



Streaming Target-Speaker ASR with Neural Transducer

Takafumi Moriya^{1,2*}, Hiroshi Sato^{1*}, Tsubasa Ochiai¹, Marc Delcroix¹, Takahiro Shinozaki²

¹NTT Corporation, Japan; ²Tokyo Institute of Technology, Japan

{takafumi.moriya.nd,hiroshi.satou.bh}@hco.ntt.co.jp

Abstract

Although recent advances in deep learning technology have boosted automatic speech recognition (ASR) performance in the single-talker case, it remains difficult to recognize multi-talker speech in which many voices overlap. One conventional approach to tackle this problem is to use a cascade of a speech separation or target speech extraction front-end with an ASR back-end. However, the extra computation costs of the front-end module are a critical barrier to quick response, especially for streaming ASR. In this paper, we propose a target-speaker ASR (TS-ASR) system that implicitly integrates the target speech extraction functionality within a streaming end-to-end (E2E) ASR system, i.e. recurrent neural network-transducer (RNNT). Our system uses a similar idea as adopted for target speech extraction, but implements it directly at the level of the encoder of RNNT. This allows TS-ASR to be realized without placing extra computation costs on the front-end. Note that this study presents two major differences between prior studies on E2E TS-ASR; we investigate streaming models and base our study on Conformer models, whereas prior studies used RNN-based systems and considered only offline processing. We confirm in experiments that our TS-ASR achieves comparable recognition performance with conventional cascade systems in the offline setting, while reducing computation costs and realizing streaming TS-ASR.

Index Terms: target-speaker speech recognition, neural transducer, end-to-end, streaming inference, noise robust

1. Introduction

End-to-end automatic speech recognition (E2E ASR), which can directly map acoustic features to output tokens, has recently attracted increasing interest from the ASR research community. Recent E2E ASR systems have achieved high recognition performance for single speaker conditions [1–8]. However, it remains challenging to recognize speech under severe conditions such as in presence of overlapping speakers [9]. This paper addresses the challenging problem of recognizing a target speaker in a mixture, i.e., target-speaker ASR (TS-ASR), in a streaming manner to support applications such as voice search or user-dependent voice interactive systems.

We can tackle this problem from the viewpoint of separation [10–14] or target speech extraction (TSE) [15–20], using a cascade speech enhancement front-end/ASR back-end system or an E2E system. Speech separation separates a signal into all of its sources [21, 22]. It allows ASR of all speakers in a mixture. However, the problem of global permutation ambiguity means that the mapping between output and speaker is arbitrary. Moreover, decoding the speech of all speakers can be computationally involving. In many applications such as voice search, we are interested in recognizing the speech of a specific user and not the interfering speakers. In this case, TSE is a particularly attractive solution. TSE extracts only the speech of a

target speaker given an enrollment of that speaker’s voice. It naturally avoids the output-speaker ambiguity and can limit the decoding to just the target speaker.

Both separation and TSE approaches can be implemented as a cascade or within an E2E system [12, 16–20, 23–25]. Cascade systems use a separation or TSE front-end that estimates the speech signals and then performs recognition. Cascade systems are modular, making it easy to visualize the different processing steps. However, although they can achieve high performance in offline modes, this comes at the expense of a high computational cost. Moreover, the performance of separation and TSE front-ends degrade severely in streaming mode [13, 17]. In contrast, E2E approaches directly process the mixture and output the recognition results without explicitly estimating the speech signals. They can achieve a similar level of performance at lower computational cost.

Many ASR applications require fast and accurate responses in real-time. Recently, there has been increased interest in developing streaming E2E ASR systems. Recurrent neural network-transducer (RNNT) [26] is a promising technology for streaming E2E ASR applications, and has been extensively investigated for single speaker ASR [2–8]. Several works have extended the ideas of RNNT for multi-talker ASR [27–29]. These works are E2E extensions of separation-based approaches. They recognize the speech of each speaker in a mixture, which is computationally demanding [27, 28]. Moreover, they do not identify the speakers and thus cannot be used directly for TS-ASR [27–29]. To the best of our knowledge, no study has investigated streaming E2E TS-ASR.

