

Speakers Unembedded: Embedding-free Approach to Long-form Neural Diarization

Xiang Li, Vivek Govindan, Rohit Paturi, Sundararajan Srinivasan

AWS AI Labs

{xiangzai, vvekg, paturi, sundarsr}@amazon.com

Abstract

End-to-end neural diarization (EEND) models offer significant improvements over traditional embedding-based Speaker Diarization (SD) approaches but falls short on generalizing to long-form audio with large number of speakers. EEND-vector-clustering method mitigates this by combining local EEND with global clustering of speaker embeddings from local windows, but this requires an additional speaker embedding framework alongside the EEND module. In this paper, we propose a novel framework applying EEND both locally and globally for long-form audio without separate speaker embeddings. This approach achieves significant relative DER reduction of 13% and 10% over the conventional 1-pass EEND on Callhome American English and RT03-CTS datasets respectively and marginal improvements over EEND-vector-clustering without the need for additional speaker embeddings. Furthermore, we discuss the computational complexity of our proposed framework and explore strategies for reducing processing times.

Index Terms: speaker diarization, end-to-end diarization, spectral clustering

1. Introduction

Speaker diarization addresses the “who spoken when” problem by partitioning an audio stream containing multiple speakers into homogeneous segments associated with each speaker. Conventional diarization systems [1–6] typically consist of a cascade of several separate modules: voice activity detection to detect the speech frames, speaker embedding extraction to transform the speech segments into discriminative representations, and clustering to group speech regions by speaker identity. While effective for long-form audio with an arbitrary number of speakers, these cascaded multi-module approaches face challenges in handling overlapping speech and can suffer from error propagation across the modules.

To overcome the limitations of cascaded approaches, end-to-end neural diarization (EEND) was proposed in [7] which formulates speaker diarization as a frame-wise multi-label classification task with permutation invariant training [8]. EEND can naturally handle overlapping speech by allowing multiple speakers to be active simultaneously and is also fully supervised compared to the unsupervised clustering component of the cascaded approach. However, despite its theoretical promise, EEND and its variants like EEND-SA [9], EEND-EDA [10], etc have struggled to generalize to larger numbers of speakers and arbitrarily long conversations.

In order to apply EEND models to longer audios and larger number of speakers, recent works [11, 12] have proposed hybrid frameworks that integrate EEND with conventional clustering-based approaches. These methods leverage the strong diariza-

tion capability of EEND for speaker labeling over short local windows while performing global clustering on speaker embeddings computed across the local windows. This hybrid approach can handle both overlapping speech locally and long conversations with an arbitrary number of speakers globally. Most of the recent EEND improvements have focused on integrating additional embedding [11–17] or attractor modules [18–24], which requires specialized model architectures, loss functions and data requirements. Moreover, in some real-world scenarios, creating and storing speaker embeddings may need to be avoided where possible due to privacy considerations [25].

In this paper, we propose a novel embedding-free approach that doesn’t require any speaker embeddings and can still leverage the benefits of EEND and scale it to long-form audios with arbitrary number of speakers. We achieve this by utilizing a vanilla EEND model for both local diarization within the short local windows as well as global diarization across local windows, hence named local-global EEND. The proposed method consists of three steps: local EEND, global EEND, and clustering. In the local step, long audio is split into fixed-size windows, and EEND performs diarization within each window. The global step solves the inter-window label permutation by re-applying EEND to chunks formed by pairing speaker chunks across local windows. This generates pairwise speaker scores which are used to build an affinity matrix for the final clustering and global speaker labeling, without requiring any speaker embeddings.

The rest of the paper details the local-global EEND approach, experimental setup, results compared to the baselines, and a discussion on potential computation improvements for the global step.

2. Local-global EEND

2.1. Local EEND

Figure 1 shows the schematic diagram of the proposed embedding-free approach which can be divided into local EEND, global EEND and clustering steps.

