

# T-SOT FNT: STREAMING MULTI-TALKER ASR WITH TEXT-ONLY DOMAIN ADAPTATION CAPABILITY

Jian Wu, Naoyuki Kanda, Takuya Yoshioka, Rui Zhao, Zhuo Chen, Jinyu Li

Microsoft, One Microsoft Way, Redmond, WA, USA

## ABSTRACT

Token-level serialized output training (t-SOT) was recently proposed to address the challenge of streaming multi-talker automatic speech recognition (ASR). T-SOT effectively handles overlapped speech by representing multi-talker transcriptions as a single token stream with  $\langle cc \rangle$  symbols interspersed. However, the use of a naive neural transducer architecture significantly constrained its applicability for text-only adaptation. To overcome this limitation, we propose a novel t-SOT model structure that incorporates the idea of factorized neural transducers (FNT). The proposed method separates a language model (LM) from the transducer’s predictor and handles the unnatural token order resulting from the use of  $\langle cc \rangle$  symbols in t-SOT. We achieve this by maintaining multiple hidden states and introducing special handling of the  $\langle cc \rangle$  tokens within the LM. The proposed t-SOT FNT model achieves comparable performance to the original t-SOT model while retaining the ability to reduce word error rate (WER) on both single and multi-talker datasets through text-only adaptation.

**Index Terms**— factorized neural transducer, multi-talker speech recognition, token-level serialized output training, text-only adaptation

## 1. INTRODUCTION

Multi-talker speech recognition continues to pose a significant challenge because of the serious performance drop on overlapping speech for a conventional single-talker ASR model [1]. The impact of overlapping speech is significant even with a small ratio of speech overlaps [2, 3]. With the advances of the end-to-end (E2E) ASR technique [4, 5, 6, 7], several efforts have been made for development of the E2E streaming multi-talker ASR model, e.g., SURT [8, 9], MS-RNN-T [10], MT-RNN-T [11] and t-SOT [12]. Compared with the earlier modular systems [13, 14, 15], where a speech separation module was employed to address overlapping speech, those E2E ASR solutions transcribe the multi-speaker audio directly, which makes the model simple for both optimization and deployment and potentially brings the better performance.

Among the recent studies, t-SOT models with a Transformer [16] transducer [17] structure achieved the state-of-the-art (SOTA) recognition results in multi-talker recognition on several datasets including LibriCSS [2, 12] and AMI [18, 19]. Unlike the prior models [8, 9, 11] that used two branches for the simultaneous transcriptions of the overlapping speech, t-SOT has only a single output branch to generate the token sequence from multiple speakers. To distinguish the token streams from different speakers, a special “channel change” token,  $\langle cc \rangle$ , is inserted when adjacent tokens belong to different speakers. This simple framework achieved a streaming multi-talker ASR with a simpler model architecture and lower decoding cost. Meanwhile, it also achieved comparable WER on

non-overlapping speech with a single-speaker ASR model, evading the performance degradation witnessed in other ASR [9] and earlier cascaded approaches [20, 21].

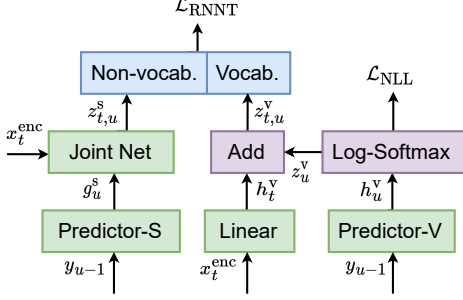
While t-SOT models achieved promising results, it still faces a challenge when we adapt it to a specific domain by using only text data. The challenge stems from its E2E architecture as well as the introduction of special token  $\langle cc \rangle$  to distinguish overlapping speakers. Firstly, it is known that the E2E ASR model is difficult to adapt by only using only text data, and many researches have been conducted. Among them, the most popular approach is the language model (LM) fusion [22, 23, 24, 25] that incorporates an external LM score on the target domain. However, the performances of these approaches are sensitive to the weight tuning on a development set. Recently, factorized neural transducer (FNT) [26, 27] was proposed to address this issue. In the FNT model, the prediction network for the regular vocabulary tokens acts as a standard LM. Therefore, various LM adaptation techniques could be applied with the text-only corpus. However, it is not straightforward to integrate neither the LM fusion nor FNT with t-SOT. Specifically, a t-SOT model generates tokens spoken by all the speakers in a chronological order together with  $\langle cc \rangle$  token. The intermingled word sequences from various speakers, along with the  $\langle cc \rangle$  token, disrupt the inherent order of natural language. Consequently, it leads challenges for the standard LM to handle the decoding sequence of the t-SOT model properly, either through LM fusion or FNT.