In this work, we propose an E2E TS-ASR framework that integrates the TSE functionality within an RNNT. We call this framework target-speaker RNNT (TS-RNNT). The system operates as follows. First, we extract a target speaker embedding from the enrollment using a speaker encoder module. The speaker embedding is used as an auxiliary input to the encoder of RNNT to inform the system which speaker to recognize in the mixture. We can incorporate the embedding within the RNNT encoder with a simple element-wise operation at an intermediate layer. Note that the embedding can be computed in advance. Therefore, our proposed TS-RNNT does not impact the computational complexity nor the latency of a vanilla RNNT model. This study presents two major differences from prior studies about E2E TS-ASR [24, 25]; we investigate streaming models and base our study on the recently proposed state-of-the-art Conformer models [1], whereas prior studies used long short-term memory (LSTM)-based systems and dealt only with offline processing.

We conduct experiments to compare our proposed TS-RNNT with a cascade combination of TSE and RNNT (TSE+RNNT) for offline and streaming modes. Experiments show that our proposed TS-RNNT achieves competitive performance with TSE+RNNT for offline conditions at significantly lower computational cost. More importantly, it greatly outperforms a TSE+RNNT system in streaming mode.

* Equal contribution.

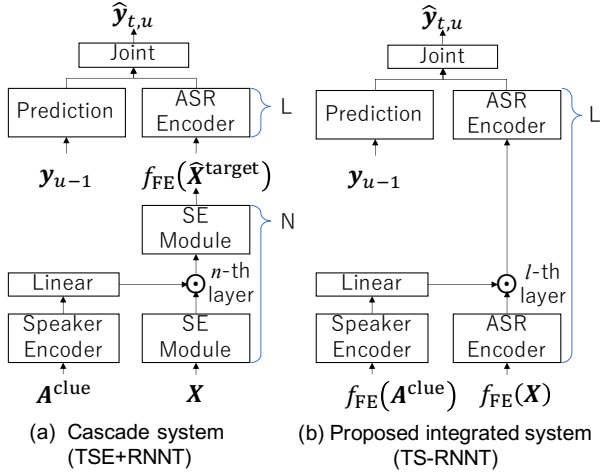


Figure 1: The overview of TS-ASR systems.

2. Target-Speaker ASR

In this section, we introduce the TS-ASR systems used in our experiments. We first explain the TSE front-end and RNNT-based ASR back-end modules, which are the foundations of this work. Then we introduce our proposed TS-RNNT system.

Let $\mathbf{X} = [x_1, \dots, x_{T'}] \in \mathcal{R}^{T'}$ be the single microphone input speech mixture of duration T' that includes the target and interfering speaker's voices and noise. $Y = [y_1, \dots, y_U]$ is the sequence of tokens of length U associated with the utterance spoken by target speaker, where $y_u \in \{1, \dots, K\}$. K is the number of tokens.

2.1. Overview of cascade TS-ASR (TSE+RNNT)

Fig. 1 (a) is a diagram of the cascade TS-ASR system, which is composed of the TSE and ASR modules described below.

2.1.1. TSE front-end

The TSE module extracts the target-speaker's speech $\mathbf{X}^{\text{target}}$ from the mixture \mathbf{X} using an enrollment speech \mathbf{A}^{clue} . We use the time-domain TSE system we developed in [18, 19], which is similar to [30, 31].

First, we extract the embedding of the target-speaker, $\mathbf{h}^{\text{target}}$, from the enrollment speech, \mathbf{A}^{clue} , using a speaker encoder $f^{\text{Spk-Enc}}(\cdot)$ which consists of a multi-layer neural network followed by a linear layer and a time-average pooling layer. Then, we use a speaker extraction network, $f^{\text{SE}}(\cdot)$ to extract the target speech given the embedding, $\mathbf{h}^{\text{target}}$. The above operations, which yield the estimated target speech signal $\hat{\mathbf{X}}^{\text{target}}$, are defined as follows:

$$\mathbf{H}^{\text{target}} = f^{\text{Spk-Enc}}(\mathbf{A}^{\text{clue}}; \theta^{\text{Spk-Enc}}), \quad (1)$$

$$\mathbf{h}^{\text{target}} = \frac{1}{T'} \sum_{t'=1}^{T'} \mathbf{h}_{t'}^{\text{target}}, \quad (2)$$

$$\hat{\mathbf{X}}^{\text{target}} = f^{\text{SE}}(\mathbf{X}, \mathbf{h}^{\text{target}}; \theta^{\text{SE}}). \quad (3)$$

In this work, we insert the embedding $\mathbf{h}^{\text{target}}$ at the n -th layer of $f^{\text{SE}}(\cdot)$ by performing the Hadamard product between the embedding and the output of that layer, $\mathbf{h}_{t'}^{(n)}$, i.e. $\mathbf{h}_{t'}^{(n)} \odot \mathbf{h}^{\text{target}}$. The learnable parameters $\theta^{\text{TSE}} \triangleq [\theta^{\text{Spk-Enc}}, \theta^{\text{SE}}]$ are jointly optimized by using the source-to-distortion ratio (SDR) loss [13].

