



Robust End-to-end Speaker Diarization with Generic Neural Clustering

Chenyu Yang¹, Yu Wang^{1,2*}

¹Cooperative Medianet Innovation Center, Shanghai Jiao Tong University ²Shanghai AI Laboratory
{yangcy, yuwangsjtu}@sjtu.edu.cn

Abstract

End-to-end speaker diarization approaches have shown exceptional performance over the traditional modular approaches. To further improve the performance of the end-to-end speaker diarization for real speech recordings, recently works have been proposed which integrate unsupervised clustering algorithms with the end-to-end neural diarization models. However, these methods have a number of drawbacks: 1) The unsupervised clustering algorithms cannot leverage the supervision from the available datasets; 2) The K-means-based unsupervised algorithms that are explored often suffer from the constraint violation problem; 3) There is unavoidable mismatch between the supervised training and the unsupervised inference. In this paper, a robust generic neural clustering approach is proposed that can be integrated with any chunk-level predictor to accomplish a fully supervised end-to-end speaker diarization model. Also, by leveraging the sequence modelling ability of a recurrent neural network, the proposed neural clustering approach can dynamically estimate the number of speakers during inference. Experimental show that when integrating an attractor-based chunk-level predictor, the proposed neural clustering approach can yield better Diarization Error Rate (DER) than the constrained K-means-based clustering approaches under the mismatched conditions.

Index Terms: Speaker diarization, end-to-end neural diarization, neural clustering

1. Introduction

Speaker diarization, which is often referred to as "who speaks when" problem, is an important speech processing task. The goal of a speaker diarization system is to estimate the temporal boundary of each talking speakers in real audio recordings [1, 2]. Traditional speaker diarization approaches first segment the audio streams into speech and non-speech regions, then the embeddings representing speaker characteristics are extracted from each speech frame. Popular used speaker embeddings include i-vectors [3], d-vectors [4, 5], and x-vectors [6]. After the speaker embeddings are extracted, the embeddings belonging to the same speaker are clustered using unsupervised clustering methods to yield the diarization results, where agglomerative hierarchical clustering (AHC), K-means clustering, and spectral clustering are commonly used methods. Various speaker embeddings and clustering techniques have been explored for speaker diarization tasks in [7, 8]. Although the traditional approaches can give good performance on various datasets, the main disadvantages of them are twofolds: first, in these approaches multiple components need to be opti-

mized using separate criteria and second, it is difficult for them to deal with overlapped speech as the unsupervised clustering normally assumes that each segment can only be assigned to one speaker. Although recently some approaches are proposed leveraging the supervised neural clustering models to improve the diarization performance, such as the unbounded interleaved-state recurrent neural network (UIS-RNN) [9] and Discriminative Neural Clustering (DNC) [10] approaches, they still suffer from the aforementioned two problems.

To solve those problems, end-to-end joint optimization-based models have been proposed which form the speaker diarization problem as a multi-label classification problem. The goal of these end-to-end framework is to achieve the diarization by a single neural network that is trained to directly optimize the diarization performance. The self-attentive EEND approach proposed in [11] was optimized to calculate diarization results for every speaker in a mixture from input audio features using permutation invariant training (PIT) [12]. EEND can only handle scenarios where numbers of speakers are fixed. It was later extended to deal with variable number of speakers by the end-to-end diarization model with encoder-decoder based attractor (EEND-EDA) [13] and speaker-wise conditional EEND (SC-EEND) [14]. Although previous experiments have shown that these end-to-end models can have a better or similar performance as the traditional methods, they fail to handle some complicated situations, for example, long-form speech. In order to address these issues, some approaches have been proposed to combine the advantages of both end-to-end models and unsupervised clustering algorithms as a two-stage framework [15, 16]. These two-stage approaches normally divide the complete audio sequence into chunks. Each chunk is processed by a neural-network-based model separately. An unsupervised clustering algorithm is then applied to determine the speaker correspondence among chunks. EEND-vector [15, 16] extends the EEND so that it can both diarize the subsequence and extract corresponding chunk-level speaker embeddings. These embeddings are clustered using the COP-K-means algorithm [17]. This approach enables the original model to process long overlapped speech. Thus, it combines the advantages of both clustering and the EEND-based method. The approach proposed in [18] incorporates local attractors with the EEND-EDA model which splits the sequence after feeding it into the encoder and replaces the linear decoder with a local encoder-decoder-based attractor calculation module, in order to deal with the mismatch of the number of speakers between training and evaluation.