In this paper, we proposed a novel factorized neural transducer structure named t-SOT FNT to enable the text-only adaptation while maintaining the advantages of the t-SOT based multi-talker ASR. Our changes include two aspects. Firstly, we change the joint network of FNT to output not only the probability of  $\langle blank \rangle$  token but also that of  $\langle cc \rangle$  token. Secondly,  $N$  hidden states are maintained within the vocabulary predictor for switching among  $N$  concurrent speakers, where  $N = 2$  in this work. When  $\langle cc \rangle$  token is emitted, vocabulary predictor will turn to use the other hidden state for the future inference and give all-zero output for the current step. Our experiments show that, compared to a naive t-SOT model, the proposed t-SOT FNT model can achieve comparable performance on meeting conversation data, while achieving better WER on general single-talker ASR set by leveraging a better initialization of the vocabulary predictor network. Moreover, we observe that further WER reduction on both single and multi-talker datasets can be achieved through the text-only adaptation on vocabulary prediction network.

## 2. RELATED WORKS

### 2.1. Token-level Serialized Output Training (t-SOT)

The basic idea of the t-SOT is to generate the tokens of all the speakers in a single sequence with a streaming fashion, according to the order of the token emission time in the audio. To achieve this, the



**Fig. 1.** Illustration of the factorized neural transducer structure.  $x_t^{\text{enc}}$  is the  $t$ -th frame feature from encoder and  $y_{u-1}$  is the predicted token of the previous step. Only green blocks have the trainable parameters.

serialized transcriptions are generated as the supervision of the t-SOT model. In the case of up to 2 concurrent speakers, a special token  $\langle \text{cc} \rangle$  is inserted between the consecutive two tokens once they are spoken by different speakers, indicating the change of the virtual channel, i.e. active speaker. During inference, a post-processing step can be employed to form the transcriptions of two channels from the t-SOT decoding result by switching the output channel index when  $\langle \text{cc} \rangle$  is emitted. For example, a decoding sequence of “hello how are  $\langle \text{cc} \rangle$  i am  $\langle \text{cc} \rangle$  you  $\langle \text{cc} \rangle$  fine thank  $\langle \text{cc} \rangle$  good  $\langle \text{cc} \rangle$  you” can be reformatted to “hello how are you good” and “I am fine thank you” for the further processing. As the t-SOT only changes the supervised label of the ASR model, with the network structure, training loss function consistent as in conventional ASR. Various improvements from one speaker ASR can be easily integrated into t-SOT framework for further performance improvement, such as better network architecture, training scheduler etc. Meanwhile, we can generalize the t-SOT to support the mixed audio of up to more than 2 concurrent speakers by adding more special “channel change” tokens. However, given the overlaps of two speakers are the most common case in the real environment, we continue to prioritize our efforts on the scenario with up to 2 concurrent speakers.

## 2.2. Factorized Neural Transducer (FNT)


FNT decomposes the posterior prediction of the output token set  $\mathcal{O}$  into two parts, i.e.,  $\mathcal{O} = \{\mathcal{V} \cup \mathcal{S}\}$ , where  $\mathcal{V}$  and  $\mathcal{S}$  refer to regular vocabulary tokens and special non-vocabulary token (blank), respectively. As shown in Figure.1, given the  $t$ -th frame output from the acoustic encoder  $x_t^{\text{enc}}$  and the predicted token  $y_{u-1}$  in the previous step, the prediction of the (blank) token follows the standard transducer framework by estimating  $z_{t,u}^s$  as

$$z_{t,u}^s = \text{Joint}(x_t^{\text{enc}}, g_u^s), \quad (1)$$

where  $g_u^s$  is the output of the special prediction network with an input of  $y_{u-1}$ . The regular vocabulary prediction  $z_{t,u}^v$  is computed from  $x_t^{\text{enc}}$  and the vocabulary predictor output  $h_u^v$  as followings.

$$\begin{aligned} h_t^v &= \text{Linear}(x_t^{\text{enc}}), \\ z_u^v &= \text{LogSoftmax}(h_u^v), \\ z_{t,u}^v &= h_t^v + z_u^v. \end{aligned} \quad (2)$$