The input audio is first split into W windows with a fixed window length. In each window i , frame-level acoustic features are extracted, denoted as $\mathbf{X}_i = \{\mathbf{x}_{i,t}\}_{t=1}^T$, $\mathbf{x}_{i,t} \in \mathbb{R}^F$ where t is the frame index, T is the total number of frames in a window and F is the feature dimension of Mel-filterbank features in this work. Speaker label $\mathbf{y}_{i,t} = \{y_{i,t,s}\}_{s=1}^{S_{local}}$ denotes speech activities for S_{local} speakers at frame t within window i and is defined as

$$y_{i,t,s} = \begin{cases} 0 & \text{(Speaker } s \text{ is inactive at } t) \\ 1 & \text{(Speaker } s \text{ is active at } t) \end{cases} \quad (1)$$

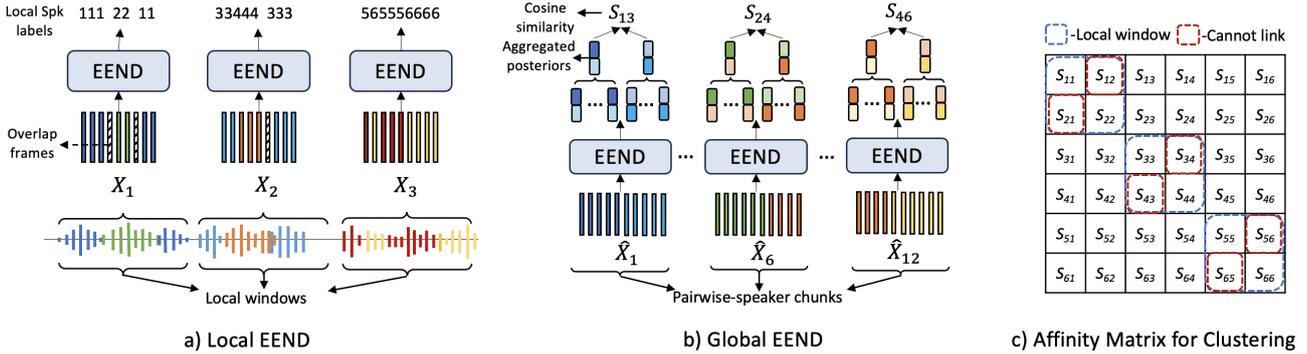


Figure 1: *Local-global EEND framework. This assumes 3 local windows with 2-speaker local EEND, i.e. $W=3$, $S_{local}=2$ resulting in $C=12$ pairwise-speaker chunks for global EEND.*

The local EEND estimates frame-wise posteriors $P(y_{i,t,s}|\mathbf{X}_i)$ in each window using a vanilla EEND model. These posteriors are binarized using a threshold Th_{local} and median filtered [9] to obtain the local speaker labels $\mathbf{y}_{i,t}$.

2.2. Global EEND

In order to perform global SD, the global EEND step computes the speaker similarities across the local windows using the same EEND model. In order to compute these, the overlapping speaker frames within each local window are first filtered out and the remaining frames of each speaker in a window are paired with the frames of speakers in subsequent windows, resulting in new chunks $\{\hat{\mathbf{X}}_i\}_{i=1}^C$

$$\hat{\mathbf{X}}_i = \text{concat}(\mathbf{x}_{j,t}|t = \{m_{j,s}\}_{s=1}^M, \mathbf{x}_{k,t}|t = \{n_{k,s}\}_{s=1}^N) \quad (1)$$

where $\mathbf{x}_{j,t}$ and $\mathbf{x}_{k,t}$ represent frame-level acoustic features of window j and window k ($j \neq k$), respectively. $\{m_{j,s}\}_{s=1}^M$ represents the M frame indices corresponding to speaker m in window j and $\{n_{k,s}\}_{s=1}^N$ represents the N frame indices corresponding to speaker n in window k . C is the total number of pairwise-speaker chunks processed by global EEND, where

$$C \leq W \times (W - 1) / 2 \times S_{local}^2 \quad (2)$$

In the case where a speaker has limited or no non-overlapping frames, we leverage the overlapping frames similarly to the EEND-vector clustering approach.

