

BW-EDA-EEND: STREAMING END-TO-END NEURAL SPEAKER DIARIZATION FOR A VARIABLE NUMBER OF SPEAKERS

Eunjung Han Chul Lee Andreas Stolcke

Amazon Alexa Speech, Sunnyvale, CA, USA
{cehan, chulle, stolcke}@amazon.com

ABSTRACT

We present a novel online end-to-end neural diarization system, BW-EDA-EEND, that processes data incrementally for a variable number of speakers. The system is based on the Encoder-Decoder-Attractor (EDA) architecture of Horiguchi et al., but utilizes the incremental Transformer encoder, attending only to its left contexts and using block-level recurrence in the hidden states to carry information from block to block, making the algorithm complexity linear in time. We propose two variants: For unlimited-latency BW-EDA-EEND, which processes inputs in linear time, we show only moderate degradation for up to two speakers using a context size of 10 seconds compared to offline EDA-EEND. With more than two speakers, the accuracy gap between online and offline grows, but the algorithm still outperforms a baseline offline clustering diarization system for one to four speakers with unlimited context size, and shows comparable accuracy with context size of 10 seconds. For limited-latency BW-EDA-EEND, which produces diarization outputs block-by-block as audio arrives, we show accuracy comparable to the offline clustering-based system.

Index Terms— Speaker diarization, end-to-end neural diarization, encoder-decoder attractor, online inference.

1. INTRODUCTION

Speaker diarization is the process of partitioning an audio stream into homogeneous segments according to speaker identities. Thus, diarization determines “who spoke when” in a multi-speaker environment, with a variety of applications to conversations involving multiple speakers, such as meetings, television shows, medical consultations, or call center conversations. In particular, the speaker boundaries produced by a diarization system can be used to map transcripts generated by a multi-speaker automatic speech recognition (ASR) system into speaker-attributed transcripts [1, 2, 3]. Moreover, speaker embeddings inferred by diarization can help the ASR system adapt to, or focus on the speech of a targeted speaker [4].

Conventional speaker diarization systems are based on clustering of speaker embeddings. In this approach, several components are integrated into a single system: speech segments are determined by voice activity detection (VAD); these speech segments are further divided into smaller chunks of fixed size; speaker embeddings are then extracted by speaker embedding extractors for each chunk; finally, those speaker embeddings are clustered to map each segment to a speaker identity [5, 6, 7, 8, 9, 10]. For embeddings, i-vectors [11], x-vectors [12], or d-vectors [13] are commonly used. Clustering methods typically used for speaker diarization are agglomerative hierarchical clustering (AHC) [5, 6, 8], k-means clustering [9], and spectral clustering [10]. Recently, neural network-based clustering has been explored [14]. Clustering-based speaker diarization

achieves good performance but has several shortcomings. First, it relies on multiple modules (VAD, embedding extractor, etc.) that are trained separately. Therefore, clustering-based systems require careful joint calibration in the building process. Second, systems are not jointly optimized to minimize diarization errors; clustering in particular is an unsupervised process. Finally, clustering does not accommodate overlapping speech naturally, even though recent work has proposed ways to handle regions with simultaneously active speakers in clustering [15].

End-to-end neural diarization (EEND) with self-attention [16] is one of the approaches that aim to model the joint speech activity of multiple speakers. It integrates voice activity and overlap detection with speaker tracking in end-to-end fashion. Moreover, it directly minimizes diarization errors and has demonstrated excellent diarization accuracy on two-speaker telephone conversations. However, EEND as originally formulated is limited to a fixed number of speakers because the output dimension of the neural network needs to be prespecified. Several methods have been proposed recently to overcome the limitations of EEND. One approach uses a speaker-wise chain rule to decode a speaker-specific speech activity iteratively conditioned on previously estimated speech activities [17]. Another approach proposes an encoder/decoder-based attractor calculation [18]. The embeddings of multiple speakers are accumulated over the time course of the audio input, and then disentangled one-by-one, for speaker identity assignment by speech frame. However, all these state-of-the-art EEND methods only work in an offline manner, which means that the complete recording must be available before diarization output is generated. This makes their application impractical for settings where potentially long multi-speaker recordings need to be processed incrementally (in streaming fashion).