Finally,  $z_{t,u}^s$  and  $z_{t,u}^v$  are concatenated to form the distribution over  $\mathcal{O}$  for the training with the transducer loss  $\mathcal{L}_{\text{RNNT}}$ . Besides, a negative log likelihood (NLL) loss is also applied on  $z_u^v$  to enforce the



**Fig. 2.** The behavior of the vocabulary predictor in proposed integrated t-SOT FNT.  $\{y_0, \dots, y_{14}\}$  are input tokens and  $c_0, c_1$  are two hidden states. The green block indicates one processing step of the vocabulary predictor.  $\langle \text{cc} \rangle$  serves as the “switch” of two hidden states used for the vocabulary predictor and once it’s emitted, the output of the vocabulary predictor is assigned as all-zero vector.

vocabulary predictor to act as a standalone LM<sup>1</sup>. As a result, the final objective function for FNT is defined as

$$\mathcal{L}_{\text{FNT}} = \mathcal{L}_{\text{RNNT}} + \lambda \cdot \mathcal{L}_{\text{NLL}}, \quad (3)$$

where  $\lambda$  is a hyper-parameter to control the weight of the LM loss. During the text only adaptation, the vocabulary predictor is adapted to a target domain based on the available text data. As prior work [27] has shown that adding Kullback-Leibler (KL) divergence loss between the outputs of the adapted vocabulary predictor and original ones can help to avoid the performance degradation on the general domain, the final objective function for text-only adaptation of the vocabulary predictor is

$$\mathcal{L}_{\text{adapt}} = \mathcal{L}_{\text{NLL}} + \omega \cdot \mathcal{L}_{\text{KLD}}, \quad (4)$$

where  $\omega$  is the weight of the KL divergence loss.

## 3. T-SOT FNT

A naive way to incorporate t-SOT framework into FNT architecture is training FNT by using t-SOT based transcription by treating  $\langle \text{cc} \rangle$  token as a member of regular vocabulary tokens  $\mathcal{V}$ . We refer this combination as “naive t-SOT FNT”. The naive t-SOT FNT can achieve multi-talker ASR while keep using the original FNT architecture. However, there are two obstacles in this naive combination. Firstly, t-SOT transcription includes an additional channel switching token, i.e.  $\langle \text{cc} \rangle$ , which does not contain semantic meaning and will potentially introduce disruption to the LM. Secondly, as the serialized outputs in t-SOT mixes transcription from multiple speakers, which breaks the inherent order of natural language and conflicts with the standard LM optimized for single speaker data. As such, the naive combination of FNT and t-SOT will result in inferior performance and difficulty to incorporate external LM into the vocabulary predictor.

To address the above limitations, we propose a new variant of FNT, named “integrated t-SOT FNT”. In the integrated t-SOT FNT framework, we treat  $\langle \text{cc} \rangle$  as a member of special non-vocabulary tokens  $\mathcal{S}$ , and predict it by the joint network instead of the vocabulary predictor. In addition, to enable effective integration of external LM into the vocabulary predictor, we introduce the special procedure to handle multiple LM states in the vocabulary predictor.

The behavior of the proposed vocabulary predictor is exemplified in Figure 2. Unlike the conventional FNT as well as the naive t-SOT FNT, the integrated t-SOT FNT maintains multiple hidden

<sup>1</sup>It’s equivalent to the way adopted in previous FNT work [26, 27] that applied cross entropy (CE) loss on vocabulary predictor’s output  $h_u^v$ .

**Algorithm 1** Inference process of the vocabulary predictor in t-SOT FNT with up to two concurrent speakers

```
1: Input:  $\mathcal{Y} = \{y_0, \dots, y_{U-1}\}$ 
2: Initialize:  $c_0 = \emptyset, c_1 = \emptyset, \mathcal{H} = \{\}, j = 0$ 
3: for  $u$  in  $\{0, \dots, U - 1\}$  do
4:   if  $u = 0$  then
5:      $h_u^v, c_0 = \text{PredictorV}(y_u, c_0)$ 
6:      $c_1 = c_0$ 
7:   else
8:     if  $y_u = \langle \text{cc} \rangle$  then
9:        $j = 1 - j$ 
10:       $h_u^v = \emptyset$ 
11:     else
12:        $h_u^v, c_j = \text{PredictorV}(y_u, c_j)$ 
13:     Append  $h_u^v \rightarrow \mathcal{H}$ 
14: return  $\mathcal{H}$ 
```

states within the vocabulary predictor. The number of states to maintain is equal to the maximum number of concurrently active speaker, which is pre-defined to be 2 in this work. The  $\langle \text{cc} \rangle$  from the decoded sequence serves as the “switch” of two hidden states used for the vocabulary predictor, which enables each hidden state to capture the semantic transition of tokens from single speaker, as in the standard LM. In this way, an external LM can be smoothly integrated into the t-SOT framework.

