$$

Here, o_u is the u -th estimated token, and h_u is a hidden representation of the prediction network at step u . L is the number of layers in the TT encoder. The resultant $y_{t,u}$ represents a probability distribution of the tokens at time t given $o_{[1:u-1]}$. In inference, the token o_u is estimated by performing time-synchronous beam search on $y_{t,u}$.

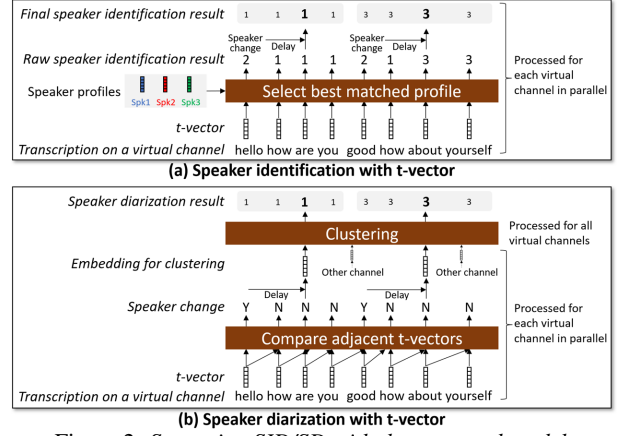


Figure 2: Streaming SID/SD with the proposed model.

3. Streaming SA-ASR with t-vector

3.1. t-vector: Token-level speaker embedding

The key challenge of SID/SD with streaming multi-talker ASR is estimating clean speaker representations *even from overlapping speech* in a way synchronous with the ASR output tokens *whose emission times can have arbitrary delays* from their actual spoken times in the input audio. To overcome this challenge, we propose an encoder-decoder based speaker embedding extraction method that works synchronously with the emission of token o_u from TT.

The proposed speaker embedding extraction model consists of a speaker encoder and a speaker decoder as shown by the green blocks in Fig. 1. The speaker encoder has the same number of layers (i.e. L layers) with the transformer encoder of TT. It cooperatively works with the TT encoder and generates raw speaker representation $z_{t,L}^{\text{spk}}$ from the final (i.e. L -th) layer for each time frame t . Specifically, the proposed speaker encoder is formulated as follows.

$$z_{[1:T],0}^{\text{spk}} = \text{DvecExtractorNet}(x_{[1:T']}) \quad (6)$$

$$z_{t,l}^{\text{spk}} = z_{t,l-1}^{\text{spk}} + \text{MHA}_l^{\text{spk}}(z_{t,l-1}^{\text{asr}}, z_{m(t),l-1}^{\text{asr}}, z_{m(t),l-1}^{\text{spk}}). \quad (7)$$

Here, $\text{DvecExtractorNet}()$ is a neural network with almost the same architecture as a d-vector extractor [34] (Res2Net [35, 36] in our implementation) except that all look-ahead weights in the convolution layers are zeroed out to prevent the speaker encoder from adding further latency. It generates the speaker representation $z_{[1:T],0}^{\text{spk}}$ at the bottom (i.e. 0-th) layer. For each layer l , MHA is applied by using the query $z_{t,l-1}^{\text{asr}}$ and key $z_{m(t),l-1}^{\text{asr}}$ from TT while using the value $z_{m(t),l-1}^{\text{spk}}$ inside the speaker encoder. The masking function $m()$ is shared with TT such that the speaker encoder works with limited latency.

The speaker decoder generates speaker embedding e_u , named t-vector, for each non-blank token, i.e., o_u as follows.

$$e_u, h'_u = \text{LSTM}(z_{t_u,L}^{\text{spk}} + \text{Embed}(o_u), h'_{u-1}), \quad (8)$$

where t_u is the emission time frame of o_u based on TT, $\text{Embed}()$ is the embedding function, and h'_u is the hidden state of the long short-term memory (LSTM) at step u .

For training the speaker encoder and decoder, we use a simple softmax loss based on speaker identification using the cosine similarity for e_u as

$$\mathcal{L}^{\text{spk}} = \sum_{u|o_u \neq \langle \text{cc} \rangle} -\log \frac{e^{\cos(e_u, d_u)}}{e^{\cos(e_u, d_u)} + \sum_{d' \in \Phi(\mathcal{D})} e^{\cos(e_u, d')}} \quad (9)$$

Table 1: Evaluation result for LibriSpeechMix with ASR+SID task given 8 speaker profiles, each of which was extracted from 2 utterances. The average duration for SID latency is calculated based on the reference alignment.

