



End-to-End Joint Target and Non-Target Speakers ASR

Ryo Masumura, Naoki Makishima, Taiga Yamane, Yoshihiko Yamazaki, Saki Mizuno, Mana Ihori, Mihiro Uchida, Keita Suzuki, Hiroshi Sato, Tomohiro Tanaka, Akihiko Takashima, Satoshi Suzuki, Takafumi Moriya, Nobukatsu Hojo, Atsushi Ando

NTT Computer and Data Science Laboratories, NTT Corporation, Japan

ryo.masumura@ntt.com

Abstract

This paper proposes a novel automatic speech recognition (ASR) system that can transcribe individual speaker's speech while identifying whether they are target or non-target speakers from multi-talker overlapped speech. Target-speaker ASR systems are a promising way to only transcribe a target speaker's speech by enrolling the target speaker's information. However, in conversational ASR applications, transcribing both the target speaker's speech and non-target speakers' ones is often required to understand interactive information. To naturally consider both target and non-target speakers in a single ASR model, our idea is to extend autoregressive modeling-based multi-talker ASR systems to utilize the enrollment speech of the target speaker. Our proposed ASR is performed by recursively generating both textual tokens and tokens that represent target or non-target speakers. Our experiments demonstrate the effectiveness of our proposed method.

Index Terms: automatic speech recognition, target-speaker, non-target speakers, autoregressive model, enrollment speech

1. Introduction

Single-talker automatic speech recognition (ASR) systems limit their applicability because multiple utterances are often overlapped in many practical scenes. Therefore, recognizing multi-talker overlapped monaural speech signals has been the focus of much attention [1, 2]. In particular, target-speaker ASR (TS-ASR) systems that can transcribe only a target speaker's speech from the overlapped ones is known to be an effective way when the target speaker is specified. The systems are informed about the target speaker using an enrollment speech of that speaker, and can extract the same speaker's speech as the enrollment one. In this paper, we aim to improve ASR systems that use the speaker enrollment to make it suitable for conversational ASR applications such as meetings.

Toward a better performance of TS-ASR, several methods have been developed with the progress of deep learning technology. Initial TS-ASR systems have been built by cascading single-talker ASR with target-speaker speech extraction [3–7], which is a process to separate a target speaker's speech from a multi-talker overlapped one. The cascading systems are comparatively feasible to build because individual components are separately trained, but their performance is limited because overall optimization of TS-ASR cannot be achieved. Therefore, previous studies challenge to directly build TS-ASR in an end-to-end manner [8–10]. The main advantage of the end-to-end TS-ASR systems is that overall optimization to transcribe a target speaker's speech into text can be performed. In addition, since there is no need to separate the target speaker's speech as a signal, computation complexity can be reduced compared

with the cascading systems.

TS-ASR systems are really effective to execute personalized speech commands in a cocktail-party situation or record personalized life-logs; however, ignoring non-target speakers' speech is unsuitable for understanding interactions between the target speaker and non-target speakers (e.g., interactions between a target salesperson and customers in a store). In conversational ASR applications, both the target speaker's speech and non-target speakers' ones often need to be transcribed. In addition, we suspect that eliminating non-target speakers' speech is ineffective to perform accurate TS-ASR. Considering non-target speakers' information will explicitly enable us to model the dependency among the speakers and lead to improving TS-ASR performance. In fact, considering non-target interference speaker's information has achieved better performance in senone-based acoustic modeling [11]. One challenge is how we naturally consider both the target speaker and non-target speakers in an end-to-end ASR framework.

In this paper, we propose end-to-end joint target and non-target speakers ASR (TS-NTS-ASR) systems that jointly transcribe a target speaker's and non-target speakers' speech using a unified autoregressive model. Our key idea is to extend autoregressive modeling-based multi-talker ASR (MT-ASR) systems (detailed in Sections 2 and 3.2) [12, 13] to utilize the enrollment speech of the target speaker. In existing MT-ASR systems, transcriptions of multiple speakers are generated as a serialized token sequence, but it does not determine whether each transcription is spoken by the target speaker or not. Therefore, our proposed method utilizes one unified autoregressive model for not only transcribing all speakers' speech but also identifying whether they are the target speaker or non-target speakers using the enrollment speech. The identification is achieved by handling both textual tokens and tokens that represent the target or non-target speakers within the serialized token sequence. Our proposed TS-NTS-ASR systems have two advantages: they can transcribe both the target speaker's speech and non-target speakers' speech in a simple beam search decoding, and they improve TS-ASR performance compared with conventional TS-ASR systems that ignore non-target speakers' speech.