More concretely, the inference process of the vocabulary predictor on the intermingled token sequence  $\mathcal{Y} = \{y_0, \dots, y_{U-1}\}$  is defined in the Algorithm 1, where the current hidden state  $c_j$  alternates between  $c_0$  and  $c_1$ . In transducer, as  $\mathcal{Y}$  always starts with the  $\langle \text{blank} \rangle$ , i.e.,  $y_0 = \langle \text{blank} \rangle$ , we reset the hidden states  $c_0, c_1$  at the initial inference step (line 5 and 6). Later when  $\langle \text{cc} \rangle$  is emitted, the output of the vocabulary predictor  $h_u^v$  is assigned as all-zero vector and the current hidden state  $c_j$  for the future inference is switched to  $c_{1-j}$  (line 9). For regular vocabulary token,  $h_u^v$  is calculated through vocabulary predictor based on the current token  $y_u$  and hidden state  $c_j$ , as shown in line 12. With this design, vocabulary predictor still acts like a standard LM even with the intermingled token sequence that consists of  $\langle \text{cc} \rangle$ .

In the training stage of the t-SOT FNT, we still use equation (3) as the objective function except that all-zero  $h_u^v$  with the  $\langle \text{cc} \rangle$  input are masked in the calculation of LM loss, i.e.,  $\mathcal{L}_{\text{NLL}}$ . Thanks to the multi-state design, the text-only adaptation scheme of the integrated t-SOT FNT is same as the standard FNT, which only adapts the vocabulary predictor on the text corpus following the equation (4).

## 4. EXPERIMENTS SETUP

### 4.1. Model Structure

We used the same encoder structure for all the ASR models in this work, which contained 2 convolution layers and an 18-layer Conformer [6] encoder with the chunk-wise streaming mask, resulting in a latency of 160 msec. The attention dimension of the multi-head self-attention (MHSA) layer in each Conformer block was set to 512 with 8 heads and the 2048-dim feed-forward network (FFN) layer was adopted with the Gaussian error linear unit (GELU).

The prediction network of the single-talker ASR model (referred as “CT”) and t-SOT model (referred as “t-SOT CT”) consisted of a 2-layer 1024 dimensional long short-term memory (LSTM). In both naive and integrated t-SOT FNT models, the non-vocabulary prediction network was a 2-layer 512-dim LSTM and vocabulary predictor was a 2-layer 1536-dim LSTM. We applied a dropout rate of 0.1 to those prediction networks. The dimension of the joint network were

**Table 1.** WER (%) comparison between t-SOT FNT and other models on general ASR set and AMI and ICSI test set, respectively.

Model	Seed		General	AMI		ICSI	
	Enc.	Pred.		dev	eval	dev	eval
CT	-	-	11.9	32.9	35.9	31.7	30.9
t-SOT CT	CT	CT	12.6	21.8	<b>24.6</b>	19.7	17.4
Naive t-SOT FNT	t-SOT CT	-	12.8	21.9	24.8	19.6	17.5
Integrated t-SOT FNT	CT	-	13.0	22.5	25.3	19.9	17.7
	t-SOT CT	-	12.7	21.8	24.9	<b>19.2</b>	17.2
	CT	LM	12.8	22.6	25.6	20.1	18.3
	t-SOT CT	LM	<b>12.2</b>	<b>21.8</b>	<b>24.6</b>	19.3	<b>17.1</b>

all set to 512. The regular vocabulary contained 4003 word pieces thus the output size of single talker model and t-SOT models were 4004 and 4005, respectively.