In this study, we propose a novel method to perform EEND in a blockwise *online* fashion so that speaker identities are tracked with low latency soon after new audio arrives, without much degradation in accuracy compared to the offline system. We utilize the incremental Transformer encoder, where we attend to only its left contexts and ignore its right contexts, thus enabling blockwise online processing. Furthermore, the incremental Transformer encoder uses block-level recurrence in the hidden states to carry over information block by block, reducing computation time while attending to previous blocks. To our knowledge, ours is the first method that uses the incremental Transformer encoder with block-level recurrence to enable online speaker diarization.

2. PRIOR WORK

Horiguchi et al. [18] proposed a version of EEND that estimates the number of speakers by aggregating multiple speaker embeddings into a small number of attractors, one per speaker (EEND

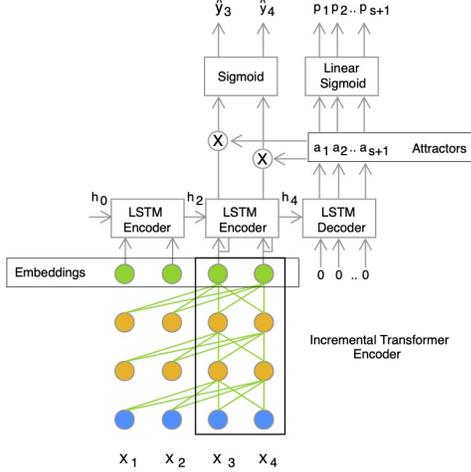


Fig. 1. Schematic diagram of BW-EDA-EEND

with Encoder Decoder Attractor, EDA-EEND). To calculate a flexible number of attractors from variable lengths of speaker embedding sequences, they utilize an LSTM-based encoder-decoder [19]. Given a T -length sequence of F -dimensional audio features, $X = [x_1, \dots, x_T]$, $x_t \in R^F$, a Transformer encoder computes D -dimensional diarization embeddings, $E = [e_1, \dots, e_T]$, $e_t \in R^D$. Then, these embeddings are fed into a unidirectional LSTM encoder, producing the final hidden and cell states (h_0, c_0) . Assuming there are S speakers in X , ideally a set of S time-invariant D -dimensional attractors, $A = [a_1, \dots, a_S]$, $a_s \in R^D$, are then decoded, also using a unidirectional LSTM, with initial states (h_0, c_0) as follows:

$$\begin{aligned}
 E &= \text{TransformerEncoder}(X), \\
 (h_0, c_0) &= \text{LSTMEncoder}(E), \\
 a_s, (h_s, c_s) &= \text{LSTMDecoder}(0, (h_{s-1}, c_{s-1})) \text{ for } 1 \leq s \leq S + 1, \\
 p_s &= \sigma(\text{Linear}(a_s)) \text{ for } 1 \leq s \leq S + 1, \\
 \hat{Y} &= \sigma(EA^T), A = [a_1, \dots, a_S]
 \end{aligned}$$

where 0 is the D -dimensional zero vector and used as a constant input for the decoder. Attractor existence probabilities p_s are estimated using a fully connected layer with a sigmoid function. A probability below a chosen threshold indicates that all attractors have been decoded, thereby implicitly determining the number of speakers. Joint speech activities for S speakers (\hat{Y}) are estimated by a matrix multiplication of embeddings and attractors (EA^T), followed by a sigmoid function.

3. BLOCKWISE EDA-EEND

The original EDA-EEND algorithm [18] performs an *offline* inference operation by first encoding all input frames, aggregating all speaker embeddings, and then finally decoding the attractors. This means that it is not suitable to perform *online* (responsive) speaker diarization. We propose a family of new blockwise EDA-EEND variants that enable online speaker diarization when (1) latency needs to be limited or (2) computation needs to be linear in the input length. To this end, we perform blockwise EDA and EEND operations by incrementally performing various encoding and decoding operations per input audio block, each of only a few seconds in duration. First, to enable incremental processing, we limit the transformer encoder to attend only to a block's left contexts (ignoring

right contexts). Second, we perform the blockwise recurrence computation throughout the hidden states of the transformer encoder, carrying over information from one block to the next block. This approach was motivated by the segment-level recurrence with state reuse technique used in Transformer XL [20]. These blockwise recursive operations in the transformer encoder can significantly reduce the overall inference time while also saving memory, since the hidden states from the previous blocks are cached and then reused instead of being re-computed from scratch for every block. More importantly, with a limited context size, the computational complexity of the given Transformer becomes linear-time with respect to audio length, instead of having quadratic complexity. Although each block can only model local dependency information, the top node of the encoder can still model long-term dependencies via the recurrence. Also, this is shown to reduce speaker permutations between blocks in our experiments.