We detail our contributions as follows.

- This paper is the first study of end-to-end joint target and non-target speakers ASR systems. For the proposed autoregressive modeling, we show three serialized token sequence patterns to take the target speaker and non-target speakers' spoken text into consideration. We also show a network structure that can effectively utilize the speaker enrollment in autoregressive modeling.
- This paper also shows a method to implement non-target speakers ASR (NTS-ASR) systems that transcribe an individual speaker's speech except for the target speaker's speech.

- This paper demonstrates that our proposed method effectively transcribes both the target speaker’s speech and non-target speakers’ speech in experiments using a Japanese multi-talker overlapped speech dataset.

2. Related Work

Speech processing with target speaker enrollment: This paper is related to speech processing that uses speaker enrollment. In previous studies, target-speaker speech extraction [3–7], TS-ASR [8–10], and target-speaker voice activity detection [14–19] have been mainly investigated. These studies use the enrollment speech of a target speaker to extract the target speaker’s information. In contrast, our proposed method examines speech processing for both the target speaker and non-target speakers using the target speaker enrollment.

Multi-talker ASR: This paper is related to end-to-end MT-ASR systems that can transcribe an individual speakers’ speech from a multi-talker overlapped one. Initial end-to-end MT-ASR systems introduce modeling with multiple output branches, in which each branch generates a transcription for one speaker by considering all possible permutations of speakers [20–26]. A recent hopeful approach for end-to-end modeling is autoregressive modeling in which transcriptions of multiple speakers are recursively generated from one output branch [12, 13]. Our proposed method is regarded as an extended modeling of the latter MT-ASR to identify a target speaker or non-target speakers using the target speaker enrollment.

3. Preliminaries

This section describes target-speaker ASR (TS-ASR) and multi-talker ASR (MT-ASR) systems based on autoregressive modeling. These two systems can handle monaural multi-talker overlapped speech. The former generates only the target speaker’s spoken text and the latter generates all speakers’ spoken text.

3.1. Target-speaker ASR

TS-ASR based on autoregressive modeling predicts a generation probability of a target speaker’s spoken text $\mathbf{W} = \{w_1, \dots, w_N\}$ from monaural multi-talker overlapped speech \mathbf{X} and the target speaker’s enrollment speech \mathbf{E} , where $w_n \in \mathcal{V}$ is the n -th token in the spoken text, N is the number of tokens in the spoken text, and \mathcal{V} is the vocabulary set. In autoregressive modeling, the generation probability of \mathbf{W} is defined as

$$P(\mathbf{W}|\mathbf{X}, \mathbf{E}; \Theta_{\text{ts}}) = \prod_{n=1}^N P(w_n|w_{1:n-1}, \mathbf{X}, \mathbf{E}; \Theta_{\text{ts}}), \quad (1)$$

where Θ_{ts} represents the trainable model parameter sets and $w_{1:n-1} = \{w_1, \dots, w_{n-1}\}$. The loss function to optimize the model parameter sets is defined as

$$\mathcal{L}(\Theta_{\text{ts}}) = - \sum_{(\mathbf{X}, \mathbf{E}, \mathbf{W}) \in \mathcal{D}_{\text{ts}}} \log P(\mathbf{W}|\mathbf{X}, \mathbf{E}; \Theta_{\text{ts}}), \quad (2)$$

where \mathcal{D}_{ts} represents a paired dataset of input multi-talker speech, the target speaker’s enrollment speech, and the target speaker’s spoken text. Note that spoken text becomes empty when the target speaker is not included in the multi-talker overlapped speech.