Model name	Model type	# of param.	Latency on ASR	Latency on SID (Avg. duration)	1-speaker			2-speaker-mixed		
					SER	WER	SAWER	SER	WER	SAWER
SOT LSTM SA-ASR [19]	Non-streaming	146M	∞	∞	0.2	4.2	4.5	2.5	8.7	9.9
SOT Conformer SA-ASR [28]	Non-streaming	142M	∞	∞	0.6	3.3	3.9	2.0	4.3	6.4
SURIT [27]	Streaming	90M	0.15 sec	$0.85T^\dagger$ (8.25 sec)	-	-	-	6.7^\ddagger	10.1	-
t-SOT TT-18 w/ t-vector	Streaming	100M	0.16 sec	+ 2 words (0.68 sec)	2.2	5.5	5.9	3.8	7.9	8.5
t-SOT TT-18 w/ t-vector	Streaming	100M	0.16 sec	+ 4 words (1.36 sec)	1.0	5.2	5.5	2.0	7.3	8.0
t-SOT TT-18 w/ t-vector	Streaming	100M	0.16 sec	+ 8 words (2.72 sec)	0.6	5.0	5.4	1.2	7.0	7.8
t-SOT TT-36 w/ t-vector	Streaming	160M	0.16 sec	+ 8 words (2.72 sec)	0.7	4.4	4.9	1.5	6.5	7.4
t-SOT TT-36 w/ t-vector	Streaming	160M	2.56 sec	+ 8 words (2.72 sec)	0.7	3.3	4.0	1.5	4.7	5.8

$^\dagger T$ is the length of the audio segment whose average duration was 9.70 sec for the test set. ‡ Each speaker profile was extracted by using 10 utterances.

where d_u is a reference d-vector for the speaker of token o_u . $\Phi(\mathcal{D})$ is a random subset of the d-vectors \mathcal{D} in the training data, where $\Phi(\mathcal{D})$ does not include the d-vector of the speaker of o_u . In training, the value of t_u , which is necessary to extract e_u , is determined by Viterbi time-alignment of the reference tokens based on TT. Note that \mathcal{L}^{spk} is computed only for normal tokens except for $\langle cc \rangle$ because e_u for $\langle cc \rangle$ is not used for inference.

3.2. Streaming ASR+SID with t-vector

In the context of SA-ASR, SID is a task to determine the speaker of each word given a set of speaker profiles (i.e. d-vectors) of candidate speakers. The proposed procedure is shown in the upper diagram of Fig. 2. Here, for each virtual output channel, a t-vector for each word¹ is compared with each speaker profile based on the cosine similarity, and the index of the speaker profile that has the largest cosine similarity is selected as a raw SID result of the word. A speaker change is then detected based on the change of the raw SID result. A final SID result is determined after a delay of a fixed number of words (e.g., 2 words in Fig. 2) from the speaker change point. Note that, once a speaker change is detected, the next speaker change is not detected until the final SID result is determined as exemplified in Fig. 2 where speaker change is not detected at “how” in the second segment.

The delay parameter mentioned above is introduced because it usually takes a few iterations for the speaker decoder to reliably estimate the speaker representation. On the other hand, the longer the delay is, the higher the chance of missing speaker change points becomes. The impact of the delay will be discussed in the experiment section.

3.3. Streaming ASR+SD with t-vector

In the context of SA-ASR, SD is a task to assign a speaker cluster index for each spoken word without having speaker profiles. The proposed procedure is shown in the lower diagram of Fig. 2. Here, for each virtual output channel, we first detect the speaker change point if the cosine similarity of adjacent t-vectors is smaller than a pre-determined threshold. The t-vector after a fixed number of words from the speaker change point is used as the speaker embedding for the segment until the next speaker change is detected. Every time a new speaker embedding is generated, a clustering procedure is applied to the speaker embeddings estimated for all the output channels up to the current time. This allows an intermediate SD result to be presented with low latency, where the algorithmic latency is determined by the delay in finalizing the speaker embedding for the segment.²

¹When one word consists of multiple subwords, the t-vector corresponding to the last subword in the word is used.