2.1.2. ASR back-end

We adopt an RNNT-based ASR back-end [26] that can perform streaming ASR. RNNT learns the mapping between sequences of different lengths. It consists of an ASR encoder and a pre-

dition network, which allows the posterior probabilities to be jointly conditioned on not only the ASR encoder outputs but also on previous predictions. The processing is described as follows.

First, the target speech signal, $\mathbf{X}^{\text{target}}$, is transformed into a sequence of acoustic features, i.e. log Mel-filterbank, using a feature extractor $f^{\text{FE}}(\cdot)$. Then, the features are encoded into length- T sequence, $\mathbf{H}^{\text{ASR}} = [\mathbf{h}_1^{\text{ASR}}, \dots, \mathbf{h}_T^{\text{ASR}}]$, via an ASR encoder network $f^{\text{ASR-Enc}}(\cdot)$. Next, the tokens, Y , are also encoded into $\mathbf{H}^{\text{Pred}} = [\mathbf{h}_1^{\text{Pred}}, \dots, \mathbf{h}_U^{\text{Pred}}]$ via a prediction network $f^{\text{Pred}}(\cdot)$. These two encoded features are fed to a feed-forward network, $f^{\text{Joint}}(\cdot)$, to compute the token posterior probabilities, $\hat{y}_{t,u}$. The above operation can be expressed as follows:

$$\mathbf{h}_t^{\text{ASR}} = f^{\text{ASR-Enc}}(f^{\text{FE}}(x_{t'}^{\text{target}}); \theta^{\text{ASR-Enc}}), \quad (4)$$

$$\mathbf{h}_u^{\text{Pred}} = f^{\text{Pred}}(y_{u-1}; \theta^{\text{Pred}}), \quad (5)$$

$$\hat{y}_{t,u} = \text{Softmax}\left(f^{\text{Joint}}(\mathbf{h}_t^{\text{ASR}}, \mathbf{h}_u^{\text{Pred}}; \theta^{\text{Joint}})\right), \quad (6)$$

where $\text{Softmax}(\cdot)$ indicates a softmax operation. All the learnable parameters $\theta^{\text{RNNT}} \triangleq [\theta^{\text{ASR-Enc}}, \theta^{\text{Pred}}, \theta^{\text{Joint}}]$ are optimized by using RNNT loss with the forward-backward algorithm [26].

For training, we use the target speech signals, $\mathbf{X}^{\text{target}}$ corrupted with noise to make a robust system. At test time, we use the extracted signal estimated with TSE, $\hat{\mathbf{X}}^{\text{target}}$. Note that we do not optimize the TSE and ASR modules jointly. Joint optimization of TSE front-end and ASR for streaming applications will be part of our future works.

2.2. Overview of integrated TS-ASR with RNNT

In this paper, we propose an integrated modeling approach for TS-ASR using RNNT, which operates in a fully E2E manner, and allows streaming ASR.

2.2.1. Proposed TS-RNNT architecture

Fig. 1 (b) is a schematic diagram of our proposed TS-RNNT system. The architecture is almost the same as that of the vanilla RNNT described in 2.1.2. The difference is that the encoder of the TS-RNNT inputs directly the speech mixture and uses the target speaker embedding to inform which speaker in the mixture to decode. This is performed using a similar processing as the TSE front-end but within the encoder, i.e., a similar speaker encoder module and a fusion mechanism at an intermediate layer using the Hadamard product.