Our algorithm is illustrated in Figure 1. Given a T -length sequence of F -dimensional audio features, $X = [x_1, \dots, x_T]$, let $X_b = [x_{b,1}, \dots, x_{b,W}]$ denote the b th block input sequence of size W . Given the Transformer encoder with n layers, let $E_b^i = [e_{b,1}^i, \dots, e_{b,W}^i]$ denote its hidden states in the i th layer. The Transformer encoder computes the sequence of D -dimensional diarization embeddings of size W for the b th block, $E_b = [e_{b,1}, \dots, e_{b,W}]$ as:

$$\begin{aligned}
 Q_b^i &= E_b^{i-1} \text{ for } 1 \leq i \leq n \ (E_b^0 = X_b), \\
 K_b^i &= V_b^i = \text{Concat}(E_{b-L:b-1}^{i-1}, E_b^{i-1}) \text{ for } 1 \leq i \leq n, \\
 E_b^i &= \text{TransformerEncoderLayer}(Q_b^i, K_b^i, V_b^i) \text{ for } 1 \leq i < n, \\
 E_b &= E_b^n
 \end{aligned}$$

where $E_{b-L:b-1}^{i-1}$ is a sequence of hidden states that are produced from the previous L blocks (i.e., going back L blocks with respect to the current b th block) and Concat is a concatenation function of two hidden block sequences along the time dimension. Motivated by Transformer-XL [20], both key K_b^i and value V_b^i vectors use the same hidden states from the current block (E_b^{i-1}) along with other cached hidden states from the previous blocks $E_{b-L}^{i-1}, \dots, E_{b-1}^{i-1}$. When $L = \infty$, we use all hidden states produced from all previous blocks to E_b^{i-1} .

Having made the Transformer encoder operation incremental and linear, we now propose two variants of BW-EDA-EEND that differ in how the attractors are computed. The first variant, denoted limited-latency BW-EDA-EEND (BW-EDA-EEND-LL), computes attractors for each block and produces outputs with limited-latency (according to the block size), suitable for generating outputs online.

The second variant, denoted unlimited-latency BW-EDA-EEND (BW-EDA-EEND-UL), computes attractors at the end of the inputs. It still has unlimited latency, but limits embedding computation to be linear in input length with a limited context size.

For BW-EDA-EEND-LL, block dependent D -dimensional attractors, $A_b = [a_{b,1}, \dots, a_{b,S}]$ are decoded using a blockwise unidirectional LSTM as follows:

$$\begin{aligned}
 I_b &= \text{Concat}(E_{b-L:b-1}, E_b), \\
 (h_b, c_b) &= \text{LSTMEncoder}(I_b, (h_{b-L}, c_{b-L})), \\
 a_{b,s}, (h_{b,s}^*, c_{b,s}^*) &= \text{LSTMDecoder}(0, (h_{b,s-1}^*, c_{b,s-1}^*)) \\
 &\text{ for } 1 \leq s \leq S + 1, \\
 p_{b,s} &= \sigma(\text{Linear}(a_{b,s})) \text{ for } 1 \leq s \leq S + 1, \\
 \hat{Y}_b &= \sigma(E_b A_b^T), A_b = [a_{b,1}, \dots, a_{b,S}]
 \end{aligned}$$

where (h_b, c_b) are the hidden and cell states of the LSTM encoder at the end of the b th block (i.e., the LSTM encoder runs unidirection-

ally along a time window of length W -length), and $E_{b-L:b-1}$ is the sequence of diarization embeddings produced from the previous L blocks. Note that $(h_{b,0}^*, c_{b,0}^*) = (h_b, c_b)$, meaning that the last hidden state from the LSTM encoder is used as an initial hidden state for the LSTM decoder. When $L = \infty$, we use (h_0, c_0) instead of (h_{b-L}, c_{b-L}) in the LSTM encoder.

