



Attention-based Encoder-Decoder Network for End-to-End Neural Speaker Diarization with Target Speaker Attractor

Zhengyang Chen¹, Bing Han¹, Shuai Wang², Yanmin Qian^{1†}

¹MoE Key Lab of Artificial Intelligence, AI Institute
X-LANCE Lab, Department of Computer Science and Engineering, Shanghai Jiao Tong University
²Shenzhen Research Institute of Big Data, Shenzhen, China
{zhengyang.chen, hanbing97, yanminqian}@sjtu.edu.cn, wangshuai@cuhk.edu.cn

Abstract

This paper proposes a novel Attention-based Encoder-Decoder network for End-to-End Neural speaker Diarization (AED-EEND). In AED-EEND system, we incorporate the target speaker enrollment information used in target speaker voice activity detection (TS-VAD) to calculate the attractor, which can mitigate the speaker permutation problem and facilitate easier model convergence. In the training process, we propose a teacher-forcing strategy to obtain the enrollment information using the ground-truth label. Furthermore, we propose three heuristic decoding methods to identify the enrollment area for each speaker during the evaluation process. Additionally, we enhance the attractor calculation network LSTM used in the end-to-end encoder-decoder based attractor calculation (EEND-EDA) system by incorporating an attention-based model. By utilizing such an attention-based attractor decoder, our proposed AED-EEND system outperforms both the EEND-EDA and TS-VAD systems with only 0.5s of enrollment data.

Index Terms: End-to-End speaker diarization, EEND-EDA, TS-VAD, Target Speaker Attractor

1. Introduction

The speaker diarization task is defined to address the “Who spoke when?” problem. Traditional speaker diarization systems [1, 2] usually use a stage-wise paradigm: 1) A speaker embedding extractor is adopted to get the speaker embedding for each segment. 2) A clustering algorithm [3, 4, 5] merges segments based on the speaker embedding similarity and assigns the speaker label to each segment. 3) Optionally, some compensation algorithm like Variational-Bayesian refinement [6, 3, 7] is used to calibrate the clustering results. However, the stage-wise system can not deal with the overlap problem due to the frame-wise one-to-one assignment of the clustering algorithm.

To address this issue, many end-to-end methods have been proposed in recent years. For example, EEND [8, 9] models the diarization task as a multi-class classification problem, allowing frames to be classified into multiple speakers. However, a limitation of EEND is its fixed number of output heads, which makes it difficult to handle varying numbers of speakers. To address this limitation, Horiguchi et al. [10] proposed EEND-EDA, which uses an LSTM encoder-decoder to predict the attractor for each speaker. The number of attractors can be flexible, allowing the system to generalize to sessions where the

number of speakers varies. Fujita et al. [11] and Takashima et al. [12] have also proposed methods that sequentially output diarization results for each speaker using a chain rule. However, all of these methods suffer from the speaker permutation problem [13], making system training challenging.

The target-speaker voice activity detection (TS-VAD) [14, 15] system has attracted significant attention due to its outstanding performance in speaker diarization challenges [15, 16]. This has led to the development of various TS-VAD variants [17, 18, 19, 20] by different researchers. In the TS-VAD system, a separate diarization system first predicts the single-speaker speaking area for each speaker. Subsequently, the speaker embeddings extracted for each speaker are concatenated with the acoustic features and fed into the diarization system. Because the TS-VAD system uses pre-acquired speaker embeddings, it does not suffer from the speaker permutation problem.

In this paper, we propose an advanced attention-based encoder-decoder network for end-to-end neural speaker diarization (AED-EEND) by integrating the target speaker enrollment information used in the TS-VAD system into the original EEND-EDA system. Our system overcomes the limitations of both TS-VAD and EEND-EDA while leveraging their advantages. Unlike the TS-VAD system, which requires enrollment information from an external diarization system and speaker embedding extractor, our system obtains the enrollment information directly from the AED-EEND system, thereby making our system fully end-to-end. Additionally, we replace the LSTM-based model used in EEND-EDA to generate attractors with an attention-based decoder, which has shown to be more effective. More importantly, our AED-EEND system only uses an attention-based encoder-decoder network, which is much simpler than the EEND-EDA and TS-VAD systems.