3.2. Multi-talker ASR

MT-ASR based on autoregressive modeling predicts a generation probability of all speakers’ spoken text $\mathbf{W}^{1:T} = \{\mathbf{W}^1, \dots, \mathbf{W}^T\}$ from monaural multi-talker overlapped speech \mathbf{X} , where $\mathbf{W}^t = \{w_1^t, \dots, w_{N^t}^t\}$ is the t -th speaker’s spoken text, N^t is the number of tokens in the t -th speaker’s spoken text, and T is the number of speakers in the multi-talker overlapped speech. Multiple spoken texts in one autoregressive model are serialized into a single token sequence. Thus, the generation probability of $\mathbf{W}^{1:T}$ is defined as

$$\begin{aligned} P(\mathbf{W}^{1:T}|\mathbf{X}; \Theta_{\text{mt}}) &= P(\mathcal{S}|\mathbf{X}; \Theta_{\text{mt}}) \\ &= \prod_{l=1}^{|\mathcal{S}|} P(s_l|s_{1:l-1}, \mathbf{X}; \Theta_{\text{mt}}), \end{aligned} \quad (3)$$

where Θ_{mt} represents the trainable model parameter sets, $\mathcal{S} = \{s_1, \dots, s_{|\mathcal{S}|}\}$ is the serialized token sequence, and $s_l \in \{\mathcal{V} \cup \mathcal{O}\}$ is the l -th token in the serialized token sequence. $\mathcal{O} = \{\text{[sep]}, \text{[eos]}\}$ represents the special token set, where [sep] represents the speaker change and [eos] represents the end-of-sentence. There are multiple permutations in the order of the multiple spoken texts $\mathbf{W}^{1:T}$, so they are sorted by their start times, which is called first-in first-out. When the speaker index t is ordered by the start time, the serialized token sequence is represented as

$$\begin{aligned} \mathcal{S} &= \{w_1^1, \dots, w_{N^1}^1, \text{[sep]}, w_1^2, \dots, w_{N^2}^2, \\ &\quad \dots, w_{N^{T-1}}^{T-1}, \text{[sep]}, w_1^T, \dots, w_{N^T}^T, \text{[eos]}\}. \end{aligned} \quad (4)$$

Thus, the serialized token sequence is composed by concatenating multiple spoken texts while inserting [sep] between them and [eos] at the end of the entire sequence.

The loss function to optimize the model parameter sets is defined as

$$\mathcal{L}(\Theta_{\text{mt}}) = - \sum_{(\mathbf{X}, \mathcal{S}) \in \mathcal{D}_{\text{mt}}} \log P(\mathcal{S}|\mathbf{X}; \Theta_{\text{mt}}), \quad (5)$$

where \mathcal{D}_{mt} represents a paired dataset of the serialized token sequence and the multi-talker overlapped speech.

4. Proposed Method

This paper proposes end-to-end joint target and non-target speaker ASR (TS-NTS-ASR) systems that can transcribe an individual speaker’s speech while identifying whether they are the target speaker or non-target speakers from multi-talker overlapped speech and the target speaker’s enrollment speech.

4.1. Modeling

We construct a TS-NTS-ASR model by extending autoregressive modeling-based MT-ASR (described in Section 3.2) to utilize the target speaker’s enrollment speech. The TS-NTS-ASR model predicts a joint generation probability of a target speaker’s spoken text $\mathbf{W} = \{w_1, \dots, w_N\}$ and non-target speakers’ spoken texts $\bar{\mathbf{W}}^{1:T-1} = \{\bar{\mathbf{W}}^1, \dots, \bar{\mathbf{W}}^{T-1}\}$ from monaural multi-talker overlapped speech \mathbf{X} and the target speaker’s enrollment speech \mathbf{E} , where $\bar{\mathbf{W}}^t = \{\bar{w}_1^t, \dots, \bar{w}_{N^t}^t\}$ is the t -th non-target speaker’s spoken text. To handle these outputs within an autoregressive model, the target speaker’s spoken text and non-target speakers’ spoken texts are serialized as a single token sequence $\mathbf{Z} = \{z_1, \dots, z_{|\mathbf{Z}|}\}$. In this case, the joint

generation probability is defined as

$$\begin{aligned} P(\mathbf{W}, \bar{\mathbf{W}}^{1:T-1} | \mathbf{X}, \mathbf{E}; \Theta) &= P(\mathbf{Z} | \mathbf{X}, \mathbf{E}; \Theta) \\ &= \prod_{l=1}^{|\mathbf{Z}|} P(z_l | z_{1:l-1}, \mathbf{X}, \mathbf{E}; \Theta), \end{aligned} \quad (6)$$

where Θ represents the trainable model parameter sets. $z_l \in \{\mathcal{V} \cup \mathcal{U}\}$ is the l -th token in the serialized token sequence. $\mathcal{U} = \{[\text{t}], [\text{nt}], [\text{eos}]\}$ represents the special token set, where $[\text{t}]$ represents the target speaker's section and $[\text{nt}]$ represents non-target speaker's section.