For BW-EDA-EEND-UL, all blockwise LSTM encoding and decoding operations are similar to those of BW-EDA-EEND-LL, except that block-independent D -dimensional attractors $A = [a_1, \dots, a_S]$ are computed at the end of the last block $m = \lfloor T/W \rfloor$.

During inference, the number of speakers S is not known. We estimate S , so that the first S attractor existence probabilities are greater or equal to a chosen threshold τ (typically, $\tau = 0.5$), similar to the original EDA-EEND [18]. That is, $S = \max\{s | s \in Z^+ \wedge p_{b,s} \geq \tau\}$. We use the first S attractors for the subsequent downstream computation in BW-EDA-EEND-UL. For the blockwise attractor estimation in BW-EDA-EEND-LL, we force the number of speakers at each block S_b to be non-decreasing, $S_b \geq S_{b-1}$, thereby preventing the model from forgetting the speakers that occurred previously even though they are not active in the current block. Forcing S_b to be non-decreasing helps all attractors to be kept in memory for possible matching when the corresponding speakers reappear later.

Since we re-decode all attractors after each block, we need to ensure their consistency in both order and values. We use several heuristics in the inference phase to encourage consistency of attractors, and therefore speaker labels, across blocks. In ablation experiments (Table 1), we show that the individual heuristics incrementally improve the overall performance of our BW-EDA-EEND algorithms. First, we reorder (permute) attractors (in a greedy search) so that the cosine similarities between previous and current block are maximized. Additionally, we average the values of attractors from the previous block and the current block to make speaker representations vary more smoothly as we process more audio information.

Finally, we allow the shuffling of diarization embeddings (outputs of the Transformer encoder, used as inputs to the LSTM encoder) across blocks. The use of LSTMs as a Markovian memory mechanism and the shuffling heuristic comes from the original EDA algorithm. The purpose is to incrementally add speaker embeddings to the memory representation. Shuffling enforces invariance of the memory mechanism with respect to the order in which speakers appear in the input. We shuffle diarization embeddings randomly within the current and L_s context blocks before sending them to the LSTM encoder. We set $L_s = L$ if L is finite, and $L_s = 0$ if $L = \infty$.

4. EXPERIMENTS

For EDA-EEND, we used the simulation recipe in [16] to generate meeting mixtures of one to four speakers. We simulated meeting mixtures based on Switchboard-2 (Phase I, II, III), Switchboard Cellular (Part 1, 2), and the 2004-2008 NIST Speaker Recognition Evaluation (SRE) datasets. All recordings are telephone speech sampled at 8 kHz. Since there are no time annotations in these corpora, we extracted speech segments using speech activity detection (SAD) on the basis of a time-delay neural network and statistics pooling from a Kaldi recipe. We added noises from 37 background noise recordings from the MUSAN corpus [21]. We used a set of 10,000 room impulse responses (RIRs) from the Simulated Room Impulse Response Database [22]. We also prepared real conversations from the CALLHOME corpus [23]. We divided the CALLHOME data into two subsets: an adaptation set of 250 recordings and a test set of 250 recordings, randomly split according to the recipe in [16].

Table 1. DER (%) of offline and blockwise online diarization methods on CALLHOME two-speaker data. Block size $W = 10s$ and context size $L = \infty$. Variables controlling attractor computation: by **Block** (versus per utterance), **Reorder** (to match previous block), **Average** (across blocks), **Shuffle** (embeddings across blocks).

	Train	Inference	Attractor computation				DER
			Block	Reorder	Average	Shuffle	
1	offline	offline	no				9.02
2	offline	causal	no				12.73
3	causal	causal	no				10.05
4	offline	causal	yes	no	no	no	24.24
5	causal	causal	yes	no	no	no	18.07
6	causal	causal	yes	yes	no	no	14.26
7	causal	causal	yes	yes	yes	no	12.91
8	causal	causal	yes	yes	yes	yes	11.82

The input audio features are 23-dimensional log-scaled Mel filterbanks with a 25-ms frame length and 10-ms frame shift. Each feature vector is concatenated with both the previous and following seven frames (for a total of 15 frames). We then subsample the concatenated features by a factor of ten. Consequently, (23×15) -dimensional input features sampled every 100 ms are fed into the Transformer encoder block for all EDA-EEND systems.