TS-RNNT encoder, $f^{\text{ASR-Enc}'}(\cdot)$, is a modified version of ASR encoder $f^{\text{ASR-Enc}}(\cdot)$ that receives the speaker embedding extracted from speaker encoder $f^{\text{Spk-Enc}}(\cdot)$. $f^{\text{ASR-Enc}'}(\cdot)$ is defined as follows:

$$\mathbf{H}^{\text{target}'} = f^{\text{Spk-Enc}'}(f^{\text{FE}}(\mathbf{A}^{\text{clue}}); \theta^{\text{Spk-Enc}'}), \quad (7)$$

$$\mathbf{h}^{\text{target}'} = \frac{1}{T} \sum_{t=1}^T \mathbf{h}_t^{\text{target}'}, \quad (8)$$

$$\mathbf{h}_t^{\text{ASR}'} = f^{\text{ASR-Enc}'}(f^{\text{FE}}(x_{t'}), \mathbf{h}_t^{\text{target}'}; \theta^{\text{ASR-Enc}'}), \quad (9)$$

where $\mathbf{H}^{\text{target}'}$ with length- T is the speaker encoder outputs given the input \mathbf{A}^{clue} as the enrollment, which is averaged with time axis and embedded into $\mathbf{h}^{\text{target}'}$. $\mathbf{h}_t^{\text{target}'}$ and l -th layer intermediate output $\mathbf{h}_t^{(l)}$ of the ASR encoder are multiplied as a Hadamard product, i.e. $\mathbf{h}_t^{(l)} \odot \mathbf{h}_t^{\text{target}'}$, in $f^{\text{ASR-Enc}'}(\cdot)$. Therefore, the prediction and joint networks are the same as the vanilla RNNT in 2.1.2. All networks with learnable parameters $\theta^{\text{TS-RNNT}} \triangleq [\theta^{\text{Spk-Enc}'}, \theta^{\text{ASR-Enc}'}, \theta^{\text{Pred}}, \theta^{\text{Joint}}]$ are jointly opti-

Table 1: *Data generation setup. The number of utterances in (a) is double of (b) due to single-speaker case of (b).*

	dataset	mixture type	SIR [dB]	SNR [dB]	#speakers (train set / dev set)	#mixtures or utterances (train set / dev set)
(a)	training data for ASR backend (RNNT)	1 speaker and noise	-	0 - 20	3054 / 160	400,000 / 10,000
(b)	training data for TSE and TS-RNNT	2 speakers and noise	-5 - 5	0 - 20	3054 / 160	200,000 / 5,000
(c)	evaluation data	2 speakers and noise	-5 - 5	0, 5, 10, 15, 20	30	$6,000 \times 5 = 30,000$

mized by using RNNT loss.

When decoding with the TS-RNNT, we register the speaker embedding, $\mathbf{h}^{\text{target}}$, extracted from the enrollment with the speaker encoder in advance. Then, the acoustic features of the mixture signal, $f^{\text{FE}}(x_t)$, is directly input to ASR encoder given also $\mathbf{h}^{\text{target}}$, and the ASR decoder yields the ASR results, \hat{Y} . Therefore, TS-RNNT can perform TS-ASR faster than the cascade system. Moreover, TS-RNNT does not need clean target-speaker’s speech as reference for TSE module training, only the target-speaker’s transcription for TS-RNNT training, which arguably can be easier to collect for real recordings.

2.2.2. Streaming target-speaker ASR by TS-RNNT

In order to change TS-RNNT into a streaming model, we replace the ASR encoder of TS-RNNT with left-to-right encoding module as in [2, 5]. Note that the speaker encoder module is not changed to streaming mode because it only needs to operate once to register the target-speaker’s clues $\mathbf{h}^{\text{target}}$ in advance. The performance of current TSE front-ends tend to degrade severely when operating in streaming mode. On the other-hand, ASR models can model well speech sequences with streaming models. By including the TSE functionality within the encoder of the RNNT, we expect to improve performance over cascade models for streaming mode. In addition, the tuning of the cascade system is more involving than that of the integrated system because we must separately tune the trade-offs of each module, i.e. TSE and ASR, between recognition performance and the latency. On the other hand, the integrated modeling, i.e. TS-RNNT, needs to tune just the ASR encoder to address the trade-off. Thus TS-RNNT would be easier to tune than cascade systems.

3. Experiments

3.1. Experimental setups

3.1.1. Data

We evaluated our proposal on speech recordings simulated using speech from the Corpus of Spontaneous Japanese (CSJ) [32], sampled at 16 kHz. It consists of two-speaker mixtures with background noise taken from the CHiME-4 corpus [33] at signal-to-noise ratio (SNR) between 0 and 20 dB.

We used almost the same data as in previous works [18, 19]. The details are shown in Table 1. The total training data amounts for 800 hours of speech. The overlap ratio of both training and evaluation datasets was about 89% on average. The speakers between training/development and evaluation datasets did not duplicate. We evaluated performance in terms of character error rate (CER) due to the ambiguity of Japanese word boundaries. We also measured inference speed in terms of real time factor (RTF: *decoding time/data time*). We measured the RTF using a Python implementation of the algorithm running on an Intel Xeon 2.40GHz CPU.