Besides, we propose a teacher forcing strategy for the training process that uses ground-truth labels to acquire enrollment information. For evaluation, we propose three heuristic decoding methods to output the diarization results for each speaker iteratively. Our proposed attention-based attractor decoder is particularly noteworthy as it produces robust results with only 0.5 seconds of enrollment speech, which significantly outperforms both the EEND-EDA and TS-VAD systems on the CALLHOME evaluation set.

2. Method

2.1. Attention-based Encoder-Decoder Network for EEND

In this section, we introduce the architecture of our proposed AED-EEND system. The overall architecture of our system is shown in the upper part of Figure 1. The AED-EEND system contains two parts: an embedding encoder and an attractor decoder. The encoder-decoder architecture is very sim-

[†]Corresponding Author

This work was supported in part by China STI 2030-Major Projects under Grant No. 2021ZD0201500, in part by China NSFC projects under Grants 62122050 and 62071288, and in part by Shanghai Municipal Science and Technology Major Project under Grant 2021SHZDZX0102.

ilar to the original transformer model proposed in [21] except that we neglect the positional encoding in our system. The embedding encoder accepts the acoustic feature sequence $X = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T] \in \mathbb{R}^{T \times F}$ as input, and outputs the frame level speaker embedding sequence $E = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_T] \in \mathbb{R}^{T \times D}$. The attractor decoder takes the enrollment sequence $E_{\text{enroll}} = [\mathbf{e}_{\text{non}}, \mathbf{e}_{\text{sgl}}, \mathbf{e}_{\text{ovl}}, \mathbf{e}_{\text{spk}_1}, \dots, \mathbf{e}_{\text{spk}_S}] \in \mathbb{R}^{(S+3) \times D}$ as input, where \mathbf{e}_{non} , \mathbf{e}_{sgl} and \mathbf{e}_{ovl} correspond to the enrollment embeddings for non-speech, single-speaker speaking speech and overlap speech, respectively, and $\mathbf{e}_{\text{spk}_i}$ corresponds to the enrollment embedding for i -th speaker in the recording. The enrollment embedding sequence E_{enroll} is fed into the attractor decoder to get the corresponding attractor sequence $A = [\mathbf{a}_{\text{non}}, \mathbf{a}_{\text{sgl}}, \mathbf{a}_{\text{ovl}}, \mathbf{a}_{\text{spk}_1}, \dots, \mathbf{a}_{\text{spk}_S}] \in \mathbb{R}^{(S+3) \times D}$.

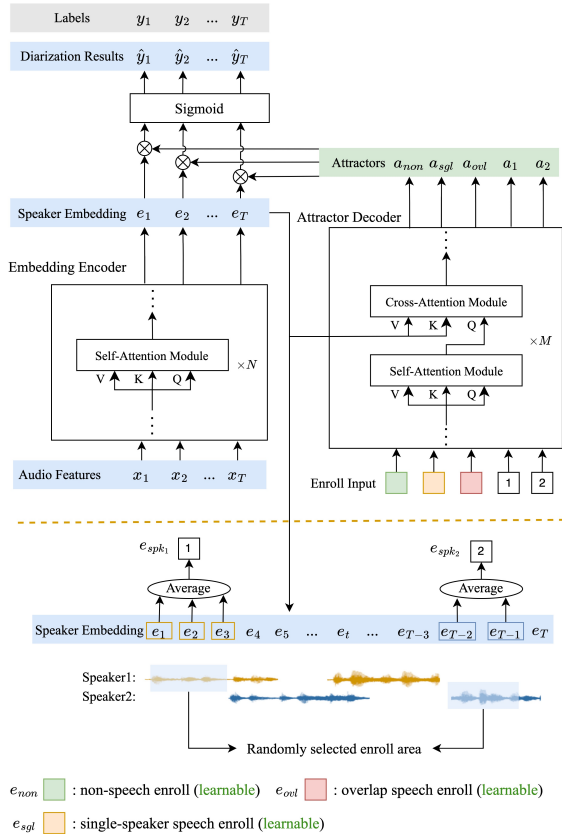


Figure 1: *AED-EEND system architecture when the speaker number is two.*

In contrast to the embedding encoder that solely incorporates the self-attention (SA) module, the attractor decoder is equipped with both the self-attention (SA) module and the cross-attention (CA) module. The SA module in the attractor decoder enables the enrollment inputs to attend to one another, whereas the CA module facilitates incorporating additional relevant information from the embedding sequence into the enrollment embedding. Specifically, the output from the SA module serves as the query, while the speaker embedding sequence output from the embedding encoder is treated as the key and value. This mechanism is expected to enhance the quality of the enrollment embedding.