EEND is applied to $\hat{\mathbf{X}}_i$ to generate inter-window frame-level speaker posteriors

$$[\mathbf{z}_1, \dots, \mathbf{z}_M, \mathbf{z}_{M+1}, \dots, \mathbf{z}_{M+N}] = \text{EEND}(\hat{\mathbf{X}}_i) \quad (3)$$

where $\{\mathbf{z}_t\}_{t=1}^{M+N}$, $\mathbf{z}_t \in \mathbb{R}^{S_{local}}$ are the inter-window frame-level posteriors. $[\mathbf{z}_1, \dots, \mathbf{z}_M]$ and $[\mathbf{z}_{M+1}, \dots, \mathbf{z}_{M+N}]$ are the posteriors corresponding to the M frames of speaker m and N frames of speaker n respectively. This process is repeated on every speaker pair across local windows as shown in Figure 1.

2.3. Embedding-free clustering

The frame-wise posteriors $\{\mathbf{z}_t\}_{t=1}^{M+N}$ are aggregated on frames belonging to the same speaker, resulting in speaker-level posteriors $\bar{\mathbf{z}}_m$ and $\bar{\mathbf{z}}_n$. Pairwise-speaker similarity S_{mn} is then calculated as

$$\bar{\mathbf{z}}_m = \text{mean}([\mathbf{z}_1, \dots, \mathbf{z}_M]) \quad (4)$$

$$\bar{\mathbf{z}}_n = \text{mean}([\mathbf{z}_{M+1}, \dots, \mathbf{z}_N]) \quad (5)$$

$$S_{mn} = \text{cosine_similarity}(\bar{\mathbf{z}}_m, \bar{\mathbf{z}}_n) \quad (6)$$

Each S_{mn} is an entry of the affinity matrix $\mathbf{S} \in \mathbb{R}^{S_{Global} \times S_{Global}}$ which will be used for the final clustering, where S_{Global} is the sum of number of speakers detected in each local window with the local EEND, upper bound of which will be $W \times S_{local}$.

In order to enhance the clustering performance as well as to save on additional computations, we incorporate cannot-link constraints among different speakers identified within the same local window obtained in the local EEND step. This constraint is enforced by assigning a speaker similarity of 0 between local speaker pairs. Spectral clustering is then employed to group the speaker frames into D speaker sets using the max eigengap heuristic similar to [6, 12].

3. Experiments

In this section, we go over the datasets used, model architecture, settings and techniques followed for efficient inference.

3.1. Data and metrics

For training the EEND model, we generate simulated mixtures by mixing Switchboard-2 (Phase I & II & III), Switchboard Cellular (Part 1 & 2), and the NIST Speaker Recognition Evaluation (2004 & 2005 & 2006 & 2008) with MUSAN corpus [26], following the data generation procedure in [7]. Mixtures with up to 3 speakers were created, with $\beta = [2, 2, 9]$ for mixture with 1, 2 and 3 speakers, respectively.

For model adaptation and evaluation, real telephone conversation dataset CALLHOME [27], i.e., NIST SRE2000 (LDC2001S97, Disk-8) is used. It is widely used as the benchmark for existing EEND-based approaches. The CALLHOME dataset contains 500 sessions, each with 2 to 6 speakers. There are mostly two dominant speakers in each conversation. We split the data into two subsets according to [10] for adaptation (CALLHOME1) and evaluation (CALLHOME2).

As local-global EEND framework is designed for dealing with long conversations, to showcase the effectiveness of this framework, evaluations on other benchmarks with longer audios are reported as well, such as CALLHOME American English (CHAE) [28] and RT03-CTS [29] which have an average duration of 30 and 10 minutes respectively. We use the official eval splits for evaluation on these datasets.

For evaluation metrics, we use the standard Diarization Error Rate (DER) [30] with a collar tolerance of 250ms and included the overlapping speech segments while scoring.

Table 1: Effect of sequence concatenation on CALLHOME2. $EEND_{vanilla}$ is the EEND model adapted on unmodified utterances. $EEND_{concat}$ is the EEND model adapted on speaker concatenated sequences described in 3.3. $EEND_{vanilla} + EEND_{concat}$ uses $EEND_{vanilla}$ for the local step and $EEND_{concat}$ for the global step of our local-global approach.