3.1.2. System configuration of TSE module

We adopted a time-domain SpeakerBeam structure as an TSE module [18, 31]. It uses blocks of stacked 1D-convolution layers as proposed in ConvTasNet [13]. The details of the implementation and training follows our prior work [18]. We also trained a streaming/causal TSE model by changing the convolu-

Table 2: *CER as a function of the layer- l where the fusion between the ASR encoder and speaker encoder output operates. The results are for offline TS-RNNTs. “All” indicates that the fusion was applied to all layers of the encoder.*

Fusion in layer- l	CER [%] on each SNR					
	20dB	15dB	10dB	5dB	0dB	Avg.
1	8.6	9.3	11.4	17.3	32.7	15.8
5	9.8	10.4	12.5	18.0	33.1	16.8
1-5	9.3	9.8	11.6	17.6	32.9	16.4
All	9.7	10.2	12.5	18.9	35.1	17.3

Table 3: *Comparisons of cascade and proposed TS-RNNT systems in terms of CERs and RTFs at each SNR. All models performed offline decoding. Note that RTFs exclude computation time of speaker encoder in TSE and TS-RNNT.*

System	CER [%] on each SNR [dB]						RTF
	20	15	10	5	0	Avg.	
RNNT	76.0	75.7	75.4	75.7	78.0	76.1	0.40
TSE+RNNT	7.9	8.8	11.1	18.3	36.3	16.5	1.22
TS-RNNT	8.6	9.3	11.4	17.3	32.7	15.8	0.40

tion and global layer normalization (LN) to causal convolution and channel-wise LN, respectively. The algorithmic latency of the streaming TSE model is 1.25ms (= 20samples/16k), which is negligible for the latency of ASR decoding. All TSE models were trained with dataset (b) in Table 1.

3.1.3. System configuration of ASR module

The input feature for ASR models was an 80-dimensional log Mel-filterbank. We used SpecAugment [34] during training. In this paper, we adopted 3262 characters as the output tokens. The training and evaluation data were preprocessed following the Kaldi and ESPnet toolkits [35, 36]. The minibatch size was set to 64 in all experiments. In this work, we adopted Conformer-based encoders as detailed below.

We investigated three version of (TS-) RNNT. First, we experimented with an offline system it consists of the same encoder architecture as Conformer (L) [1] with a kernel size of 15. We used two-layer 2D-convolutional neural networks (CNNs) followed by 17 Conformer blocks, with the stride sizes of both max-pooling layers at each layer set to 2. The prediction network had a 768-dimensional uni-directional LSTM layer followed by a 640-dimensional feed-forward layer. The speaker encoder for TS-RNNT had the same architecture as the ASR encoder while the number of blocks was reduced from 17 to 6 due to the memory usage during training.

We compare this system with two streaming systems that use a similar configuration as the offline system except that the encoder is replace with a causal Conformer (Uni-Conformer) [2] and latency-controlled Conformer (LC-Conformer) [5]. All ASR encoders of the streaming systems replaced the depthwise convolution and batch normalization with the causal equivalent and LN, respectively, and were trained using an attention mask as in [5]. The latency of Uni-Conformer with infinite history and few look-ahead frames, i.e. CNN module, was 30ms. The left and current chunk sizes of LC-Conformer were set to ∞ /680ms and 600ms, respectively, so the average latency was 330ms (= 600ms/2 + 30ms). Here ∞ means that the encoder sees all past frames.

Table 4: Comparisons of cascade and proposed TS-RNNT systems. “✓” and “✗” indicate streaming or offline mode, respectively. Systems that have a “✓” for the “Streaming All” column perform fully streaming inference. TS-RNNT does not use any TSE (“N/A”).