We used four stacked Transformer encoders with 256 attention units, containing four heads each. We first trained all three systems using simulation data with two speakers for 100 epochs, using the Adam optimizer with a learning rate schedule with 100,000 warm-up steps. We then finetuned the two-speaker models with simulated data with one to four speakers, for 100 epochs. Finally, we finetuned the models using CALLHOME adaptation data (250 recordings with two to seven speakers per call) and evaluated performance on CALLHOME test data (with two to six speakers per call). In Tables 1 and 2, we show results separately for subsets with different numbers of speakers (148 recordings with 2 speakers, 74 recordings with speakers, 20 recordings with 4 speakers).

We evaluated all systems by their diarization error rates (DER). As is standard, a tolerance of 250 ms when comparing hypothesized to reference speaker boundaries was used.

5. RESULTS

Table 1 shows the effect of various features of the BW-EDA-EEND algorithms. As described in Section 3, we trained the Transformer encoder to only attend to its left contexts and reuse the hidden states from the previous blocks (Train = causal) instead of using entire utterances to compute attention weights (Train = offline). During inference, we reuse the hidden states from the previous blocks (Inference = causal) or recompute them from scratch for every block (Inference = offline). As shown in Table 1, incremental training of the Transformer encoder improved accuracy since training and inference mechanisms now are consistent. For unlimited-latency BW-EDA-EEND, accuracy improved by 2.7% absolute (row 2 vs. 3 in Table 1) and for limited-latency BW-EDA-EEND accuracy improved by 6.2% absolute (row 4 vs. 5), when the algorithms are evaluated for two-speaker CALLHOME data using a block size of 10 seconds.

Next, during attractor inference in BW-EDA-EEND, we compare several heuristics to improve accuracy (rows 4-8 in Table 1). By reordering attractors, accuracy is improved by 3.8% absolute (row 5 vs. 6). Through attractor averaging, accuracy improves by 1.3% absolute (row 6 vs. 7). By shuffling embeddings across blocks, accuracy improves by 1.1% absolute (row 7 vs. 8).

Table 2. DERs (%) of offline and online diarization methods on simulated and CALLHOME data. Speaker number is estimated by the algorithm. As a baseline, the x-vector clustering result is included as reported in [16, 17, 18]. We use BW-EDA-EEND-UL when attractors are computed at the end of the utterances and BW-EDA-EEND-LL when attractors are computed block by block.

Model	Attractor computation	Embedding shuffling	Number of speakers							
			1	2	3	4	2	3	4	
			Simulated				CALLHOME			
Offline x-vector			37.42	7.74	11.46	22.45	15.45	18.01	22.68	
Offline EDA-EEND	by utterance	within utterance	0.27	4.18	9.66	14.21	9.02	13.78	20.69	
Online $W = 10, L = \infty$	by utterance	within block	0.28	4.22	11.17	21.04	10.05	16.59	25.50	
Online $W = 10, L = 1$	by utterance	within block	0.30	4.42	13.35	22.71	11.73	19.87	28.03	
Online $W = 10, L = \infty$	by block	within block	1.72	8.46	20.60	37.17	12.91	22.04	30.62	
Online $W = 10, L = 1$	by block	within block	1.85	10.99	22.23	36.72	16.84	26.01	28.91	
Online $W = 10, L = \infty$	by block	across blocks	1.03	6.10	12.58	19.17	11.82	18.30	25.93	
Online $W = 10, L = 1$	by block	across blocks	2.49	7.53	16.65	24.50	16.18	19.35	27.52	

We evaluated the proposed BW-EDA-EEND on simulation data with a variable number of speakers. Note that we trained the initial model with simulated mixtures of two speakers, then finetuned it with more simulation data, now mixing up to four speakers. The results for various context sizes are shown in Table 2 (left half). We use the algorithm of Table 1, row 3 to report accuracy for BW-EDA-EEND-UL (attractors are computed at ends of utterances) and the algorithm in Table 1, rows 7 and 8, for BW-EDA-EEND-LL (attractors are computed block-by-block). First, when attractors are computed at the ends of the utterances, accuracy degrades only moderately up to two speakers, using either unlimited or 10 seconds ($L = 1$) context. With more than two speakers in a conversation, the accuracy gap between online and offline increases, but BW-EDA-EEND-UL with unlimited left context still outperforms the baseline clustering-based system [16, 17, 18] for one to four speakers, or shows comparable accuracy with a single block of context.