Unlike the EEND-EDA architecture [10], which only generates attractors for the speakers present in the recording, our proposed AED-EEND system also outputs attractors for three

distinct speech types in the recording. These speech types include the single-speaker, speaker-overlap, and non-speech regions. The resulting attractors are then used to compute frame-level posterior probabilities for each speaker and speech type by performing a dot product operation with the corresponding embedding sequence.

$$\hat{Y} = \sigma(AE^T) \in (0, 1)^{(S+3) \times T} \quad (1)$$

where $\sigma(\cdot)$ is the element-wise sigmoid function. Here, we denote the posterior probabilities at each frame as $\hat{y}_t = [\hat{y}_t^{\text{non}}, \hat{y}_t^{\text{sgl}}, \hat{y}_t^{\text{ovl}}, \hat{y}_t^1, \dots, \hat{y}_t^S] \in (0, 1)^{(S+3)}$, and denote the ground truth label at each frame for speaker and speech activity as $y_t = [y_t^{\text{non}}, y_t^{\text{sgl}}, y_t^{\text{ovl}}, y_t^1, \dots, y_t^S] \in \{0, 1\}^{(S+3)}$. The label 1 corresponds to the specific speaker or specific kind of speech existing in the corresponding frame, and vice versa. We calculate the loss by averaging the cross entropy between the prediction and ground truth for each speaker or speech type at each frame:

$$\mathcal{L} = \frac{1}{T(S+3)} \sum_{t=1}^T \sum_{s \in \mathbb{S}} [-y_t^s \log \hat{y}_t^s - (1 - y_t^s) \log (1 - \hat{y}_t^s)] \quad (2)$$

where $\mathbb{S} = \{\text{non}, \text{sgl}, \text{ovl}, 1, \dots, S\}$.

2.2. System Training with Teacher Forcing Enrollment

Similar to the TS-VAD method [14], our proposed AED-EEND system also requires enrollment embedding for each speaker, but we obtain it directly from the embedding sequence E , eliminating the need for a pre-trained speaker embedding extraction system. As depicted in the lower part of Figure 1, given the single-speaker speaking area of each speaker in the recording, we can obtain the corresponding embeddings and average them to obtain the enrollment embedding for each speaker.

In the training of automatic speech recognition (ASR) systems, the teacher forcing strategy [22] is a widely adopted technique that leverages ground-truth input sequences to predict subsequent outputs during the training process. This strategy is known to improve training stability and efficiency. In the training process of our proposed AED-EEND, we similarly employ the teacher-forcing strategy. Specifically, we obtain the single-speaker speaking area for each speaker from the ground-truth label and randomly select a consecutive area with an enrollment length L_{enroll} as the enrollment area.

In contrast to speaker enrollment embeddings, which are obtained from the embedding sequence E , we directly set the enrollment embeddings for non-speech, single-speaker speech, and overlapping speech as three learnable embeddings.

2.3. System Inference by Iterative Decoding

In this section, we propose an iterative decoding method for AED-EEND in the evaluation process, and the pipeline is shown in algorithm 1. The first step is to predict the results for single-speaker, overlap, and non-speech with learnable embeddings. Then, the prediction for each speaker is obtained iteratively. At each iteration, we first find the un-predicted single-speaker speaking area I and then split the area into some time-continuous segments. If the longest segment is shorter than a pre-defined threshold L_{stop} , we will stop the decoding. To get the prediction for each speaker in one iteration, we propose three heuristic strategies to enlarge the probability that there is only one speaker in the enrollment area:

Algorithm 1: Iterative Decoding Pipeline

```
Data:  $I = [t_1, t_2, \dots]$ : frame indexes list for embedding in  $E$   
 $C = [I_1, I_2, \dots, I_K]$ :  $K$  time-continuous segments  
from  $I$ , and the indexes in  $I_k$  are sorted in order  
// get the active frame indexes list for three kinds of speeches  
1  $I_{\text{non}}, I_{\text{sgl}}, I_{\text{ovl}} = \text{AED-EEND}(\mathbf{e}_{\text{non}}, \mathbf{e}_{\text{sgl}}, \mathbf{e}_{\text{ovl}})$   
// active frame indexes list for speakers; enroll embedding set  
2  $I_{\text{spk}} = []$ ;  $\mathcal{E} = \{\}$   
3 while True do  
4 // get the un-predicted single-speaker area  
 $I = I_{\text{sgl}} - I_{\text{sgl}} \cap I_{\text{spk}}$   
5 // get segments with length  $\geq L_{\text{Enroll}}$   
 $C' = [I'_1, I'_2, \dots, I'_{K'}] = \text{filter\_segs}(C, L_{\text{Enroll}})$   
6 // get the segment with the longest length  
 $I_{\text{longest}} = \text{get\_longest\_seg}(C)$   
7  $C'.\text{add}(I_{\text{longest}})$   
8  $L_{\text{Enroll\_tmp}} = \min(\text{length}(I_{\text{longest}}), L_{\text{Enroll}})$   
9 if  $L_{\text{stop}} > \text{length}(I_{\text{longest}})$  then  
10 | break;  
11 case Init-Decode do  
12 |  $I_{\text{enroll}} = [t_1, t_2, \dots, t_{L_{\text{Enroll\_tmp}}}] \in I'_1$   
13 case Rand-Decode do  
14 | // randomly select a segment from  $C'$   
 $I'_k = \text{random\_select\_seg}(C')$   
15 | // randomly select a consecutive sub-segment with length  
 $L_{\text{Enroll\_tmp}}$   
 $I_{\text{enroll}} = \text{random\_select\_sub\_seg}(I'_k, L_{\text{Enroll\_tmp}})$   
16 case SC-Decode do  
17 | // The index in  $I$  is clustered based on corresponding  
embedding  
 $[I_1^{\text{cls}}, I_2^{\text{cls}}, \dots] = \text{spectral\_cluster}(I)$   
18 |  $I_k^{\text{cls}} = \text{get\_longest\_seg}([I_1^{\text{cls}}, I_2^{\text{cls}}, \dots])$   
19 |  $I_{\text{enroll}} = \text{random\_select\_sub\_seg}(I_k^{\text{cls}}, L_{\text{Enroll\_tmp}})$   
20 | // average the embeddings with indexes in  $I_{\text{enroll}}$   
 $\mathbf{e} = \text{average}(E_{I_{\text{enroll}}})$   
21 |  $\mathcal{E}.\text{add}(\mathbf{e})$   
22 |  $I_{\text{spk}} = (\text{AED-EEND}(\mathcal{E}))$   
Output:  $I_{\text{spk}}$ 
```

- **Init-Decode:** Use the initial L_{enroll} frames of the first time-continuous segments in I as the enrollment area.
- **Rand-Decode:** Randomly select a time-continuous segment I' from I , and then randomly select L_{enroll} length consecutive frames from I' as the enrollment area.
- **SC-Decode:** Cluster the embeddings in I using spectral clustering algorithm [4, 5], and then randomly select L_{enroll} length consecutive frames from the biggest cluster as the enrollment area.

Besides, for comparison, we also follow the strategy in the training process to get the enrollment area and denote this strategy as **GT-Decode**.

3. Experimental Setup

To ensure comparability with previous studies, we adopted most of the configurations used in EEND-EDA [10] except for the model architecture. In [10], the authors utilized an LSTM encoder-decoder to output the attractors. As outlined in section 2.1, we replaced the LSTM with a four-layer transformer model. The attention unit number is set to 256 in our experiment. Besides, our experiment also utilized the 345-dimensional acoustic features in [10] as input to the embedding encoder.