ID	ASR-Enc type	#frames of look-back		System	Streaming			CER [%] on each SNR					
		back	ahead		TSE	ASR	All	20dB	15dB	10dB	5dB	0dB	Avg.
1	Uni-Conformer	∞	3	TSE+RNNT	✗	✓	✗	11.3	12.5	16.0	24.6	43.6	21.6
2				✓	✓	✓	16.9	18.9	24.1	36.2	56.8	30.6	
3				TS-RNNT	N/A	✓	✓	13.1	14.2	17.3	25.1	42.8	22.5
4	LC-Conformer	∞	[0, 63]	TSE+RNNT	✗	✓	✗	10.0	11.0	14.2	22.5	41.6	19.9
5				✓	✓	✓	15.3	17.1	21.8	32.9	53.5	28.1	
6				TS-RNNT	N/A	✓	✓	11.4	12.4	15.0	21.9	38.6	19.8
7	LC-Conformer	71	[0, 63]	TSE+RNNT	✗	✓	✗	11.0	12.3	15.8	24.7	44.2	21.6
8				✓	✓	✓	16.8	18.6	23.8	35.6	56.3	30.2	
9				TS-RNNT	N/A	✓	✓	11.9	13.0	16.0	23.8	41.2	21.2

Offline Conformer model parameters were randomly initialized, while streaming Conformer parameters were initialized with trained offline Conformer parameters. “RNNT” and “TS-RNNT” were trained with dataset (a) and (b) in Table 1, respectively. All models were trained using RNNT loss by using the Adam optimizer with 25k warmup for a total of 100 epochs. For decoding, we used alignment-length synchronous decoding with beam size of 8 [37].

3.2. Experimental results

3.2.1. Ablation study

First, we investigate the best position for the fusion of ASR encoder and speaker encoder outputs using offline TS-RNNTs. Table 2 shows the CERs at each SNR. We can see that applying the speaker encoder output to just the first ASR encoder output, $l = 1$, attained the best performance under all SNR conditions. Hereafter, we adopted $l = 1$ for all TS-RNNT models.

3.2.2. Cascade vs. integrated system using offline models

Next, we compare the baseline cascade systems, TSE+RNNT, with the proposed integrated system, TS-RNNT for offline decoding. Table 3 shows the CERs and RTFs for each SNR condition. The RNNT baseline shows the CER obtained when recognizing the mixture without performing TSE using the RNNT back-end. As expected, this system performs poorly as it cannot identify the target speaker in a mixture. By using the TSE model as front-end, TSE+RNNT could recognize the target-speaker’s speech and achieve CER of 16.5% on average. This result confirms the importance of the TSE module. However, the RTF of TSE+RNNT was much larger than that of RNNT decoding. We could reduce the RTF of the TSE+RNNT system by adopting a smaller TSE model such as [17], but this would increase the CER. Regardless of how efficient this TSE front-end could be, it would inevitably increase the RTF.

On the other hand, TS-RNNT achieves competitive or better CERs than the cascade system while equaling the RTF of RNNT. In particular, the CERs of TS-RNNT are better than those of the cascade system under severely noisy conditions, i.e. SNR 5 and 0 dB. This is probably because the TS-RNNT system is trained in E2E manner, and thus is not affected by processing artifacts that may limit performance of the TSE+RNNT system under severe conditions. We could improve the performance of the TSE+RNNT cascade system by jointly training both modules or retraining the ASR back-end on processed speech [20]. However, the RTF of such a system would remain higher than our proposed TS-RNNT.

3.2.3. Cascade vs. integrated system using streaming models

Finally, we investigate the effectiveness of TS-RNNT for streaming models. Table 4 shows the CERs under each SNR

condition. “ASR-Enc type” indicates the type of streaming Conformer, and “#frames” is the number of look-back and look-ahead frames for the ASR encoders of the RNNT and TS-RNNT models. We investigated two types of LC-Conformer; LC-Conformer with infinite history context and LC-Conformer with the history context limited to 71 frames. A “✓” and “✗” in the “Streaming” column indicates that the module is operating in streaming mode or offline mode, respectively. All systems are fully streaming systems except system ID 1, 4 and 7 that operate offline due to the TSE front-end. The SDR improvements in offline and streaming TSE models were 15.1dB and 11.1dB, respectively, which mirrors the tendency reported in prior separation studies [13].

From the table, we observe that the proposed streaming TS-RNNT models (systems 3, 6 and 9) offer comparable performance to cascade TSE+RNNT with the offline TSE module (systems 1, 4 and 7), and greatly outperform streaming cascade systems, i.e., TSE+RNNT with the streaming TSE module (systems 2, 5 and 8) for all three encoder types. Comparing the performance of the offline TS-RNNT system of Table 3 and the best streaming system of Table 4 (system 6), we observe a relative performance gap of about 20%, which may be slightly worse than the gap observed for tasks with clean speech [8]. Closing this gap will be part of our future work.