TS-NTS-ASR systems can be associated with those described in Section 3. By disregarding \mathbf{E} from Eq. (6) to not identify the target or non-target speakers, the system can be regarded as the same modeling as MT-ASR defined by Eq. (3). In addition, by disregarding $\bar{\mathbf{W}}^{1:T-1}$ from Eq. (6), the system can be regarded as the same modeling as TS-ASR defined by Eq. (1). On the other hand, by disregarding \mathbf{W} from Eq. (6), the system can be regarded as a non-target speakers ASR (NTS-ASR) system that transcribes an individual speaker's speech except for the target speaker's speech.

The loss function to optimize the model parameter sets is defined as

$$\mathcal{L}(\Theta) = - \sum_{(\mathbf{X}, \mathbf{E}, \mathbf{Z}) \in \mathcal{D}} \log P(\mathbf{Z} | \mathbf{X}, \mathbf{E}; \Theta), \quad (7)$$

where \mathcal{D} represents a paired dataset of input multi-talker speech, the target speaker's enrollment speech, and the serialized token sequence.

4.2. Serialization of tokens

For the proposed autoregressive modeling, we show three serialized token sequence patterns. In each serialization, we can transcribe an individual speaker's speech while identifying whether they are the target or non-target speaker.

Target-speaker first: We put the target speaker's spoken text first, then the non-target speakers' spoken texts. The serialized token sequence is composed by

$$\mathbf{Z} = \{[\text{t}], w_1, \dots, w_N, [\text{nt}], \bar{w}_1^1, \dots, \bar{w}_{N1}^1, \dots, [\text{nt}], \bar{w}_1^{T-1}, \dots, \bar{w}_{N^{T-1}}^{T-1}, [\text{eos}]\}. \quad (8)$$

Note that we can perform TS-ASR by stopping when the first $[\text{nt}]$ is generated.

Non-target speaker first: We put the non-target speakers' spoken texts first, then the target speaker's spoken text. The serialized token sequence is composed by

$$\mathbf{Z} = \{[\text{nt}], \bar{w}_1^1, \dots, \bar{w}_{N1}^1, \dots, [\text{t}], w_1, \dots, w_N, [\text{eos}]\}. \quad (9)$$

Note that we can perform NTS-ASR by stopping when the first $[\text{t}]$ is generated.

First-in first-out: We sort each transcription by their start times regardless of whether the text is spoken by the target or non-target speaker. For example, when the target speaker starts speaking a second earlier among all speakers, the serialized token sequence is composed by

$$\mathbf{S} = \{[\text{nt}], \bar{w}_1^1, \dots, \bar{w}_{N1}^1, [\text{t}], w_1, \dots, w_N, [\text{nt}], \bar{w}_1^2, \dots, \bar{w}_{N2}^2, \dots, [\text{nt}], \bar{w}_1^{T-1}, \dots, \bar{w}_{N^{T-1}}^{T-1}\}. \quad (10)$$

Note that this is the same arrangement as that for MT-ASR.

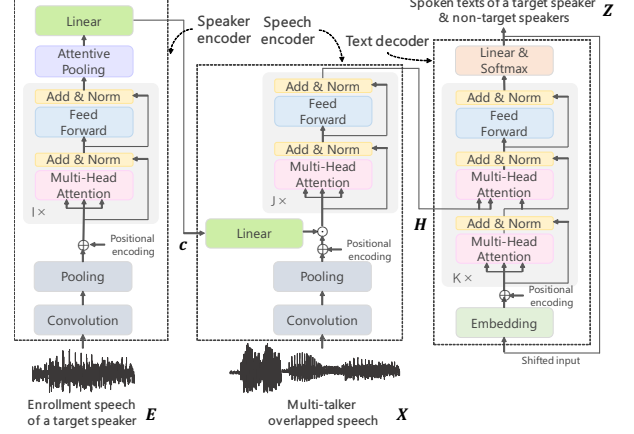


Figure 1: Network structure of TS-NTS-ASR.

4.3. Network structure

We construct the TS-NTS-ASR system from two speaker encoders and a text decoder. Figure 1 shows the network structure of TS-NTS-ASR.