However, these methods use a constrained unsupervised clustering algorithm as the post-processing, which have some limitations: 1) Some constrained clustering algorithms like COP-K-means are not robust enough. They can not handle conflicts between constraints. These conflicts may be caused by incorrect estimation of the number of speakers. 2) The unsupervised clustering algorithms are only used in the inference stage. Although extra loss functions have been applied to make

*Corresponding author.

This work was supported by National Natural Science Foundation of China under Grant No. 62106140, Shanghai Science and Technology Committee under Grant No. 21511101100. and Shanghai Key Lab of Digital Media Processing and Transmission (No. 18DZ2270700)

the training to be consistent with the inference, there is still unavoidable mismatch between the supervised training and unsupervised inference stages. This mismatch is especially obvious when estimating the number of clusters, which is required by most clustering algorithms.

To mitigate the problems of using unsupervised clustering algorithms and achieve reliable diarization for real speech, in this paper we propose a generic supervised neural clustering algorithm for speaker diarization based on recurrent neural network. The proposed approach can leverage the advantages of both the Transformer-based end-to-end framework and the neural clustering approaches that have been explored in [9, 10]. First, the proposed neural clustering network can be jointly optimized with the front end processing and the embedding extraction modules, thereby accomplishing a fully supervised the end-to-end optimization and achieving more efficient use of the annotated data. Second, during inference it is able to dynamically estimate the number of speakers in the same way as training. Finally, the proposed neural clustering approach is generic in the sense that it can be integrated with any chunk-level predictor such as the EEND-EDA-based or EEND-vector models.

2. Proposed Framework: EEND with Neural Clustering

2.1. Overview of two-stage frameworks

A neural end-to-end diarization model receives a sequence of acoustic features $F \in \mathbb{R}^{D \times T}$ as the input, and outputs the active probabilities of each speaker $Y \in (0, 1]^{S \times T}$, where D is the dimension of features, T is the length of input frames and S is the number of speakers. A two-stage diarization framework normally consists of two parts: the chunk-level predictor processes the chunk-level subsequences and calculates local speaker representations called attractors. These attractors are then fed into the clustering module which concatenates the diarization result of each chunk with the speakers being reordered correctly. Such two stages are jointly optimized with a loss function L :

$$L = L_{pre} + L_{post}.$$

Here L_{pre} is the loss function of the predictor, which mainly focuses on the internal diarization results of blocks. And L_{post} aims to train the clustering module among different blocks.

Such a proposed neural clustering approach can be integrated with any chunk-level predictor. Without loss of generality, in this paper we will focus on its integration with the EEND-EDA predictor with local attractors [13, 18] to deal with the mismatch of number of speakers. In the EEND-EDA framework, the input frames are firstly fed into a stack of Transformers, and the output of encoder E is then split into N chunks, $E = \{E_n\}_{n=1}^N$, $E_n \in \mathbb{R}^{D \times L}$, where L is the size of chunk. A chunk-level predictor is applied on each chunk separately. The details of this predictor will be given in Section 2.2. After the chunk-level prediction stage, the proposed neural clustering module then re-concatenates the outputs together, which will be described in Section 2.3.

2.2. Chunk-level diarization and attractor extraction

Similar to [18], the chunk-level predictor is a local encoder-decoder-based attractor (EDA) calculator aiming to diarize subsequences and to calculate speaker attractors. It consists of an EDA module and a Transformer decoder layer.