²An idea to repeatedly apply clustering was proposed in [37].

4. Experiments

We first conducted a preliminary experiment for ASR and SID task by using LibriSpeechMix [24], where the number of speakers in each audio is limited to two. We then extended the experiment by using LibriCSS [30] to evaluate the SA-ASR performance, with both SID and SD, for long-form audio containing many utterances.

4.1. Experiment with LibriSpeechMix

4.1.1. Experimental settings

We used the single speaker and two speaker subsets of LibriSpeechMix [19] as a preliminary experiment for ASR and SID task. The dataset is made by mixing up to two utterances randomly sampled from LibriSpeech [41] with a randomly-determined delay. For each test sample, an SA-ASR model/system must generate a multi-talker transcription while identifying the speaker of each word from 8 candidate speakers given their speech profiles, which were extracted by ResNet-based d-vector [35, 36]. We used SER, WER and speaker-attributed WER (SAWER), as defined in [19], as our evaluation metrics. Here, SAWER is our primary metric, where the estimated word is regarded as correct if and only if an estimated word and speaker both match the reference.

For the ASR block, we used a 18-layer or 36-layer TT (TT-18 or TT-36 in short) with the chunk-wise look-ahead proposed in [33] by using exactly the same configuration as in [29]. Each transformer block consisted of a 512-dim MHA with 8 heads and a 2048-dim point-wise feedforward layer. The prediction network consisted of 2 layers of 1024-dim LSTM. 4,000 word pieces plus blank and $\langle cc \rangle$ tokens were used as the recognition units. The input feature is an 80-dim log mel-filterbank extracted every 10 msec. Refer to [29] for further details on the ASR block. For the proposed speaker embedding extraction, Res2Net [35, 36] was used for DvecExtractorNet and a 128-dim 8-head MHA was used at each layer of the speaker encoder. The speaker decoder consists of a 2-layer 512-dim LSTM. As proposed in [33], we controlled the algorithmic latency of the model based on the chunk size of the attention mask.

For the model training, we simulated the training data by randomly mixing up to two utterances from “train_960” of LibriSpeech by following the procedure in [29]. For each training sample, we randomly selected up to 8 candidate speakers to calculate \mathcal{L}^{spk} . We initialized the ASR block by using the parameters of the t-SOT ASR trained in our prior work [29], and the ASR block was frozen during the training. We performed 60K training iterations with 16 GPUs, each of which consumed mini-batches of 12,000 frames. We used an AdamW optimizer with a linear decay learning rate schedule with a peak learning rate of 1.5e-3 after 10K warm up iterations.

Table 2: *cpWER (%) for LibriCSS with ASR+SID and ASR+SD tasks. Algorithmic latency is shown in the latency columns. 0L and 0S are 0% overlap conditions with long and short inter-utterance silences.*

(a) ASR+SID task given 8 speaker profiles									
System	Latency on ASR	Latency on SID (avg. duration)	cpWER (%) for different overlap ratio						Total
			0L	0S	10	20	30	40	
SOT LSTM SA-ASR [38]	∞^*	∞^*	8.0	15.7	12.5	17.5	24.3	27.6	18.6
SOT Conformer SA-ASR [28]	∞^*	∞^*	7.9	12.1	9.6	10.8	13.4	15.7	11.9
t-SOT TT-18 w/ t-vector	0.16 sec	+ 2 words (0.70 sec)	10.3	11.8	12.1	14.6	19.5	20.2	15.3
t-SOT TT-36 w/ t-vector	0.16 sec	+ 2 words (0.70 sec)	10.0	11.2	11.6	14.7	18.2	19.1	14.6
t-SOT TT-36 w/ t-vector	2.56 sec	+ 2 words (0.70 sec)	6.6	11.2	9.9	11.6	14.4	15.5	12.0

* Latency is determined by the segment length based on VAD. The average length of speech segments in LibriCSS is 13.10 sec.