Regarding training data, we followed the pipeline in [10] to simulate 1-, 2-, 3-, and 4-speaker datasets. However, unlike [10], we only created the simulation evaluation set with a sim-

ilar overlap ratio to the training set. We used the CALLHOME [23] dataset to evaluate our system on real recordings and split the dataset into two parts - part1 for adaptation and part2 for evaluation, in accordance with the Kaldi recipe ¹.

Following [10], we evaluated our systems under two conditions: the fixed number of speakers condition and the flexible number of speakers condition. For the fixed number of speakers condition, we trained the systems on 2-speaker or 3-speaker datasets for 100 epochs. To evaluate the systems on real recordings, we further fine-tuned the models on the corresponding number of speakers adaptation set from CALLHOME for another 100 epochs. For the flexible number of speakers condition, we first fine-tuned the pre-trained model on the 2-speaker simulation set on the concatenation of 1-, 2-, 3-, and 4-speaker datasets for another 25 epochs. To evaluate the model on CALLHOME, we further fine-tuned the model on the entire CALLHOME adaptation set for another 100 epochs.

During our training process, we set the enrollment length L_{enroll} , as introduced in section 2, to 10-30 frames, corresponding to 1s-3s. During the evaluation process, we set the enrollment length L_{enroll} to 5 frames (0.5s) and the stop decoding length L_{stop} to 10 frames (1s). Additionally, we randomly set some enrollment speaker embeddings to zero vectors during the training process to ensure the model works correctly even when not all the enrollment embeddings are available.

We evaluated each system’s performance using the diarization error rate (DER) with a 0.25s collar tolerance to mitigate potential labeling errors.

4. Results and Analysis

4.1. Comparison Among Different Decoding Methods

Table 1: *DER (%) results for different decoding methods on 2-speaker evaluation set.*

Dataset	Decoding Methods			
	GT-Decode	Init-Decode	Rand-Decode	SC-Decode
Simulation				
Estimated #spk	-	3.40	3.50	3.25
Oracle #spk	3.22	3.14	3.13	3.14
CALLHOME				
Estimated #spk	-	10.7	8.58	8.34
Oracle #spk	8.03	10.9	8.32	7.75

This section presents the evaluation results of different decoding methods proposed in section 2.3. The results are summarized in Table 1. In the experiments, we either use the oracle speaker number or estimate the speaker number using the stop decoding criterion proposed in section 2.3. Surprisingly, the performance gap between estimated speaker number and oracle speaker number is not significant, which demonstrates the effectiveness of the stop decoding criterion in accurately predicting the speaker number. Furthermore, the SC-Decode method exhibits the best performance on most evaluation conditions, even surpassing the GT-Decode method. Therefore, we adopt the SC-Decode method in the subsequent experiments.

4.2. Results on Fixed Number of Speakers

This section evaluates our system when the speaker number is fixed and known in advance. The corresponding results are

¹https://github.com/kaldi-asr/kaldi/tree/master/egs/callhome_diarization/v2

Table 2: *DER (%) results on the fixed number of speakers condition.*

Method	Simulation DER(%)		CALLHOME DER(%)	
	2-spk	3-spk	2-spk	3-spk
x-vector clustering [10]	28.77	31.78	11.53	19.01
BLSTM-EEND [8]	12.28	-	26.03	-
SA-EEND [9, 10]	4.56	8.69	9.54	14.00
EEND-EDA [10]	2.69	8.38	8.07	13.92
TS-VAD [17]	-	-	9.51	-
AED-EEND (ours)	3.14	5.16	7.75	12.87

shown in Table 2. Apart from the 2-spk simulation dataset evaluation, our proposed method outperforms all the other methods. It is worth noting that even though our proposed AED-EEND method performs worse than the EEND-EDA on the 2-spk simulation dataset, our method outperforms the EEND-EDA on the 2-spk CALLHOME dataset, which demonstrates the better generalizability of our proposed method.