Speaker encoder: The speaker encoder converts acoustic features of the target speaker's enrollment speech into a speaker vector. This conversion is defined as

$$\mathbf{c} = \text{SpeakerEnc}(\mathbf{E}; \theta_e), \quad (11)$$

where $\theta_e \in \Theta$ is its parameters. We implement this function using convolution and pooling layers, a positional encoding layer, transformer encoder blocks, an attentive pooling layer, and a linear layer.

Speech encoder: The speech encoder converts acoustic features of the multi-talker overlapped speech and the speaker vector into hidden representations. This conversion is defined as

$$\mathbf{H} = \text{SpeechEnc}(\mathbf{X}, \mathbf{c}; \theta_x), \quad (12)$$

where $\theta_x \in \Theta$ is its parameters. We implement this function using convolution and pooling layers to generate subsampled speech representations, a positional encoding layer, an element-wise multiplication layer with a linear layer, and transformer encoder blocks. The element-wise multiplication is performed between the speaker vector and subsampled speech representations to consider the target speaker's information [5].

Text decoder: The token decoder computes the generation probability of a token given preceding tokens and the hidden representations produced in the speech encoder. The generation probability is computed from

$$P(z_l | z_{1:l-1}, \mathbf{X}, \mathbf{E}, \Theta) = \text{TextDec}(z_{1:l-1}, \mathbf{H}; \theta_z), \quad (13)$$

where $\theta_z \in \Theta$ is its parameters. We implement this function using a linear embedding layer, a positional encoding layer, transformer decoder blocks, and a softmax layer with a linear transformation.

5. Experiments

In the experiments, we used the Corpus of Spontaneous Japanese (CSJ) [27], which is a single-talker speech dataset. We first split the dataset into training set (518.4 h), validation set (1.3 h), and test set (1.9 h) and, following procedures were conducted for each set. To generate multi-talker overlapped speech

Table 1: *Experimental results in terms of character error rate (%)*.

System	Speaker enrollment	Serialization of tokens	Single-talker speech		Multi-talker overlapped speech		
			Target	Non-target	Target	Non-target	Multi-talker
MT-ASR	-	First-in first-out	-	-	-	-	12.75
TS-ASR	✓	-	6.95	-	12.54	-	-
NTS-ASR	✓	-	-	8.36	-	13.80	-
TS-NTS-ASR (proposed)	✓	Target first	6.15	7.87	12.36	13.05	12.60
TS-NTS-ASR (proposed)	✓	Non-target first	6.52	7.40	13.18	13.24	13.57
TS-NTS-ASR (proposed)	✓	First-in first-out	5.29	7.14	11.80	12.53	12.15

dataset, we mixed multiple audio signals as a monaural signal. To this end, we randomly chose multiple audio signals so as not to select the same speakers. We set the number of speakers in the mixed signals as two or three. When mixing the audio signals, the original volume of each utterance was kept unchanged, resulting in an average signal-to-interference ratio of about 0 dB. For the delay applied to each utterance, the start times of the individual utterances differ by 0.5 s or longer. In addition, every utterance in each mixed audio sample has at least one speaker-overlapped region with other utterances. In addition, we prepared enrollment speech for the single-talker speech and the multi-talker overlapped speech sets. For the single-talker speech dataset, we prepared two cases: a different utterance of the target speaker was used as the speaker enrollment, and an utterance of a different speaker from the target speaker was used as the speaker enrollment. For the multi-talker overlapped speech dataset, an utterance spoken by one of the talkers was used as the enrollment. In training, all paired patterns (about 2,500 h) were jointly used. In testing, each of paired patterns were individually evaluated.

5.1. Setups

We constructed an MT-ASR, TS-ASR, NTS-ASR, and three TS-NTS-ASR systems. Except for MT-ASR, which does not use speaker enrollment, we introduced the network structure described in Section 4.3 for each system. MT-ASR used a network structure that excludes a speaker encoder and its connection. For these systems, the transformer blocks were composed under the following conditions: the dimensions of the output continuous representations were set to 512, the dimensions of the inner outputs were set to 2,048, and the number of heads in the multi-head attentions was set to 4. In the nonlinear transformational functions, the Swish activation was used. For both the speaker and the speech encoders, we used 80 log mel-scale filterbank coefficients as acoustic features. The frame shift was 10 ms. The acoustic features passed two convolution and max pooling layers with a stride of 2, so we down-sampled them to 1/4 along with the time axis. After these layers, we stacked 4-layer transformer encoder blocks. For the text decoder, we used 512-dimensional character embeddings where the vocabulary size was set to 3,262. We also stacked 3-layer transformer decoder blocks. For the training, we used the RAdam optimizer [28]. For the systems that used speaker enrollment, the speaker encoder was first trained using VoxCeleb2 dataset [29] with ArcFace criterion [30], and then the speech encoder and text decoder were trained while freezing the speaker encoder. We set the mini-batch size to 64 utterances and the dropout rate in the transformer blocks to 0.1. We introduced label smoothing where its smoothing parameter was set to 0.1. In addition, we applied SpecAugment [31] and enrollment-less training [17]. For testing, we used a beam search algorithm in which the beam size was set to 4.