The details of an EDA module can be found in [13]. It can be defined as $\hat{a}^n, \hat{z}^n = \text{EDA}(E_n)$, where $\hat{a}^n \in \mathbb{R}^{D \times S_n}$ is the

speaker representation and $\hat{z}^n \in (0, 1]^{S_n}$ is the corresponding existence probabilities. The chunk-level active probability is computed by $y^n = \sigma(E_n^T \cdot \hat{a}^n) \in (0, 1]^{L \times S_n}$.

The chunk-level predictor is trained with two losses, namely the attractor existence loss and diarization loss. The attractor existence loss optimizes the estimation of each speaker's existence and the diarization loss optimizes the active probabilities with a permutation-invariant training method [19]:

$$L_{attr} = \frac{1}{N} \sum_{n=1}^N \text{BCE}(\hat{z}^n, z^n), \quad z^n = \underbrace{[1, \dots, 1, 0]}_{S_n}, \quad (1)$$

$$L_{diar} = \frac{1}{T} \sum_{n=1}^N \frac{1}{S_n} \min_{\phi \in \Phi^n} (\text{BCE}(y^n, l_\phi^n)). \quad (2)$$

Here Φ^n is the set of all possible permutations of S_n speakers in the current chunk and $l_\phi^n \in \{0, 1\}^{L \times S_n}$ is the corresponding diarization label. The total loss of chunk-level predictor are calculated as:

$$L_{pre} = L_{diar} + L_{attr}. \quad (3)$$

Since the local attractors are optimized to minimize the diarization error, they can not be used as the input for the neural clustering model directly. As in [18], the local attractors, \hat{a}^n , are converted using a Transformer decoder before clustering:

$$a^n = \text{TransformerDec}(\hat{a}^n, E).$$

2.3. RNN-based clustering methods

On the basis of the chunk-level predictor, we get pairs of attractors and diarization results for each chunk $\{(a_i^n, y_i^n)\}_{i=0}^{S_n-1}$. The goal of the neural clustering module is to reorder them by the optimal correspondence among chunks.

In this paper, we explore a recurrent neural network (RNN) for the neural clustering because of its sequence modelling ability. Gated recurrent units (GRUs) [20] are leveraged for the clustering. As shown in Figure 2, the hidden states of the GRU cell can play a similar role as clusters in the unsupervised clustering algorithms [9], which are used to model the global speakers. However, instead of being fixed during each iteration, these hidden states are updated adaptively according to the estimation of the attractors for each chunk. Thus, the estimation of the diarization outputs can be further split into two steps, namely the prediction and update steps. The details of these two steps are described as follows.

The prediction step aims to associate chunk-level attractors with the global hidden states, for which the RNN clustering (RC) model is optimized to estimate the probability $p_i^n = (p_{ij}^n)_j \in (0, 1]^C$ that the input attractor a_i^n belongs to each speaker, which can be calculated as:

$$p_i^n = \text{softmax} \left((a_i^n)^T h^{n-1} \right),$$

where $h^{n-1} = (h_j^{n-1})_j \in \mathbb{R}^{D \times C}$ are the previous hidden states. C is the number of clusters present before. During the training phase, all hidden states have been initialized in the beginning, so C is equalled to $S + 1$.

Although the attractors are still out of order at this stage, we can link each attractor to a hidden state vector with the optimal permutation ϕ^n using Equation (2). Thus, the targets for training such a sequential clustering model are formed as:

$$r_{ij}^n = \begin{cases} 1 & \phi_i^n = j \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$




Figure 1: Diagram of the RNN-based sequential clustering method. Every color represents an individual speaker. h^i is the set of all GRU hidden states after the i -th block. Each hidden state corresponds to a speaker appeared in previous chunks.