System	Model	# of speakers in a session					
		2	3	4	5	6	all
1-pass EEND	$EEND_{vanilla}$	7.53	14.91	-	-	-	-
	$EEND_{concat}$	7.36	17.74	-	-	-	-
Local-global EEND	$EEND_{vanilla}$	7.99	12.21	16.39	17.10	26.12	12.48
	$EEND_{concat}$	7.29	11.85	17.83	15.76	22.38	12.16
	$EEND_{vanilla} + EEND_{concat}$	7.66	11.67	16.03	17.56	23.71	12.45

3.2. EEND model settings

To ensure a fair comparison with the existing hybrid baseline, we adopt the front-end configuration from EEND-vector-clustering [12]. This involves the extraction of 23-dimensional log-Mel-filterbank features, utilizing a frame length of 25ms and a frame shift of 10ms. The window size T is set at 300 (=30s) for both training and adaptation. The EEND architecture consists of 6 stacked self-attention-based Transformer layers, featuring eight attention heads and a hidden size of 256. This aligns with the configuration employed in [12]. In each window, the EEND model estimates the posteriors for up to 3 speakers.

During both training and adaptation, we employ the Adam optimizer [31] alongside the Noam scheduler [32], incorporating 150,000 warm-up steps for training. For adaptation, a fixed learning rate of 1×10^{-5} is utilized. Both training and adaptation phases span 100 epochs.

3.3. Sequence concatenation during adaptation

In the global EEND step, we generate the chunk-level input by concatenating frame-level acoustic features between every pair of speakers across local windows. This process results in a new pairwise-speaker sequence that has not been encountered in either the training or adaptation data. To enhance EEND’s generalization to this new input format, we incorporate this data generation procedure during adaptation. First, the frame-level acoustic features from the same speaker in each utterance are aggregated into several blocks after discarding the overlapping speech frames. Subsequently, every two blocks are concatenated to generate a new input utterance. We reformat the data using this technique for half of the samples in each batch.

3.4. Efficiency improvement for inference

Real Time Factor (RTF) is a criteria used to measure the efficiency of SD systems. It is calculated by dividing the time taken by the SD system by the total duration of the spoken audio.

In the global step during inference, each speaker within a local window is paired with every speaker in subsequent local windows, resulting in a computational complexity of $O(n^2)$. As illustrated in Figure 1, if there are 3 local windows, each containing 2 speakers, there will be 12 inference calls in the global EEND step. To enhance GPU efficiency and reduce RTF, we propose batching multiple inference requests together. Additionally, we explore different number of random frames ($N = 128, 64, 32, 16$) to minimize the number of frames required for each speaker during global EEND inference, thereby reducing computational load.

4. Results

In this section, we present the outcomes of our experiments, beginning with an evaluation of the impact of sequence concatenation in the adaptation process. Subsequently, we compare the proposed approach with existing baselines, utilizing both oracle and estimated speaker counts during clustering. Our analysis extends to additional benchmarks, such as CALLHOME American English (CHAE) and RTCTS. Finally, we delve into an examination of the efficiency improvements in global EEND inference.

4.1. Effect of sequence concatenation

We explored two types of input data formats during adaptation, resulting in $EEND_{vanilla}$ and $EEND_{concat}$. The former denotes the model adapted with original input sequences typically used in EEND model adaptation, while the latter is adapted with a combination of original sequences and concatenated sequences, as described in Section 3.3. Table 1 presents a comparison between the local-global EEND and 1-pass EEND, utilizing the different adaptation techniques. We used oracle speaker information during global clustering and selected a binarization threshold between 0.3 to 0.7 that produced the best DER on the validation set for both 1-pass and local EEND. We only evaluated 1-pass EEND on 2,3 speaker sessions since it was trained to only detect a maximum of 3 speakers.

Local-global EEND outperforms 1-pass EEND on $EEND_{vanilla}$ by 18% in the 3-speaker session but shows a marginal performance degradation in the 2-speaker session whereas $EEND_{concat}$ outperforms the best 1-pass EEND by 3% and 21% in 2-speaker and 3-speaker sessions, respectively.