Table 2: *Experimental results in terms of target speaker or non-target speaker detection error rate (%)*. *TS-NTS-ASR system introduced “first-in first-out” serialization.*

	Single-talker speech	
	Target	Non-target
TS-ASR	1.41	2.03
NTS-ASR	2.10	2.35
TS-NTS-ASR (proposed)	0.00	1.82

5.2. Results

Table 1 shows the results in terms of character error rate for each ASR system. The “Target” column represents TS-ASR performance and the “Non-target” column represents NTS-ASR performance, each of which was only evaluated when using the speaker enrollment. The “Multi-talker” column represents MT-ASR performance that evaluated all speakers’ results. First, the results show that each TS-NTS-ASR system outperformed the TS-ASR system. This indicates that considering non-target speakers’ information explicitly leads to improving TS-ASR performance. In addition, each TS-NTS-ASR system outperformed the NTS-ASR system. This also indicates that joint consideration of the target and non-target speakers is important. Next, among the TS-NTS-ASR systems, the best results were attained by “first-in first-out” serialization. It is considered that “first-in first-out,” which can attend to the input of the multi-talker overlapped speech in left-to-right order, is easier to train than “target first” and “non-target first.” Furthermore, Table 2 shows the results of evaluating the target speaker or non-target speaker detection error rate against single-talker speech to reveal the performance of considering the speaker enrollment. The error was calculated using the number of utterances that misdetects the target or non-target speaker. The results show that TS-NTS-ASR could reduce the detection error rate for both the target and non-target speakers. This induces TS-ASR and NTS-ASR performance improvements. These results demonstrate that our proposed method effectively transcribes not only the target speaker’s speech but also non-target speakers’ speech.

6. Conclusions

We presented end-to-end joint target and non-target speakers ASR systems that jointly transcribe a target speaker’s speech and non-target speakers’ speech using the target speaker’s enrollment speech. The key strength of the proposed method is to effectively transcribe all speakers’ speech while identifying whether they are the target speaker or non-target speakers using the enrollment speech. This is suitable to understand interactions between the target speaker and non-target speakers. Our experimental results showed that the proposed method effectively transcribes not only the target speaker’s speech but also non-target speakers’ speech, and outperformed the target-speaker ASR system and the non-target speaker ASR system.