4.3. Results on Flexible Number of Speakers

Table 3: *DER (%) results on the simulation dataset with flexible number of speakers.*

Method	spk#			
	1	2	3	4
x-vector clustering [10]				
Estimated #spk	37.42	7.74	11.46	22.45
Oracle #spk	1.67	28.77	31.78	35.76
EEND-EDA [10]				
Estimated #spk	0.39	4.33	8.94	13.76
Oracle #spk	0.16	4.26	8.63	13.31
AED-EEND (ours)				
Estimated #spk	0.09	2.82	6.88	12.00
Oracle #spk	0.09	2.82	6.57	11.50

4.3.1. Results on simulation dataset

In this section, we evaluate our systems when the speaker number is flexible. We first do the experiment on the simulation dataset and show the result in Table 3. From the results, we find our method outperforms both the stage-wise x-vector clustering method and the EEND-EDA method on any number of speakers condition.

Table 4: *DER (%) results on the CALLHOME dataset with flexible number of speakers.*

Method	spk#					
	2	3	4	5	6	all
x-vector clustering [10]						
Estimated #spk	15.45	18.01	22.68	31.40	34.27	19.43
Oracle #spk	8.93	19.01	24.48	32.14	34.95	18.98
EEND-EDA [10]						
Estimated #spk	8.50	13.24	21.46	33.16	40.29	15.29
Oracle #spk	8.35	13.20	21.71	33.00	41.07	15.43
AED-EEND (ours)						
Estimated #spk	6.96	12.56	18.26	34.32	44.52	14.22
Oracle #spk	6.79	12.36	19.84	32.42	37.08	14.00

4.3.2. Results on CALLHOME dataset

Then, we evaluate our system on the CALLHOME dataset and list the results in Table 4. The overall performance of our proposed method is better than both x-vector clustering and EEND-EDA, and this improvement mainly comes from the small speaker number condition. Similar to the conclusion on

the simulation dataset, whether or not to use an oracle speaker number in the decoding process has little effect on the results of CALLHOME dataset. Besides, similar to the results of EEND-EDA, our system performs worse than the x-vector clustering in the 5-speaker and 6-speaker conditions. In the follow-up work [24] of EEND-EDA, the authors improve these results by simulating the data with more speakers in each recording. We plan to try this approach in our future work.

4.4. Analysis of Non-Speech, Single-speaker Speech and Overlap Speech Prediction

Table 5: *AED-EEND result (%) on the non-speech, single-speaker speech and overlap speech prediction. All the results are based on our 2-spk systems. For the evaluation metric with * superscript, we use the total duration of the utterance as the denominator to calculate the metric.*

Dataset	Non-Speech		Single-Speaker		Overlap	
	FA	MISS	FA	MISS	FA	MISS
Simulation	3.45	3.46	6.57	6.11	8.73	9.64
CALLHOME	29.50	30.64	9.81	7.57	26.28	41.23
	FA*	MISS*	FA*	MISS*	FA*	MISS*
Simulation	0.72	0.72	3.38	3.14	2.42	2.67
CALLHOME	2.89	3.00	7.70	5.94	3.09	4.86

As mentioned in section 2.3, we first predict the non-speech, single-speaker speech, and overlap speech area for an input utterance and then use the single-speaker speech area as the prior knowledge to do the iterative decoding. This section presents an evaluation of our proposed system on these three types of predictions, as shown in Table 5. Here, we will treat each type of prediction independently, without considering the potential conflicts between different types of predictions. The simulation dataset yields accurate area predictions for all three types of speech, with false alarm and miss rates below 10%. However, for the CALLHOME dataset, our system fails to accurately predict the non-speech and overlap speech areas. Further analysis of the dataset reveals that CALLHOME has a much smaller proportion of non-speech and overlap speech compared to the simulation dataset. To better understand the error types, we compute the percentage of each type of error relative to the recording duration and present the results in the lower section of Table 5. From the results, we find that our proposed system can also accurately predict all three types of speech. However, the small proportion of non-speech and overlap areas of CALLHOME lead to higher false alarm and miss rates.