Because the optimal permutation has been estimated using Equation (2), permutation-invariant training is not required at this stage. Thus, a cross entropy (CE) loss is used to optimize the RNN clustering model:

$$L_{post} = \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{S_n} \sum_{i=0}^{S_n-1} \sum_{j=0}^{C-1} -r_{ij}^n \cdot \log(p_{ij}^n) \right).$$

After the prediction step, the goal of the update step is to adaptively adjust the hidden states at each step given the estimate of the attractors. Because r_{ij}^n provides a mapping from the i th attractor, a_i^n , to the j th hidden state, h_j^n , the hidden states can be updated using a teacher-forcing strategy [21]:

$$h_j^n = \begin{cases} h_j^{n-1} & \forall i \ r_{ij}^n = 0 \\ \text{GRU}(a_i^n, h_j^{n-1}) & \exists i \ r_{ij}^n = 1. \end{cases} \quad (5)$$

For the inference, the number of speakers is unknown. To dynamically estimate the attendance of new speakers, an initialized h_0 is concatenated with the previous states at every step. The optimal speaker permutation ϕ^n , is estimated by: $\phi^n = \arg \max_{\phi \in \Phi'} \prod_{i=0}^{S_n-1} p_{i, \phi_i}^n$, where Φ' is the set of all possible permutations. It should be noticed that a cannot-link exists between every pair of attractors in the same chunk, therefore each hidden state except for the new one can be linked to no more than one attractor. Then the hidden states are updated according to Equation (4) and (5).

After obtaining each ϕ^n , the final output of the diarization is reordered by replacing the chunk-level output y_i^n with $y_{\phi_i^n}^n$.

3. Experiments

3.1. Data

To evaluate the performance of the proposed RNN neural clustering integrated EEND-EDA approach, namely the EDA-RC approach, training and test datasets with different number of speakers were generated. The simulated datasets were generated using the data from the LibriSpeech[22] 360 hours of clean speech. The sampling frequency is 8kHz. We firstly selected n speakers from this dataset. Then for each speaker, $[30/n, 60/n]$ utterances were randomly selected. This simulating procedure was identical to the one that was introduced in [11]. It worth noting that the average length of silence intervals between utterances varied in order to maintain the overlap ratio in a reasonable range. The simulated training set includes 100,000 samples with 3 speakers. The details of the test sets for the simulation data, LS.3, LS.4 and LS.5, are given in Table 1.

The CALLHOME dataset, i.e., NIST SRE 2000 (LDC2001S97, Disk-8) [23], was used to further evaluate the performance on the real data. This dataset has been widely used for speaker diarization experiments, which contains 500 sessions of multilingual telephonic speech. Although each session has 2 to 6 speakers, there are 2 dominant speakers in each conversation. For the experiments, it was split into training and test sets using the tools provided in the Kaldi Toolkit diarization recipe [24]. Samples with ≤ 3 speakers in the training set were picked to train the models. The test set was further split with various numbers of speakers in order to evaluate the performance of approaches for both matched and unmatched conditions.

Table 1: Summary of datasets that are used in the experiments

Dataset	Source	#Spk	#Sample	Overlap(%)
LS_train	LibriSpeech	3	100,000	49.6
LS.3	LibriSpeech	3	500	42.7
LS.4	LibriSpeech	4	500	44.6
LS.5	LibriSpeech	5	500	45.2
CH_ft	CALLHOME.1	2-3	216	16.2
CH_test	CALLHOME.2	2-6	250	16.7

3.2. Experimental Setup

We use 23-dimensional log-Mel-filterbanks with a 25-ms frame length and 10-ms frame shift as the input feature. Feature vectors were concatenated with their neighboring $[-7, 7]$ frames and the concatenated features were subsampled by a factor of ten. Therefore, the total input dimension was 345 and the interval between frames was 100ms. This configuration was the same as that used in [11]. The training samples were fed into models every 500 frames while the test samples remain un-split.

The baseline approaches include both one-stage and two-stage end-to-end approaches: the EEND approach [11], the EEND-EDA approach [13] and the EEND-EDA with a COP-K-means unsupervised clustering on the local attractors (EDA+COP-K-means) [18]. For the EEND and EEND-EDA approaches, a 4-layer Transformer with 256 hidden units were used as the encoder. Each layer had 4 heads and a dropout rate of 0.1. The end-to-end models were trained for 50 epochs with a batch size of 32. An Adam optimizer [25] with the same learning rate scheduler in [26] was used during training.