7. References

- [1] T. J. Park, N. Kanda, D. Dimitriadis, K. J. Han, S. Watanabe, and S. Narayanan, "A review of speaker diarization: Recent advances with deep learning," *Computer Speech & Language*, vol. 72, p. 101317, 2022.
- [2] K. Zmolikova, M. Delcroix, T. Ochiai, K. Kinoshita, J. Cernocky, and D. Yu, "Neural target speech extraction: An overview," *arXiv:2301.13341*, 2023.
- [3] M. Delcroix, K. Zmolikova, K. Kinoshita, A. Ogawa, and T. Nakatani, "Single channel target speaker extraction and recognition with speaker beam," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 5554–5558, 2018.
- [4] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, "Voice-filter: Targeted voice separation by speaker-conditioned spectrogram masking," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2728–2732, 2019.
- [5] C. Xu, W. Rao, E. S. Chng, and H. Li, "Time-domain speaker extraction network," *In Proc. Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 327–334, 2019.
- [6] Q. Wang, I. L. Moreno, M. Saglam, K. Wilson, A. Chiao, R. Liu, Y. He, W. Li, J. Pelecanos, and M. Nika, "Voicefilter-lite: Streaming targeted voice separation for on-device speech recognition," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2677–2681, 2020.
- [7] X. Ji, M. Yu, C. Zhang, D. Su, T. Yu, X. Liu, and D. Yu, "Speaker-aware target speaker enhancement by jointly learning with speaker embedding extraction," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 7289–7293, 2020.
- [8] 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. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 451–455, 2019.
- [9] P. Denisov and N. T. Vu, "End-to-end multi-speaker speech recognition using speaker embeddings and transfer learning," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 4425–4429, 2019.
- [10] T. Moriya, H. Sato, T. Ochiai, M. Delcroix, and T. Shinozaki, "Streaming target-speaker ASR with neural transducer," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2673–2677, 2022.
- [11] N. Kanda, S. Horiguchi, R. Takashima, Y. Fujita, K. Nagamatsu, and S. Watanabe, "Auxiliary interference speaker loss for target speaker speech recognition," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 236–240, 2019.
- [12] N. Kanda, Y. Gaur, X. Wang, Z. Meng, and T. Yoshioka, "Serialized output training for end-to-end overlapped speech recognition," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2797–2801, 2020.
- [13] R. Masumura, D. Okamura, N. Makishima, M. Ithori, A. Takashima, T. Tanaka, and S. Orihashi, "Unified autoregressive modeling for joint end-to-end multi-talker overlapped speech recognition and speaker attribute estimation," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2591–2595, 2021.
- [14] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny, A. Laptev, and A. Romanenko, "Target-speaker voice activity detection: a novel approach for multi-speaker diarization in a dinner party scenario," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 274–278, 2020.
- [15] S. Ding, Q. Wang, S. yin Chang, Li Wan, and I. L. Moreno, "Personal VAD: Speaker-conditioned voice activity detection," *In Proc. Odyssey The Speaker and Language Recognition Workshop (Odyssey)*, pp. 433–439, 2020.
- [16] M. He, D. Raj, Z. Huang, J. Du, Z. Chen, and S. Watanabe, "Target-speaker voice activity detection with improved i-vector estimation for unknown number of speaker," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 3555–3559, 2021.
- [17] N. Makishima, M. Ithori, T. Tanaka, A. Takashima, S. Orihashi, and R. Masumura, "Enrollment-less training for personalized voice activity detection," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 346–350, 2021.
- [18] S. Ding, R. Rikhye, Q. Liang, Y. He, Q. Wang, A. Narayanan, T. O'Malley, and I. McGraw, "Personal VAD 2.0: Optimizing personal voice activity detection for on-device speech recognition," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 3744–3748, 2022.
- [19] W. Wang and M. Li, "Incorporating end-to-end framework into target-speaker voice activity detection," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 8362–8366, 2022.
- [20] D. Yu, X. Chang, and Y. Qian, "Recognizing multi-talker speech with permutation invariant training," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2456–2460, 2017.
- [21] X. Chang, Y. Qian, K. Yu, and S. Watanabe, "End-to-end monaural multi-speaker asr system without pretraining," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 6256–6260, 2019.
- [22] H. Seki, T. Hori, S. Watanabe, J. L. Roux, and J. R. Hershey, "A purely end-to-end system for multi-speaker speech recognition," *In Proc. Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 2620–2630, 2018.
- [23] S. Settle, J. L. Roux, T. Hori, and S. W. adn John R. Hershey, "End-to-end multi-speaker speech recognition," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 4819–4823, 2018.
- [24] X. Chang, W. Zhang, Y. Qian, J. L. Roux, and S. Watanabe, "End-to-end multi-speaker speech recognition with transformer," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 6129–6133, 2020.
- [25] I. Sklyar, A. Pionova, and Y. Liu, "Streaming multi-speaker ASR with RNN-T," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 6903–6907, 2021.
- [26] A. Tripathi, H. Lu, and H. Sak, "End-to-end multi-talker overlapping speech recognition," *In Proc. International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pp. 6124–6128, 2020.
- [27] K. Maekawa, H. Koiso, S. Furui, and H. Isahara, "Spontaneous speech corpus of Japanese," *In proc. International Conference on Language Resources and Evaluation (LREC)*, pp. 947–952, 2000.
- [28] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *In Proc. International Conference on Learning Representations (ICLR)*, 2020.
- [29] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 1086–1090, 2018.
- [30] J. Deng, J. Guo, J. Yang, N. Xue, I. Kotsia, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," *In Proc. IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR)*, pp. 4690–4699, 2018.
- [31] 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," *In Proc. Annual Conference of the International Speech Communication Association (INTERSPEECH)*, pp. 2613–2617, 2019.