When attractors are computed block-by-block, BW-EDA-EEND-LL does not perform well unless diarization embeddings are shuffled *across* blocks. This shows the importance of shuffling embeddings for the LSTM to learn order-invariance during encoding. When frame-level embeddings are shuffled across blocks, we see comparable accuracy to a offline clustering-based (x-vector) system with unlimited left context or a single block of context. We conclude that it is crucial that EDA receives embeddings from *all* the speakers to accurately compute attractors for all the speakers who have spoken so far. If only a subset of speakers are active in a given block, it seems that information about the previous speakers tends to be forgotten. In simulation data, we often observe only a subset of speakers active in the tail of conversation, hurting the accuracy when attractors are computed block-by-block (Fig. 2, top half).

Furthermore, we evaluated BW-EDA-EEND on real conversation data. For testing on real recordings, we finetuned the model first trained with simulation data, using the CALLHOME adaptation set. The results for various block sizes are shown in Table 2 (right half). Consistent with previous findings, BW-EDA-EEND-UL produces similar accuracy as that of offline EDA-EEND for up to two speakers when context is unlimited or 10 seconds. When attractors are computed block-by-block, BW-EDA-EEND-LL does not perform well if diarization embeddings are shuffled within blocks, but results in accuracy comparable to the clustering-based system when frame-level embeddings are shuffled across blocks, for up to three speakers and left context unlimited or 10 seconds ($L = 1$) context. Similar to simulation data, we often find that a subset of speakers is active only in part of a conversation, hurting the accuracy when attractors are computed block-by-block (Fig. 2, bottom half).

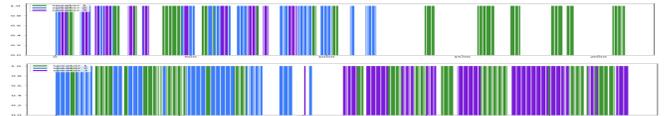


Fig. 2. Visualization of speaker turns for simulated (top) and CALLHOME data (bottom). There are three speakers in these examples.

6. CONCLUSIONS

We implemented two versions of a blockwise online variant of EDA-EEND, BW-EDA-EEND-UL (unlimited latency) and BW-EDA-EEND-LL (limited latency). Blockwise online processing is enabled by utilizing an incremental Transformer encoder that attends to only its left contexts and ignores its right contexts and uses block-level recurrence in the hidden states to carry over information between blocks, which makes algorithm complexity linear in time. BW-EDA-EEND-UL shows only moderate degradation of accuracy for up to two speakers using either unlimited or 10-second context, compared to offline EDA-EEND. BW-EDA-EEND-LL has accuracy comparable to an offline clustering-based system when frame-level embeddings are shuffled across blocks. Future algorithmic improvements should address the consistency of attractor direction and ordering over time, as blocks are processed incrementally.

We observe that multi-speaker training data as simulated in prior work [16] (and adopted here for compatibility) is not very realistic when compared to real conversational data, in terms of turn-taking behavior (as shown in Fig. 2), and suspect that this may limit the effectiveness of model training. In future work, we plan to modify the simulation algorithm to create more realistic meeting mixtures by adopting the recently proposed method that was used to create LibriCSS test data [24]. We also believe that test data for EEND needs to become more realistic, moving from mixtures of telephone speech channels (CALLHOME) to far-field recordings of multiple speakers speaking and interacting in the same room.

7. ACKNOWLEDGMENTS

We would like to thank our colleague Sundararajan Srinivasan, as well as Shinji Watanabe and other members of the 2020 JSALT workshop on “Speech Recognition and Diarization for Unsegmented Multi-talker Recordings with Speaker Overlaps” for help with data preparation and valuable discussions.