For 1-pass EEND, $EEND_{concat}$ only marginally benefits the 2-speaker session but not the 3-speaker one, as expected due to the data mismatch between adaptation and evaluation, where concatenation during adaptation only occurs on pairwise speakers. When evaluating more speakers for local-global EEND, $EEND_{concat}$ consistently performs better than $EEND_{vanilla}$, except for the 4-speaker session.

In order to exactly match the local and global input conditions, we also attempted to apply $EEND_{vanilla}$ in local EEND and $EEND_{concat}$ in global EEND. This further improved performance in 3 and 4 speaker sessions but $EEND_{concat}$ for both steps achieves the best overall DER across all speaker sessions.

4.2. Comparison with other baselines

The performance of local-global EEND is compared with other baselines in Table 2 and Table 3, with oracle and esti-

Table 2: Comparison with other baselines with oracle number of speakers on CALLHOME2. The best scores are **bolded**.

System	# of speakers in a session					
	2	3	4	5	6	all
x-vector-clustering [10]	8.93	19.01	24.48	32.14	34.95	18.98
EDA-EEND [10]	8.35	13.20	21.71	33.00	41.07	15.43
EEND-vector-clust. (T=30s) [12]	8.08	11.27	15.01	23.14	26.56	12.22
Local-global EEND	7.29	11.85	17.83	15.76	22.38	12.16

Table 3: Comparison with other baselines with estimated number of speakers on CALLHOME2. The best scores are **bolded** and the second best are underlined.

System	# of speakers in a session					
	2	3	4	5	6	all
x-vector-clustering [10]	15.54	18.01	22.68	31.40	34.27	19.43
EDA-EEND [10]	8.50	13.24	21.46	33.16	40.29	15.29
EEND-vector-clust. (T=30s) [12]	7.96	<u>11.93</u>	<u>16.38</u>	<u>21.21</u>	23.10	12.49
EEND-EDA-local-global [18]	7.11	11.88	14.37	25.95	21.95	11.84
Local-global EEND	<u>7.51</u>	12.20	17.88	16.01	<u>22.35</u>	<u>12.20</u>

ated numbers of speakers respectively. Results for EEND-vector-clustering with a window size of 30s are extracted from their paper, and the outcomes for local-global EEND with $EEND_{concat}$ are reported.

In the case of oracle numbers of speakers, local-global EEND outperforms EEND-vector-clustering [12] in all sessions except for 3 and 4-speaker sessions, with only a marginal degradation in the 3-speaker session. Particularly noteworthy is the substantial improvement of 32% and 15.7% in 5- and 6-speaker sessions, respectively. The decline in the 4-speaker session may be attributed to the sub-optimal window size, a phenomenon observed in the EEND-vector-clustering paper as well, suggesting that different window sizes may result in significant performance variations.

Table 3 presents results with estimated numbers of speakers, introducing another strong baseline, EEND-EDA-local-global [18]. Compared to EEND-vector-clustering, local-global EEND achieves superior performance across nearly all sessions, with similar exceptions in the 4-speaker session. In comparison to EEND-EDA-local-global, local-global EEND exhibits slightly inferior performance in the {2, 3, 6}-speaker sessions but significantly outperforms in the 5-speaker session. This discrepancy could stem from differences in training data volume where the local-global EEND is trained with only up to 3-speaker mixtures, whereas EEND-EDA-local-global is trained on a larger training data including 100,000 additional 4-speaker mixtures. Consequently, EEND-EDA-local-global achieves the best DER in the 4-speaker session, aligning with this matched training scenario.

4.3. Performance on other benchmarks

To provide a comprehensive evaluation of the local-global EEND system, we extend our analysis to other well-established diarization benchmarks, namely CALLHOME American English (CHAE) and RTCTS. These datasets feature longer-duration audios, offering insights into the efficacy of local-global diarization systems. Binarization is carried out using a fixed threshold ($Th_{local}=0.5$) to obtain diarization results, and clustering is performed using the estimated number of speakers.

As depicted in Table 4, local-global EEND demonstrates

Table 4: DER (%) on long-form audio datasets ($Th_{local}=0.5$).