For the EDA+COP-K-means and EDA-RC approaches, the

chunk size was set to 50. Parameters of the EDA chunk-level predictor were the same as those in EEND-EDA. The models were trained for another 20 epochs by an Adam optimizer with a learning rate of 10^{-5} . Diarization Error Rate (DER) was used as the evaluation metric. Errors less than 0.25s around boundaries were tolerated. Please note that for the DER computation, all of the errors were evaluated including the overlapping speech segments. For EDA-RC, a beam search strategy was applied during decoding and the beam size was set to 3. Models of the last 1/10 epochs were averaged to obtain a better performance.

3.3. Experimental Results

Table 2 shows the DERs for the Librispeech-based simulated datasets. Please note that the switching strategy described in [18] is not adopted here, because the main purpose is to compare the unsupervised and supervised clustering methods instead of an overall performance. As expected, it can be seen that EEND and EEND-EDA approaches can yield performance for matched conditions but get worse with the increase of speakers. On the other hand, the EDA+COP-K-means approach takes an obvious advantage over the EEND-EDA approach and it outperforms the EEND-EDA by about 10% DER under mismatched conditions. For our proposed EDA-RC approach, it further gives a DER reduction of 1.66% and 1.56% compared to the baseline approaches under mismatched conditions while the performance does not degrade much for the matched conditions.

Additionally, a refine method is applied to the EDA-RC, where an extra decoding process was conducted with the hidden states initialized as the first-pass decoded hidden states. Table 2 shows that this strategy can yield small performance gains when the number of speakers is large. Finally, the performance of the EDA-RC with the oracle permutations (oracle) is also given as a reference. This shows the performance upper bound of the clustering integrated EEND approaches as the correct clustering methods are used to reorder the speakers of each chunk.

Table 2: Diarization Error Rate (%) on simulated datasets. All models are trained with the 3-speaker training set.

Models	two-stage	match		mismatch	
		LS_3	LS_4	LS_5	
EEND [11]	✗	5.19	-	-	
EEND-EDA [13]	✗	4.93	33.79	41.82	
+ COP-K-means [18]	✓	6.93	22.48	32.36	
EDA-RC (ours)	✓	5.17	20.82	30.80	
EDA-RC+refine (ours)	✓	5.64	20.92	30.23	
EDA-RC (oracle)	✓	4.39	8.70	10.83	

In order to compare the performance on real audio recordings, models were also finetuned on the CALLHOME real datasets. Table 3 shows the DERs on CALLHOME. The proposed EDA-RC approach yields 25.46% and 35.79% DER for the 4 and 5 speakers' conditions, which corresponds to about 12.8% and 9.0% relative reduction compared to the EDA+COP-K-means approach. This validates the effectiveness and robustness of the proposed RNN-based clustering model.

To make a more comprehensive comparison, a switching strategy that can dynamically decide whether to decode at chunk level or global level according to the potential number of speakers is also applied [18] in this case. Thus, a global EDA loss is further added during the pretraining and finetuning of the

Table 3: Diarization Error Rate (%) on real dataset CALLHOME with different numbers of speakers.

Models	match		mismatch	
	2	3	4	5
EEND [11]	11.77	17.65	-	-
EEND-EDA [13]	10.91	17.05	25.36	38.58
+ COP-K-means [18]	12.34	20.23	29.21	39.34
+ switch[18]	11.51	18.51	25.19	39.01
EDA-RC (ours)	13.18	18.66	25.46	35.79
+ switch	11.78	18.20	26.57	37.04
EDA-RC (oracle)	11.91	14.50	17.69	18.93

EDA-RC model.

We notice that the gap between the clustering-integrated approaches and the conventional ones are rather small. This may be caused by the mismatch of data. As mentioned in Section 3.1, CALLHOME consists of real conversational speech where there are only two dominant speakers while all speakers have the same priority in the simulated data. Even though, with the switching method, clustering-integrated approaches can also have a better performance in some cases.