### 4.2. Data and Metric

We used 30 thousand (K) hours Microsoft in-house data, with the personally identifiable information removed, for the training of the single-talker ASR model. For t-SOT CT and t-SOT FNT model training, we used the combination of the multi-talker simulation data based on the 30K data and real meeting corpus from the training set of AMI [18] and ICSI [28] as well as the Microsoft internal meeting recordings. In the multi-talker simulation using the 30K data, we randomly mixed two utterances on-the-fly with a probability of 67%. For the rest of 33%, the original single-talker utterance was used. Finally, we used LibriSpeech text data (18 million (M) words), which was not included in the transcriptions of 30K training data, for the text-only adaptation experiments.

We evaluated our models on several datasets, including the general single-talker ASR test set, single-distant microphone audio from AMI and ICSI, and LibriSpeech-style datasets (LibriSpeech [29], LibriSpeechMix [30], LibriCSS [2]). Our general single-talker ASR test set covers various different application scenarios and consists of a total of 9.9M words. It is used to evaluate the single-speaker ASR accuracy before domain adaptation. On the other hand, AMI and ICSI were used to evaluate the multi-talker ASR accuracy before domain adaptation. We applied a causal logarithmic-loop-based automatic gain control (AGC) on AMI and ICSI to normalize audio volume. Finally, the LibriSpeech-style datasets were used to evaluate the effect of text-only adaptation in both single-talker and multi-talker ASR scenarios. We measured WER as an evaluation metric. For multi-talker test sets, we computed WER based on the algorithm proposed in [31, 19].

### 4.3. Training and Evaluation

The 80-dim log mel-filterbank using 25 msec window and 10 msec hop size was extracted as the input feature for ASR models. We applied global mean and variance normalization. All the ASR models were trained on 16 NVIDIA V100 GPUs with AdamW optimizer. For CT model, we performed 500K-step training with a linear decay learning rate scheduler. 50K warm-up steps were used and the peak learning rate was set to  $1.5e^{-3}$ . t-SOT CT model was trained for 275K steps, using the CT model as the seed. Warm-up steps were removed and the peak learning rate was set to  $2e^{-4}$ .

To simplify the training scheme of t-SOT FNT, the encoder parameters were initialized from the well-trained CT or t-SOT CT models, and the peak learning rate was set to  $4e^{-4}$  and  $2e^{-4}$ , respectively. Other configurations were kept same as t-SOT CT, including

**Table 2.** Text-only adaptation results of t-SOT FNT on LibriSpeech (LS), LibriSpeechMix (LS-Mix) and LibriCSS datasets. We used the integrated t-SOT FNT in this experiment, and refer it as t-SOT FNT.

Model	Seed		Adapt	LS		LS-Mix		LibriCSS						
	Enc.	Pred.		clean	other	dev-2spk	test-2spk	0L	0S	OV10	OV20	OV30	OV40	Avg.
CT	-	-	×	5.9	11.9	-	-	8.8	11.7	19.3	26.3	33.2	38.6	23.0
t-SOT CT	CT	CT	×	5.9	12.1	10.9	11.0	8.8	9.3	11.7	15.7	19.9	22.6	14.7
t-SOT FNT	t-SOT CT	-	×	6.2	12.6	11.6	11.8	9.1	10.0	12.2	15.6	20.5	23.1	15.1
			✓	5.0	10.7	10.3	10.4	7.9	8.6	11.1	<b>14.4</b>	18.9	22.1	13.8
		LM	×	5.5	10.9	11.0	10.6	8.6	9.0	11.2	15.4	19.9	22.5	14.5
			✓	<b>4.7</b>	<b>10.4</b>	<b>10.1</b>	<b>10.1</b>	<b>7.9</b>	<b>8.2</b>	<b>10.5</b>	14.5	<b>18.8</b>	<b>21.8</b>	<b>13.6</b>

the training steps and batch size. Following the previous FNT work [27], in order to utilize more text data, we can also initialize vocabulary predictor of the FNT from a pre-trained LM with same architecture but was trained independently on a much larger text corpus. In this work, the text corpus for LM training contains 29 Billion words. When it comes with text-only adaptation, we adapt the vocabulary predictor of the integrated t-SOT FNT model for 10K steps on 4 V100 GPUs with a learning rate decayed from  $1e^{-4}$  to 0.  $\omega$  was set to 1 to keep the performance on general ASR test set. For ASR decoding, we used a beam size of 16.

## 5. RESULTS

In this section, we first discuss the performance of the t-SOT FNT before adaptation on single and multi-talker test sets in Section 5.1. We then discuss the text-only adaptation results in Section 5.2.