The LC-conformer with infinite lookback context (system 6) achieves the best overall performance with a latency of just 63 frames, i.e., 330 ms. Note that the performance of the LC-conformer degrades when we limit the lookback context (system 7, 8 and 9), which indicates that the future context is important, but so is the past context when processing overlapped speech.

These results demonstrate the effectiveness of the integrated system for streaming decoding, as it avoids the performance degradation caused by using the streaming TSE module and achieves performance competitive with the offline TSE module while matching the inference speed of a vanilla RNNT.

4. Conclusion

We have proposed an integrated modeling for a streaming TS-ASR, called TS-RNNT. The target-speaker embedding, which is extracted from the speaker encoder, is fused with the intermediate outputs of the RNNT encoder to allow direct recognition of the speech of a target speaker in a mixture. Our proposed system offers comparable performance to a cascade TSE+RNNT system in the offline setting, with significantly lower complexity. Indeed, our system kept the low complexity of a vanilla RNNT. Moreover, we can greatly outperform cascade systems in streaming mode. The results of this work confirm the potential of TS-RNNT to achieve streaming ASR robust against interfering speakers.

5. References

- [1] A. Gulati, J. Qin, C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, "Conformer: Convolution-augmented transformer for speech recognition," in *Proc. of INTERSPEECH*, 2020, pp. 5036–5040.
- [2] T. N. Sainath, Y. He, A. Narayanan, R. Botros, R. Pang, D. Rybach, C. Allauzen, E. Variati, J. Qin, Q.-N. Le-The, S. yiin Chang, B. Li, A. Gulati, J. Yu, C.-C. Chiu, D. Caseiro, W. Li, Q. Liang, and P. Rondon, "An efficient streaming non-recurrent on-device end-to-end model with improvements to rare-word modeling," in *Proc. of INTERSPEECH*, 2021, pp. 1777–1781.
- [3] G. Kurata and G. Saon, "Knowledge distillation from offline to streaming RNN transducer for end-to-end speech recognition," in *Proc. of INTERSPEECH*, 2020, pp. 2117–2121.
- [4] B. Li, S. yiin Chang, T. Sainath, R. Pang, Y. R. He, T. Strohmam, and Y. Wu, "Towards fast and accurate streaming end-to-end asr," in *Proc. of ICASSP*, 2020, pp. 6064–6068.
- [5] 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. of ICASSP*, 2021, pp. 5904–5908.
- [6] T. Moriya, T. Ashihara, T. Tanaka, T. Ochiai, H. Sato, A. Ando, Y. Ijima, R. Masumura, and Y. Shinohara, "SimpleFlat: A simple whole-network pre-training approach for RNN transducer-based end-to-end speech recognition," in *Proc. of ICASSP*, 2021, pp. 5664–5668.
- [7] T. Moriya, T. Ashihara, A. Ando, H. Sato, T. Tanaka, K. Matsuura, R. Masumura, M. Delcroix, and T. Shinozaki, "Hybrid RNN-T/Attention-based streaming ASR with triggered chunkwise attention and dual internal language model integration," in *Proc. of ICASSP*, 2022, pp. 8282–8286.
- [8] T. Moriya, T. Tanaka, T. Ashihara, T. Ochiai, H. Sato, A. Ando, R. Masumura, M. Delcroix, and T. Asami, "Streaming end-to-end speech recognition for Hybrid RNN-T/Attention architecture," in *Proc. of INTERSPEECH*, 2021, pp. 1787–1791.
- [9] J. Barker, S. Watanabe, E. Vincent, and J. Trmal, "The fifth CHiME speech separation and recognition challenge: dataset, task and baselines," *arXiv preprint arXiv:1803.10609*, 2018.
- [10] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, "Deep clustering: Discriminative embeddings for segmentation and separation," in *Proc. of ICASSP*, 2016, pp. 31–35.
- [11] Y. Luo and N. Mesgarani, "Tasnet: Time-domain audio separation network for real-time, single-channel speech separation," in *Proc. of ICASSP*, 2018, pp. 696–700.
- [12] Y. Qian, X. Chang, and D. Yu, "Single-channel multi-talker speech recognition with permutation invariant training," *Speech Communication*, vol. 104, pp. 1–11, 2018.