System	Dataset		
	CHAE test	CHAE 109	RTCTS test
1-pass EEND	5.95	7.25	5.69
Local-global EEND	5.20	7.02	5.14

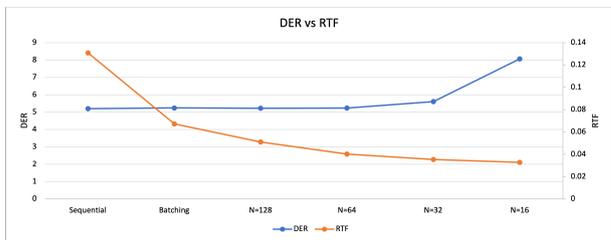


Figure 2: RTF vs DER with different strategies on efficiency improvement, including batching the inferences and minimizing the number of frames required for each speaker. N indicates a subset of N random frames.

superior performance compared to 1-pass EEND across various datasets. Specifically, it outperforms 1-pass EEND by 12.7%, 9.7% and 3%, on the CHAE test set, RTCTS test set and CHAE 109, respectively.

4.4. Inference efficiency improvements

Figure 2 shows the results on CHAE test set with different strategies to improve inference efficiency. All the experiments are performed on NVIDIA A10G GPU on AWS cloud (G5-2xLarge). Moving from sequential inference to batching (batching 500 chunks), the RTF is reduced by 50%. Further RTF reduction is gained from reducing computation by selection a subset of frames ($N=128, 64, 32, 16$) randomly for each speaker in global EEND. A subset of 64 frames can produce a desirable RTF reduction by nearly 70% with no impact on DER. Regarding the computational cost versus the input audio length, the local-global EEND produces the RTF of $\{7.3e^{-3}, 1.5e^{-2}, 2.2e^{-2}, 5.0e^{-2}\}$ for the audio length of $\{5, 10, 15, 30\}mins$, respectively.

5. Conclusion

This paper introduces a novel embedding-free diarization methodology that employs EEND in both local and global steps. The global clustering is accomplished without the need for speaker embeddings, utilizing EEND on concatenated pairwise speaker features across local windows to derive the pairwise speaker similarities. This approach achieves significant relative DER reduction of 13% and 10% over the conventional 1-pass EEND on CHAE and RT03-CTS datasets respectively and even offers a marginal 3% relative DER reduction over EEND-vector-clustering without the need for additional speaker embeddings or loss functions. The paper also includes a discussion on the computational complexity of the global EEND step and explores strategies for reducing the processing times. By batching multiple chunk-level inferences and minimizing the number of frames required for each speaker, the RTF can be reduced by nearly 70% without the impact on diarization performance.