### 5.1. Performance of single and multi-talker ASR

The WERs on the general single-talker ASR set, AMI and ICSI are reported in Table 1, where four model structures are listed and compared. With the same training configuration, our proposed integrated t-SOT FNT model outperformed the naive t-SOT FNT on ICSI dataset while achieving comparable WERs on AMI and general single-talker ASR test sets. This result illustrates that the proposed multi-state vocabulary predictor of the integrated t-SOT FNT works on par with the vocabulary predictor of the conventional FNT, while the former provides a way to naturally integrating external LM into the vocabulary predictor.

Among the integrated t-SOT FNT variants, we observed that starting training from t-SOT CT encoder leads better performance than that from CT. This is expected as the former one has been optimized on multi-talker data. On the other hand, initializing the vocabulary predictor with a pre-trained LM on the larger scale text data resulted better WER on general ASR set, but the improvement was limited on AMI and ICSI. The reason might be that the training corpus of the LM covers the scenarios of the general ASR set but lacks sufficient meeting conversation data.

On general single-talker ASR test set, the best t-SOT FNT model (last row) achieved better result than t-SOT CT, closing the WER gap from the single-talker CT from 0.7% (=12.6%-11.9%) to only 0.3% (=12.2%-11.9%). On AMI data, t-SOT FNT achieved similar performance with t-SOT CT while on ICSI, it outperformed t-SOT CT model by a 0.4% and 0.3% on development and evaluation sets, respectively. These results demonstrated the capability to convert an existing t-SOT CT model to a t-SOT FNT model by keeping the accuracy of the original t-SOT CT model.

### 5.2. Results of the text-only adaptation

We picked the top two integrated t-SOT FNT models from Table 1 and performed text-only adaptation using text data from LibriSpeech training set. The results were reported in Table 2, where three data sets, LibriSpeech, LibriSpeechMix, and LibriCSS, were evaluated. The performance of the t-SOT FNT were improved after adaptation, not only on single-talker audio, but also on the multi-talker data regardless of if the data is simulated mixture (LibriSpeechMix) or real mixture (LibriCSS). Overall, the t-SOT FNT with LM initialization achieved best performance. Compared with t-SOT CT, t-SOT FNT brought a relative WER reduction of 8.4% and 7.5% on LibriSpeechMix and LibriCSS, respectively. In addition, compared with CT that achieved 5.9% and 11.9% on LibriSpeech, t-SOT FNT achieved significantly better WERs of 4.7% and 10.4% by enjoying the text-only adaptation capability.

Among the t-SOT FNT models, we observed that the text-only adaptation closed the gap between the LM initialization on all the three test sets. For example, the relative averaged WER difference on LibriSpeechMix was reduced from 8.3% to 2.5%, and that on LibriSpeech was reduced from 14.6% to 4.0%. Overall, our results demonstrated that the proposed t-SOT FNT enjoyed the advantage of the vocabulary predictor where the general single-talker ASR accuracy was improved by utilizing a powerful LM, and the accuracy was further improved by the text-only domain adaptation.

## 6. CONCLUSIONS

In this paper, we proposed the t-SOT FNT model to incorporate the text-only adaption capability into the multi-talker ASR. A set of the hidden states were maintained within the vocabulary predictor to keep track the natural token transition from non-overlapping speakers. Compared with t-SOT CT model, the proposed t-SOT FNT achieved comparable WER on AMI and ICSI data sets and better WER on general single-talker ASR set. The experiments on LibriSpeech-style test set further demonstrated that significant WER reduction can be obtained by text-only domain adaptation on both single-talker and multi-talker audio.

## 7. REFERENCES

- [1] Jon Barker et al., “The fifth CHiME speech separation and recognition challenge: dataset, task and baselines,” *arXiv preprint arXiv:1803.10609*, 2018.
- [2] Zhuo Chen et al., “Continuous speech separation: Dataset and analysis,” in *ICASSP. IEEE*, 2020, pp. 7284–7288.
- [3] Desh Raj et al., “Integration of speech separation, diarization, and recognition for multi-speaker meetings: System descrip-