8. REFERENCES

- [1] C. Boeddeker, J. Heitkaemper, J. Schmalenstroeer, L. Drude, J. Heymann, and R. Haeb-Umbach, "Front-end processing for the CHiME-5 dinner party scenario," in *Proc. CHiME-5 Workshop*, Hyderabad, 2018.
- [2] N. Kanda, R. Ikeshita, S. Horiguchi, Y. Fujita, K. Nagamatsu, X. Wang, V. Manohar, N. E. Y. Soplin, M. Maciejewski, S.-J. Chen *et al.*, "The Hitachi/JHU CHiME-5 system: Advances in speech recognition for everyday home environments using multiple microphone arrays," in *Proc. CHiME-5 Workshop*, Hyderabad, 2018, pp. 6–10.
- [3] N. Kanda, Y. Fujita, S. Horiguchi, R. Ikeshita, K. Nagamatsu, and S. Watanabe, "Acoustic modeling for distant multi-talker speech recognition with single-and multi-channel branches," in *Proc. IEEE ICASSP*, 2019, pp. 6630–6634.
- [4] N. Kanda, S. Horiguchi, Y. Fujita, Y. Xue, K. Nagamatsu, and S. Watanabe, "Simultaneous speech recognition and speaker diarization for monaural dialogue recordings with target-speaker acoustic models," in *Proceedings IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 31–38.
- [5] G. Sell and D. Garcia-Romero, "Speaker diarization with PLDA i-vector scoring and unsupervised calibration," in *Proc. IEEE Spoken Language Technology Workshop (SLT)*, 2014, pp. 413–417.
- [6] D. Garcia-Romero, D. Snyder, G. Sell, D. Povey, and A. McCree, "Speaker diarization using deep neural network embeddings," in *Proc. IEEE ICASSP*, 2017, pp. 4930–4934.
- [7] Q. Li, F. L. Kreyssig, C. Zhang, and P. C. Woodland, "Discriminative neural clustering for speaker diarisation," *arXiv preprint arXiv:1910.09703*, 2019.
- [8] M. Maciejewski, D. Snyder, V. Manohar, N. Dehak, and S. Khudanpur, "Characterizing performance of speaker diarization systems on far-field speech using standard methods," in *Proc. IEEE ICASSP*, 2018, pp. 5244–5248.
- [9] D. Dimitriadis and P. Fousek, "Developing on-line speaker diarization system," in *Proc. Interspeech*, 2017, pp. 2739–2743.
- [10] Q. Wang, C. Downey, L. Wan, P. A. Mansfield, and I. L. Moreno, "Speaker diarization with LSTM," in *Proc. IEEE ICASSP*, 2018.
- [11] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Front-end factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2010.
- [12] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *Proc. IEEE ICASSP*, 2018, pp. 5329–5333.
- [13] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *Proc. IEEE ICASSP*, 2018, pp. 4879–4883.
- [14] A. Zhang, Q. Wang, Z. Zhu, J. Paisley, and C. Wang, "Fully supervised speaker diarization," in *Proc. IEEE ICASSP*, 2019, pp. 6301–6305.
- [15] L. Bullock, H. Bredin, and L. P. Garcia-Perera, "Overlap-aware diarization: Resegmentation using neural end-to-end overlapped speech detection," in *Proc. IEEE ICASSP*, 2020, pp. 7114–7118.
- [16] Y. Fujita, N. Kanda, S. Horiguchi, Y. Xue, K. Nagamatsu, and S. Watanabe, "End-to-end neural speaker diarization with self-attention," in *Proc. IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 296–303.
- [17] Y. Fujita, S. Watanabe, S. Horiguchi, Y. Xue, J. Shi, and K. Nagamatsu, "Neural speaker diarization with speaker-wise chain rule," *arXiv preprint arXiv:2006.01796*, 2020.
- [18] S. Horiguchi, Y. Fujita, S. Watanabe, Y. Xue, and K. Nagamatsu, "End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors," in *Proc. Interspeech*, 2020, pp. 269–273.
- [19] Z. Chen, Y. Luo, and N. Mesgarani, "Deep attractor network for single-microphone speaker separation," in *Proc. IEEE ICASSP*, 2017, pp. 246–250.
- [20] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-XL: Attentive language models beyond a fixed-length context," in *Proc. ACL*, 2019, pp. 2978–2988.
- [21] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [22] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *Proc. IEEE ICASSP*, 2017, pp. 5220–5224.
- [23] M. Przybocki and A. Martin, "2000 NIST Speaker Recognition Evaluation (LDC2001S97)," Linguistic Data Consortium, 2001.
- [24] Z. Chen, T. Yoshioka, L. Lu, T. Zhou, Z. Meng, Y. Luo, J. Wu, X. Xiao, and J. Li, "Continuous speech separation: Dataset and analysis," in *Proc. IEEE ICASSP*, 2020, pp. 7284–7288.