5. Conclusion

In this paper, we propose the AED-EEND system, which effectively integrates enrollment information into the EEND-EDA system. To enhance the model’s training, we introduce a teacher forcing technique that leverages ground-truth labels to obtain enrollment embeddings during training. We present three heuristic decoding approaches to identify the enrollment information for each speaker during the evaluation process. Furthermore, we replace the LSTM architecture for attractor calculation in EEND-EDA with an attention-based model. Our approach achieves remarkable performance with very short enrollment data. AED-EEND introduces a novel paradigm for diarization tasks. However, due to space limitations, our evaluation on real datasets was limited. In future research, we will conduct a more thorough analysis of our proposed AED-EEND system and evaluate it in a broader range of real-world datasets.

6. References

- [1] S. H. Shum, N. Dehak, R. Dehak, and J. R. Glass, “Unsupervised methods for speaker diarization: An integrated and iterative approach,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 10, pp. 2015–2028, 2013.
- [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] 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.
- [4] 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.
- [5] Q. Lin, R. Yin, M. Li, H. Bredin, and C. Barras, “Lstm based similarity measurement with spectral clustering for speaker diarization,” *arXiv preprint arXiv:1907.10393*, 2019.
- [6] G. Sell and D. Garcia-Romero, “Diarization resegmentation in the factor analysis subspace,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4794–4798.
- [7] M. Diez, L. Burget, and P. Matejka, “Speaker diarization based on bayesian hmm with eigenvoice priors,” in *Odyssey*, 2018, pp. 147–154.
- [8] Y. Fujita, N. Kanda, S. Horiguchi, K. Nagamatsu, and S. Watanabe, “End-to-end neural speaker diarization with permutation-free objectives,” *arXiv preprint arXiv:1909.05952*, 2019.
- [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,” *arXiv preprint arXiv:2005.09921*, 2020.
- [11] 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.
- [12] Y. Takashima, Y. Fujita, S. Watanabe, S. Horiguchi, P. Garcia, and K. Nagamatsu, “End-to-end speaker diarization conditioned on speech activity and overlap detection,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 849–856.
- [13] J. R. Hershey, Z. Chen, J. Le Roux, and S. Watanabe, “Deep clustering: Discriminative embeddings for segmentation and separation,” in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 31–35.
- [14] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny *et al.*, “Target-speaker voice activity detection: a novel approach for multi-speaker diarization in a dinner party scenario,” *arXiv preprint arXiv:2005.07272*, 2020.
- [15] —, “The stc system for the chime-6 challenge,” in *CHiME 2020 Workshop on Speech Processing in Everyday Environments*, 2020.
- [16] W. Wang, X. Qin, M. Cheng, Y. Zhang, K. Wang, and M. Li, “The dku-smiip diarization system for the voxceleb speaker recognition challenge 2022,” in *Voxsrc Workshop*, 2022.
- [17] C.-Y. Cheng, H.-S. Lee, Y. Tsao, and H.-M. Wang, “Multi-target extractor and detector for unknown-number speaker diarization,” *arXiv preprint arXiv:2203.16007*, 2022.
- [18] D. Wang, X. Xiao, N. Kanda, T. Yoshioka, and J. Wu, “Target speaker voice activity detection with transformers and its integration with end-to-end neural diarization,” *arXiv preprint arXiv:2208.13085*, 2022.
- [19] M. Cheng, W. Wang, Y. Zhang, X. Qin, and M. Li, “Target-speaker voice activity detection via sequence-to-sequence prediction,” *arXiv preprint arXiv:2210.16127*, 2022.
- [20] Y. Jiang, R. Tao, Z. Pan, and H. Li, “Target active speaker detection with audio-visual cues,” *arXiv preprint arXiv:2305.12831*, 2023.
- [21] 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.
- [22] A. Woodward, C. Bonnín, I. Masuda, D. Varas, E. Bou-Balust, and J. C. Riveiro, “Confidence measures in encoder-decoder models for speech recognition,” in *INTERSPEECH*, 2020, pp. 611–615.
- [23] “2000 nist speaker recognition evaluation,” 2000. [Online]. Available: <https://catalog.ldc.upenn.edu/LDC2001S97>
- [24] S. Horiguchi, Y. Fujita, S. Watanabe, Y. Xue, and P. Garcia, “Encoder-decoder based attractors for end-to-end neural diarization,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1493–1507, 2022.