Considering an unsupervised clustering algorithm does not require a temporal order, it would be interesting to investigate the sensitivity of the RNN-based neural clustering to the order of the input chunks. The results are given in Table 4 with various shuffled ratios of samples. The DER decreased by about 1% when the order is completely shuffled, which indicates the EDA-RC is not sensitive to the order of inputs. This may be resulted from the position independence of the Transformer encoder. In addition, table 4 also shows that a global EDA loss has little influence on our EDA-RC.

Table 4: Performance of EDA-RC with the order of chunks being shuffled

w/ L_{global}	shuffled ratio (%)	Test data	
		LS_3	LS_4
-	0	5.17	20.82
-	50	5.54	21.21
-	100	5.79	21.88
+	0	5.78	21.27

4. Conclusions and Future Work

In this paper, we propose a generic neural clustering method for two-stage end-to-end diarization models. We evaluate our model by solving the speaker mismatch problem of overlapped speech recordings. And it can be extended to more applications in the future, e.g., online EEND systems [27, 28]. Compared to unsupervised clustering algorithms, our method is more robust and has a better performance. However, there is still some room for improvement. For example, such a method severely relies on the quality of local speaker representations, so a good chunk-level predictor is important. And it is sensitive to data adaptation, which should be solved in the future work. Apart from this, some advanced neural architectures like self-attention mechanism can also be used to improve the performance.

5. 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] 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.
- [3] 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.
- [4] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2014, pp. 4052–4056.
- [5] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end text-dependent speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 5115–5119.
- [6] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2018, pp. 5329–5333.
- [7] G. Sell, D. Snyder, D. McCree, Alanm Garcia-Romero, and 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.
- [8] M. Diez, F. Landini, L. Burget, and et al., "BUT system for DIHARD speech diarization challenge 2018." in *INTERSPEECH*, 2018, pp. 2798–2802.
- [9] A. Zhang, Q. Wang, Z. Zhu, J. Paisley, and C. Wang, "Fully supervised speaker diarization," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6301–6305.
- [10] Q. Li, F. L. Kreyssig, C. Zhang, and P. C. Woodland, "Discriminative neural clustering for speaker diarisation," in *IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 574–581.
- [11] Y. Fujita, N. Kanda, S. Horiguchi, Y. Xue, K. Nagamatsu, and S. Watanabe, "End-to-end neural speaker diarization with self-attention," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 296–303.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] K. Kinoshita, M. Delcroix, and N. Tawara, "Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 7198–7202.
- [16] —, "Advances in integration of end-to-end neural and clustering-based diarization for real conversational speech," in *INTERSPEECH*, 2021, pp. 3565–2569.
- [17] K. Wagstaff, C. Cardie, S. Rogers, S. Schrödl *et al.*, "Constrained K-means clustering with background knowledge," in *ICML*, vol. 1, 2001, pp. 577–584.
- [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 *ASRU*, 2021.
- [19] D. Yu, M. Kolbaek, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 241–245.
- [20] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [21] R. J. Williams and D. Zipser, "A learning algorithm for continually running fully recurrent neural networks," *Neural computation*, vol. 1, no. 2, pp. 270–280, 1989.
- [22] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 2015, pp. 5206–5210.
- [23] M. Przybocki and A. Martin, "2000 NIST speaker recognition evaluation (LDC2001S97)," in *Philadelphia, New Jersey: Linguistic Data Consortium*, 2001.
- [24] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The Kaldi speech recognition toolkit," in *IEEE workshop on automatic speech recognition and understanding*, 2011.
- [25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [26] 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.
- [27] E. Han, C. Lee, and A. Stolcke, "BW-EDA-EEND: Streaming end-to-end neural speaker diarization for a variable number of speakers," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 7193–7197.
- [28] Y. Xue, S. Horiguchi, Y. Fujita, S. Watanabe, P. Garcia, and K. Nagamatsu, "Online end-to-end neural diarization with speaker-tracing buffer," in *2021 IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 841–848.