- tion, comparison, and analysis,” in *SLT*. IEEE, 2021, pp. 897–904.
- [4] Qian Zhang et al., “Transformer transducer: A streamable speech recognition model with transformer encoders and RNN-T loss,” in *ICASSP*. IEEE, 2020, pp. 7829–7833.
- [5] Xie Chen et al., “Developing real-time streaming transformer transducer for speech recognition on large-scale dataset,” in *ICASSP*. IEEE, 2021, pp. 5904–5908.
- [6] Anmol Gulati et al., “Conformer: Convolution-augmented transformer for speech recognition,” *arXiv preprint arXiv:2005.08100*, 2020.
- [7] Jinyu Li, “Recent advances in end-to-end automatic speech recognition,” *APSIPA Transactions on Signal and Information Processing*, vol. 11, no. 1, 2022.
- [8] Liang Lu et al., “Streaming end-to-end multi-talker speech recognition,” *IEEE Signal Processing Letters*, vol. 28, pp. 803–807, 2021.
- [9] Desh Raj et al., “Continuous streaming multi-talker asr with dual-path transducers,” in *ICASSP*. IEEE, 2022, pp. 7317–7321.
- [10] Ilya Sklyar et al., “Streaming multi-speaker asr with rnn-t,” in *ICASSP 2021–2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6903–6907.
- [11] Ilya Sklyar et al., “Multi-turn RNN-T for streaming recognition of multi-party speech,” in *ICASSP*. IEEE, 2022, pp. 8402–8406.
- [12] Naoyuki Kanda et al., “Streaming multi-talker ASR with token-level serialized output training,” *arXiv preprint arXiv:2202.00842*, 2022.
- [13] Jian Wu et al., “Improved Speaker-Dependent Separation for CHiME-5 Challenge,” in *Proc. Interspeech*, 2019, pp. 466–470.
- [14] Takuya Yoshioka et al., “Multi-microphone neural speech separation for far-field multi-talker speech recognition,” in *ICASSP*. IEEE, 2018, pp. 5739–5743.
- [15] Naoyuki Kanda et al., “Guided source separation meets a strong asr backend: Hitachi/paderborn university joint investigation for dinner party asr,” *arXiv preprint arXiv:1905.12230*, 2019.
- [16] Ashish Vaswani et al., “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [17] Alex Graves, “Sequence transduction with recurrent neural networks,” *arXiv preprint arXiv:1211.3711*, 2012.
- [18] Jean Carletta et al., “The AMI meeting corpus: A pre-announcement,” in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [19] Naoyuki Kanda et al., “VarArray meets t-SOT: Advancing the state of the art of streaming distant conversational speech recognition,” in *ICASSP*. IEEE, 2023, pp. 1–5.
- [20] Sanyuan Chen et al., “Continuous speech separation with conformer,” in *ICASSP*. IEEE, 2021, pp. 5749–5753.
- [21] Jian Wu et al., “Investigation of practical aspects of single channel speech separation for asr,” *arXiv preprint arXiv:2107.01922*, 2021.
- [22] Anjali Kannan et al., “An analysis of incorporating an external language model into a sequence-to-sequence model,” in *ICASSP*. IEEE, 2018, pp. 1–5828.
- [23] Erik McDermott et al., “A density ratio approach to language model fusion in end-to-end automatic speech recognition,” in *ASRU*. IEEE, 2019, pp. 434–441.
- [24] Ehsan Variani et al., “Hybrid autoregressive transducer (HAT),” in *ICASSP*. IEEE, 2020, pp. 6139–6143.
- [25] Zhong Meng et al., “Internal language model estimation for domain-adaptive end-to-end speech recognition,” in *SLT*. IEEE, 2021, pp. 243–250.
- [26] Xie Chen et al., “Factorized neural transducer for efficient language model adaptation,” in *ICASSP*. IEEE, 2022, pp. 8132–8136.
- [27] Rui Zhao et al., “Fast and accurate factorized neural transducer for text adaption of end-to-end speech recognition models,” in *ICASSP*. IEEE, 2023, pp. 1–5.
- [28] Adam Janin et al., “The ICSI meeting corpus,” in *Proc. ICASSP*. IEEE, 2003, vol. 1, pp. I–I.
- [29] Vassil Panayotov et al., “LibriSpeech: an ASR corpus based on public domain audio books,” in *ICASSP*. IEEE, 2015, pp. 5206–5210.
- [30] Naoyuki Kanda et al., “Serialized output training for end-to-end overlapped speech recognition,” in *Proc. Interspeech*, 2020, pp. 2797–2801.
- [31] Jonathan G Fiscus et al., “Multiple dimension levenshtein edit distance calculations for evaluating automatic speech recognition systems during simultaneous speech,” in *LREC*. Citeseer, 2006, pp. 803–808.