(b) ASR+SD task given the oracle number of speakers									
System	Latency on ASR	Latency on SD (avg. duration)	cpWER (%) for different overlap ratio.						Total
			0L	0S	10	20	30	40	
TS-VAD \rightarrow Transformer ASR [39]	∞^\diamond	∞^\diamond	9.5	11.0	16.1	23.1	33.8	40.9	23.9
SOT LSTM SA-ASR \rightarrow NME-SC [38]	∞^\diamond	∞^\diamond	10.3	15.8	13.4	17.1	24.4	28.6	19.2
SOT Conformer SA-ASR \rightarrow NME-SC [28]	∞^\diamond	∞^\diamond	8.6	12.7	11.2	11.3	16.1	17.5	13.3
NME-SC \rightarrow SOT Conformer SA-ASR [40]	∞^\diamond	∞^\diamond	9.0	12.2	8.7	10.9	13.7	13.9	11.6
t-SOT TT-18 w/ t-vector	0.16 sec	+ 2 words (0.70 sec)	9.0	13.5	12.1	15.4	19.4	21.4	15.8
t-SOT TT-36 w/ t-vector	0.16 sec	+ 2 words (0.70 sec)	9.1	10.8	11.9	14.2	18.3	20.4	14.7
t-SOT TT-36 w/ t-vector	2.56 sec	+ 2 words (0.70 sec)	6.8	10.2	10.5	11.6	15.4	16.7	12.4

\diamond Latency is determined by the length of the recording. The average length of the LibriCSS recordings is 10 min.

Table 3: *cpWER (%) for LibriCSS based on SA-TT-36 (2.56 sec ASR latency) with different delayed decisions on SID/SD.*

Model	Task	# of delayed words on SID/SD				
		0	1	2	4	8
t-SOT TT-36 w/ t-vector	ASR+SID	12.6	12.1	12.0	12.2	15.5
t-SOT TT-36 w/ t-vector	ASR+SD	13.1	12.7	12.4	14.3	19.7

4.1.2. Evaluation results

The evaluation results are shown in Table 1. We first observed that the proposed t-SOT TT-18 with t-vector-based SID significantly outperformed the prior streaming model, SURIT [27], on both SER and WER even with a SID latency of only 2 words (0.68 sec on average). With a longer SID delay of 8 words, which is still faster than SURIT, all metrics were further improved. We then observed that both the model size and the latency on ASR had a significant impact on WER and SAWER. The model based on TT-36 with 2.56 sec of ASR latency even achieved a comparable SAWER for the single speaker test set and a better SAWER for the two-speaker-mixed test set than the prior offline models.

4.2. Experiment with LibriCSS

4.2.1. Experimental settings

To evaluate the proposed model in a more realistic setting, we conducted an evaluation with LibriCSS [30]. LibriCSS is a set of 10-min recordings (10 hours in total), created by playing back utterances of LibriSpeech “test_clean” in a real meeting room. The original recording was made with a 7-ch microphone array, and the first channel of the recording was used in our experiment. Each 10-min recording are made of utterances from 8 speakers and categorized by the speaker overlap ratio from 0% to 40%. We evaluated the proposed model based on concatenated minimum-permutation WER (cpWER) [5], which is affected by both ASR and speaker attribution errors.

For training the model, we simulated the training data by following the procedure of [29], where each training audio consists of up to 5 utterances randomly sampled from LibriSpeech “train_960”. Up to 8 candidate speakers were randomly selected from LibriSpeech “train_960” to compute \mathcal{L}^{spk} . We initialized the ASR block with the parameters of t-SOT TT trained for LibriCSS in [29], and the speaker block with the parameters of the model trained for LibriSpeechMix. The ASR block was frozen

during the training using \mathcal{L}^{spk} . We performed 25,000 training iterations with 16 GPUs, each of which consumed a mini-batch of 12,000 frames. The AdamW optimizer was used with a linear decay learning rate schedule with a peak learning rate of $1.5\text{e-}4$.

In the inference, we stopped the beam search every time a silence was detected by WebRTC Voice Activity Detector [42] after 20 seconds of decoding. We forcibly terminated the beam search when a silence was not detected after 40 seconds of decoding. For the ASR and SID task, we supplied 8 speaker profiles corresponding to the speakers in the recordings, where each speaker profile is extracted by using 5 utterances of the speaker. For the ASR and SD task, we used the normalized maximum eigengap-based spectral clustering (NME-SC) [43] given the oracle number of speakers. The threshold for detecting the speaker change was set to 0.98.