6. References

- [1] X. Anguera, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, and O. Vinyals, "Speaker diarization: A review of recent research," *IEEE Transactions on audio, speech, and language processing*, vol. 20, no. 2, pp. 356–370, 2012.
- [2] G. Sell and D. Garcia-Romero, "Speaker diarization with plda i-vector scoring and unsupervised calibration," in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 413–417.
- [3] D. Snyder, P. Ghahremani, D. Povey, D. Garcia-Romero, Y. Carmiel, and S. Khudanpur, "Deep neural network-based speaker embeddings for end-to-end speaker verification," in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 165–170.
- [4] D. Garcia-Romero, D. Snyder, G. Sell, D. Povey, and A. McCree, "Speaker diarization using deep neural network embeddings," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 4930–4934.
- [5] G. Sell, D. Snyder, A. McCree, D. Garcia-Romero, J. Villalba, M. Maciejewski, V. Manohar, N. Dehak, D. Povey, S. Watanabe *et al.*, "Diarization is hard: Some experiences and lessons learned for the jhu team in the inaugural dihard challenge," in *Interspeech*, 2018, pp. 2808–2812.
- [6] Q. Wang, C. Downey, L. Wan, P. A. Mansfield, and I. L. Moreno, "Speaker diarization with lstm," in *2018 IEEE International conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5239–5243.
- [7] Y. Fujita, N. Kanda, S. Horiguchi, K. Nagamatsu, and S. Watanabe, "End-to-end neural speaker diarization with permutation-free objectives," in *Interspeech*, 2019, pp. 4300–4304.
- [8] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 241–245.
- [9] Y. Fujita, N. Kanda, S. Horiguchi, Y. Xue, K. Nagamatsu, and S. Watanabe, "End-to-end neural speaker diarization with self-attention," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 296–303.
- [10] 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 *Interspeech*, 2020, pp. 269–273.
- [11] K. Kinoshita, M. Delcroix, and N. Tawara, "Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7198–7202.
- [12] —, "Advances in integration of end-to-end neural and clustering-based diarization for real conversational speech," in *Interspeech*, 2021, pp. 3565–3569.
- [13] N. Zeghidour, O. Teboul, and D. Grangier, "Dive: End-to-end speech diarization via iterative speaker embedding," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 702–709.
- [14] Y. Yu, D. Park, and H. K. Kim, "Auxiliary loss of transformer with residual connection for end-to-end speaker diarization," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 8377–8381.
- [15] I. Fung, L. Samarakoon, and S. J. Broughton, "Robust end-to-end diarization with domain adaptive training and multi-task learning," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–7.
- [16] C. Wang, J. Li, X. Fang, J. Kang, and Y. Li, "End-to-end neural speaker diarization with absolute speaker loss," in *Interspeech*, 2023, pp. 3577–3581.
- [17] A. Plaquet and H. Bredin, "Powerset multi-class cross entropy loss for neural speaker diarization," in *Interspeech*, 2023, pp. 3222–3226.
- [18] S. Horiguchi, S. Watanabe, P. Garcia, Y. Xue, Y. Takashima, and Y. Kawaguchi, "Towards neural diarization for unlimited numbers of speakers using global and local attractors," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 98–105.
- [19] M. Rybicka, J. Villalba, N. Dehak, and K. Kowalczyk, "End-to-end neural speaker diarization with an iterative refinement of non-autoregressive attention-based attractors," in *Interspeech*, 2022, pp. 5090–5094.
- [20] Y. Fujita, T. Komatsu, R. Scheibler, Y. Kida, and T. Ogawa, "Neural diarization with non-autoregressive intermediate attractors," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [21] F. Hao, X. Li, and C. Zheng, "End-to-end neural speaker diarization with an iterative adaptive attractor estimation," *Neural Networks*, vol. 166, pp. 566–578, 2023.
- [22] Z. Chen, B. Han, S. Wang, and Y. Qian, "Attention-based encoder-decoder network for end-to-end neural speaker diarization with target speaker attractor," in *Interspeech*, 2023, pp. 3552–3556.
- [23] L. Samarakoon, S. J. Broughton, M. Härkönen, and I. Fung, "Transformer attractors for robust and efficient end-to-end neural diarization," in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [24] F. Landini, M. Diez, T. Stafylakis, and L. Burget, "Diaper: End-to-end neural diarization with perceiver-based attractors," *arXiv preprint arXiv:2312.04324*, 2023.
- [25] F. Teixeira, A. Abad, B. Raj, and I. Trancoso, "Privacy-oriented manipulation of speaker representations," *arXiv preprint arXiv:2310.06652*, 2023.
- [26] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [27] M. Przybocki and A. Martin, "2000 nist speaker recognition evaluation (ldc2001s97)," *Philadelphia, New Jersey: Linguistic Data Consortium*, 2001.
- [28] A. D. G. Canavan and G. Zipperlen, "Callhome american english speech ldc97s42," *Web Download. Philadelphia: Linguistic Data Consortium*, 1997.
- [29] J. G. Fiscus, G. Doddington, A. Le, G. Sanders, M. Przybocki, and D. Pallett, "nist rich transcription evaluation data ldc2007s10," *Web Download. Philadelphia: Linguistic Data Consortium*, 2007.
- [30] NIST, "The 2009 (rt-09) rich transcription meeting recognition evaluation plan," <http://www.itl.nist.gov/iad/mig/tests/rt/2009/docs/rt09-meeting-eval-plan-v2.pdf>, 2009.
- [31] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations (ICLR)*, 2015.
- [32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.