- [13] Y. Luo and N. Mesgarani, "Conv-tasnet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [14] J. Li, L. Deng, Y. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [15] K. Žmolíková, M. Delcroix, K. Kinoshita, T. Ochiai, T. Nakatani, L. Burget, and J. Černocký, "Speakerbeam: Speaker aware neural network for target speaker extraction in speech mixtures," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 800–814, 2019.
- [16] M. Delcroix, K. Žmolíková, K. Kinoshita, A. Ogawa, and T. Nakatani, "Single channel target speaker extraction and recognition with speaker beam," in *Proc. of ICASSP*, 2018, pp. 5554–5558.
- [17] Q. Wang, I. L. Moreno, M. Saglam, K. Wilson, A. Chiao, R. Liu, Y. He, W. Li, J. Pelecanos, M. Nika *et al.*, "Voicefilter-lite: Streaming targeted voice separation for on-device speech recognition," in *Proc. of INTERSPEECH*, 2020, pp. 2677–2681.
- [18] H. Sato, T. Ochiai, M. Delcroix, K. Kinoshita, T. Moriya, and N. Kamo, "Should we always separate?: Switching between enhanced and observed signals for overlapping speech recognition," in *Proc. of INTERSPEECH*, 2021, pp. 1149–1153.
- [19] H. Sato, T. Ochiai, M. Delcroix, K. Kinoshita, N. Kamo, and T. Moriya, "Learning to enhance or not: Neural network-based switching of enhanced and observed signals for overlapping speech recognition," in *Proc. of ICASSP*, 2022, pp. 6287–6291.
- [20] J. Shi, C. Zhang, C. Weng, S. Watanabe, M. Yu, and D. Yu, "Improving RNN transducer with target speaker extraction and neural uncertainty estimation," in *Proc. of ICASSP*, 2021, pp. 6908–6912.
- [21] P. Comon and C. Jutten, *Handbook of Blind Source Separation: Independent component analysis and applications*. Academic press. Elsevier, 2010.
- [22] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1702–1726, 2018.
- [23] N. Kanda, S. Horiguchi, R. Takashima, Y. Fujita, K. Nagamatsu, and S. Watanabe, "Auxiliary interference speaker loss for target-speaker speech recognition," in *Proc. of INTERSPEECH*, 2019, pp. 236–240.
- [24] P. Denisov and N. T. Vu, "End-to-end multi-speaker speech recognition using speaker embeddings and transfer learning," in *Proc. of INTERSPEECH*, 2019, pp. 4425–4429.
- [25] M. Delcroix, S. Watanabe, T. Ochiai, K. Kinoshita, S. Karita, A. Ogawa, and T. Nakatani, "End-to-end speakerbeam for single channel target speech recognition," in *Proc. of INTERSPEECH*, 2019, pp. 451–455.
- [26] A. Graves, "Sequence transduction with recurrent neural networks," in *Proc. of ICML*, 2012.
- [27] I. Sklyar, A. Piunova, and Y. Liu, "Streaming multi-speaker ASR with RNN-T," in *Proc. of ICASSP*, 2021, pp. 6903–6907.
- [28] L. Lu, N. Kanda, J. Li, and Y. Gong, "Streaming multi-talker speech recognition with joint speaker identification," in *Proc. of INTERSPEECH*, 2021, pp. 1782–1786.
- [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," *CoRR*, vol. abs/2202.00842, 2022.
- [30] C. Xu, W. Rao, E. S. Chng, and H. Li, "Time-domain speaker extraction network," in *Proc. of ASRU*. IEEE, 2019, pp. 327–334.
- [31] M. Delcroix, T. Ochiai, K. Zmolikova, K. Kinoshita, N. Tawara, T. Nakatani, and S. Araki, "Improving speaker discrimination of target speech extraction with time-domain speakerbeam," in *Proc. of ICASSP*, 2020, pp. 691–695.
- [32] K. Maekawa, H. Koiso, S. Furui, and H. Isahara, "Spontaneous speech corpus of Japanese," in *Proc. of LREC*, 2000, pp. 947–9520.
- [33] "CHiME4 challenge," http://spandh.dcs.shef.ac.uk/chime_challenge/CHiME4, Cited April 2 2021.
- [34] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," *Proc. of INTERSPEECH*, pp. 2613–2617, 2019.
- [35] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovski, G. Stemmer, and K. Veseli, "The Kaldi speech recognition toolkit," in *Proc. of ASRU*, 2011.
- [36] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proc. of INTERSPEECH*, 2018, pp. 2207–2211.
- [37] G. Saon, Z. Tuske, and K. Audhkhasi, "Alignment-length synchronous decoding for RNN Transducer," in *Proc. of ICASSP*, 2020, pp. 7799–7803.