4.2.2. Evaluation results

Table 2 shows the main result. Note that, this is the first time that a fully streaming SA-ASR system was applied to LibriCSS, and thus we don’t have the numbers for the prior streaming models in the table. The proposed t-SOT TT model with t-vector achieved cpWERs of 12.0% and 12.4% for ASR+SID and ASR+SD, respectively, which are both very close to the SOTA results of the offline SA-ASR model.

Table 3 shows the effect of the delayed decisions on SID/SD. As discussed in Section 3.2, a longer delay stabilizes the speaker embedding at the cost of increased miss errors of speaker change points. Unlike LibriSpeechMix which contains up to 2 utterances, the drawback of missing speaker change points had a larger impact for LibriCSS, resulting in a significant degradation of cpWER with the 4- and 8-word delays.

5. Conclusion

This paper presented a new streaming SA-ASR model. We proposed an encoder-decoder based speaker embedding extractor that synchronously works with the t-SOT ASR, and showed a simple streaming SID/SD procedure with estimated speaker embeddings, t-vector. In our evaluation, we showed that the proposed model achieved comparable or sometimes even superior SA-ASR accuracy to the SOTA offline models while enabling low-latency inference.

6. References

- [1] A. Janin *et al.*, “The ICSI meeting corpus,” in *ICASSP*, vol. I, 2003, pp. 364–367.
- [2] J. Carletta *et al.*, “The AMI meeting corpus: A pre-announcement,” in *International workshop on machine learning for multimodal interaction*, 2006, pp. 28–39.
- [3] J. G. Fiscus, J. Ajot, and J. S. Garofolo, “The rich transcription 2007 meeting recognition evaluation,” in *Multimodal Technologies for Perception of Humans*, 2007, pp. 373–389.
- [4] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, “The fifth ‘CHiME’ speech separation and recognition challenge: Dataset, task and baselines,” *Interspeech*, pp. 1561–1565, 2018.
- [5] S. Watanabe *et al.*, “CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings,” in *CHiME 2020*, 2020.
- [6] S. Chen, Y. Wu, Z. Chen, J. Wu, J. Li, T. Yoshioka, C. Wang, S. Liu, and M. Zhou, “Continuous speech separation with Conformer,” in *ICASSP*, 2021, pp. 5749–5753.
- [7] Z. Bai and X.-L. Zhang, “Speaker recognition based on deep learning: An overview,” *Neural Networks*, vol. 140, pp. 65–99, 2021.
- [8] T. J. Park *et al.*, “A review of speaker diarization: Recent advances with deep learning,” *arXiv:2101.09624*, 2021.
- [9] T. Yoshioka *et al.*, “Advances in online audio-visual meeting transcription,” in *ASRU*, 2019, pp. 276–283.
- [10] D. Yu, X. Chang, and Y. Qian, “Recognizing multi-talker speech with permutation invariant training,” *Interspeech*, pp. 2456–2460, 2017.
- [11] H. Seki, T. Hori, S. Watanabe, J. Le Roux, and J. R. Hershey, “A purely end-to-end system for multi-speaker speech recognition,” in *ACL*, 2018, pp. 2620–2630.
- [12] X. Chang, Y. Qian, K. Yu, and S. Watanabe, “End-to-end monaural multi-speaker ASR system without pretraining,” in *ICASSP*, 2019, pp. 6256–6260.
- [13] 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*, 2019, pp. 237–244.
- [14] N. Kanda, Y. Fujita, S. Horiguchi, R. Ikeshita, K. Nagamatsu, and S. Watanabe, “Acoustic modeling for distant multi-talker speech recognition with single-and multi-channel branches,” in *ICASSP*, 2019, pp. 6630–6634.
- [15] N. Kanda, S. Horiguchi, R. Takashima, Y. Fujita, K. Nagamatsu, and S. Watanabe, “Auxiliary interference speaker loss for target-speaker speech recognition,” in *Interspeech*, 2019, pp. 236–240.
- [16] L. El Shafey, H. Soltan, and I. Shafran, “Joint speech recognition and speaker diarization via sequence transduction,” in *Interspeech*, 2019, pp. 396–400.
- [17] N. Kanda, S. Horiguchi, Y. Fujita, Y. Xue, K. Nagamatsu, and S. Watanabe, “Simultaneous speech recognition and speaker diarization for monaural dialogue recordings with target-speaker acoustic models,” in *ASRU*, 2019.
- [18] H. H. Mao, S. Li, J. McAuley, and G. W. Cottrell, “Speech recognition and multi-speaker diarization of long conversations,” in *Interspeech*, 2020, pp. 691–695.
- [19] N. Kanda *et al.*, “Joint speaker counting, speech recognition, and speaker identification for overlapped speech of any number of speakers,” in *Interspeech*, 2020, pp. 36–40.
- [20] N. Kanda, X. Xiao, J. Wu, T. Zhou, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, “A comparative study of modular and joint approaches for speaker-attributed asr on monaural long-form audio,” in *ASRU*, 2021.
- [21] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [22] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio, “End-to-end continuous speech recognition using attention-based recurrent NN: First results,” in *NIPS Workshop on Deep Learning*, 2014.
- [23] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *NIPS*, 2015, pp. 577–585.
- [24] N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, “Serialized output training for end-to-end overlapped speech recognition,” in *Interspeech*, 2020, pp. 2797–2801.
- [25] H. Soltan, M. Wang, I. Shafran, and L. E. Shafey, “Understanding medical conversations: Rich transcription, confidence scores & information extraction,” in *Proc. Interspeech*, 2021, pp. 4418–4422.
- [26] Ö. Çetin and E. Shriberg, “Analysis of overlaps in meetings by dialog factors, hot spots, speakers, and collection site: Insights for automatic speech recognition,” in *Interspeech*, 2006.
- [27] L. Lu, N. Kanda, J. Li, and Y. Gong, “Streaming multi-talker speech recognition with joint speaker identification,” in *Interspeech*, 2021, pp. 1782–1786.
- [28] N. Kanda *et al.*, “End-to-end speaker-attributed ASR with Transformer,” in *Interspeech*, 2021, pp. 4413–4417.
- [29] 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.
- [30] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, and J. Li, “Continuous speech separation: dataset and analysis,” in *ICASSP*, 2020, pp. 7284–7288.
- [31] Q. Zhang, H. Lu, H. Sak, A. Tripathi, E. McDermott, S. Koo, and S. Kumar, “Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” in *Proc. ICASSP*, 2020, pp. 7829–7833.
- [32] A. Vaswani *et al.*, “Attention is all you need,” in *NIPS*, 2017, pp. 6000–6010.
- [33] X. Chen, Y. Wu, Z. Wang, S. Liu, and J. Li, “Developing real-time streaming transformer transducer for speech recognition on large-scale dataset,” in *Proc. ICASSP*, 2021, pp. 5904–5908.
- [34] E. Variiani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, “Deep neural networks for small footprint text-dependent speaker verification,” in *ICASSP*, 2014, pp. 4052–4056.
- [35] S. Gao, M.-M. Cheng, K. Zhao, X.-Y. Zhang, M.-H. Yang, and P. H. Torr, “Res2net: A new multi-scale backbone architecture,” *IEEE trans. on PAMI*, 2019.
- [36] X. Xiao *et al.*, “Microsoft speaker diarization system for the VoxCeleb speaker recognition challenge 2020,” in *ICASSP*, 2021, pp. 5824–5828.
- [37] D. Dimitriadis and P. Fousek, “Developing on-line speaker diarization system,” in *Proc. Interspeech*, 2017, pp. 2739–2743.
- [38] N. Kanda *et al.*, “Investigation of end-to-end speaker-attributed ASR for continuous multi-talker recordings,” in *SLT*, 2021, pp. 809–816.
- [39] D. Raj *et al.*, “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System description, comparison, and analysis,” in *SLT*, 2021, pp. 897–904.
- [40] N. Kanda, X. Xiao, Y. Gaur, X. Wang, Z. Meng, Z. Chen, and T. Yoshioka, “Transcribe-to-diarize: Neural speaker diarization for unlimited number of speakers using end-to-end speaker-attributed asr,” *arXiv preprint arXiv:2110.03151*, 2021.
- [41] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *ICASSP*, 2015, pp. 5206–5210.
- [42] “WebRTC Voice Activity Detector,” <https://github.com/wiseman/py-webrtcvad>.
- [43] T. J. Park, K. J. Han, M. Kumar, and S. Narayanan, “Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap,” *IEEE Signal Processing Letters*, vol. 27, pp. 381–385, 2019.