

O-EENC-SD: Efficient Online End-to-End Neural Clustering for Speaker Diarization

Gruttadauria, Elio; Fontaine, Mathieu; Le Roux, Jonathan; Essid, Slim

TR2025-031 March 08, 2025

Abstract

We introduce O-EENC-SD: an end-to-end online speaker diarization system based on EEND-EDA, featuring a novel RNN-based stitching mechanism for online prediction. In particular, we develop a novel centroid refinement decoder whose usefulness is assessed through a rigorous ablation study. Our system provides key advantages over existing methods: a hyperparameter-free solution compared to unsupervised clustering approaches, and a more efficient alternative to current online end-to-end methods, which are computationally costly. We demonstrate that O-EENC-SD is competitive with the state of the art in the two-speaker conversational telephone speech domain, as tested on the CallHome dataset. Our results show that O-EENC-SD provides a great trade-off between DER and complexity, even when working on independent chunks with no overlap, making the system extremely efficient.

*IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
2025*

O-EENC-SD: Efficient Online End-to-End Neural Clustering for Speaker Diarization

Elio Gruttadauria*, Mathieu Fontaine*, Jonathan Le Roux†, Slim Essid*

*LTCI, Télécom Paris, Institut polytechnique de Paris, Palaiseau, France

†Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

Abstract—We introduce **O-EENC-SD**: an end-to-end online speaker diarization system based on EEND-EDA, featuring a novel RNN-based stitching mechanism for online prediction. In particular, we develop a novel centroid refinement decoder whose usefulness is assessed through a rigorous ablation study. Our system provides key advantages over existing methods: a hyperparameter-free solution compared to unsupervised clustering approaches, and a more efficient alternative to current online end-to-end methods, which are computationally costly. We demonstrate that O-EENC-SD is competitive with the state of the art in the two-speaker conversational telephone speech domain, as tested on the CallHome dataset. Our results show that O-EENC-SD provides a great trade-off between DER and complexity, even when working on independent chunks with no overlap, making the system extremely efficient.

Index Terms—Online speaker diarization, conversational telephone speech (CTS), CallHome, EEND-EDA

I. INTRODUCTION

Speaker diarization (SD) is an important speech processing task which aims at estimating the speech activities of each speaker in a real-world recording [1]. In online SD, data is processed on-the-fly in a streaming fashion and predictions must be causal. In contrast, in traditional or offline SD, the whole recording is available as the input of the system.

While SD is traditionally framed as a clustering problem [2], [3], recent developments pushed forward a novel end-to-end paradigm [4]–[6]. The benefits of the end-to-end neural diarization (EEND) family of models include the direct optimization of the diarization error and better handling of overlapped speech. Nevertheless, clustering remains a flexible tool both for offline and online inference. In offline SD, clustering has been used to reduce the computational burden of end-to-end models [7], [8] or to solve the empirical limitations on the maximum number of speakers the models can handle [9]. Clustering has also been used to adapt models to work in an online fashion [10], [11]: rather than the entire sequence being processed altogether, it is split into chunks whose predictions are then stitched back together. As end-to-end models are commonly trained under a permutation-invariant training (PIT) paradigm [12], the order of speaker activities in the output is unknown. In this context, the role of clustering is to solve the speaker permutation problem in-between blocks.

Another approach to solve the permutation problem is to include an overlap between consecutive blocks, as the shared portion can be used to reorder the speakers. This shared portion is commonly referred to as the *buffer*. The drawback of this approach is the repeated computation on the overlap region, which can be reduced with the use of more advanced buffering strategies [13], [14]. Buffering strategies encompass different sampling techniques to compose the shared portion between chunks, reaching performances competitive to fully causal models [15]. Even so, the buffer size considered in

This work was supported by the Audible project, funded by French BPI and partly supported by ANR Project SAROUMANE (ANR-22-CE23-0011). Also, it was performed using HPC resources from GENCI-IDRIS.

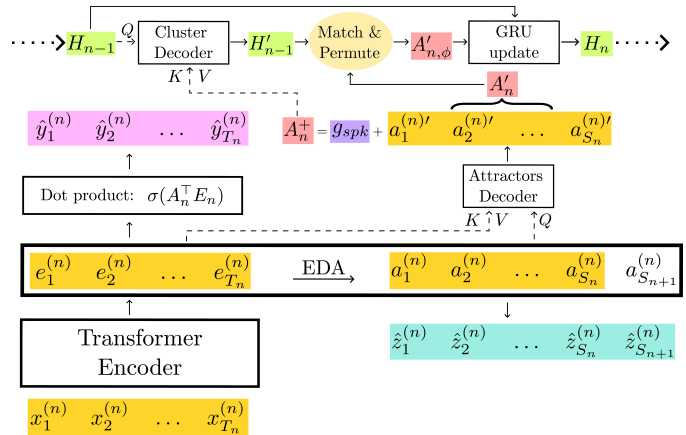


Fig. 1. Outline of O-EENC-SD, detailing the processing of a single chunk.

state-of-the-art systems tends to be quite large ($\geq 100s$), which entails a substantial computational cost (in terms of operations per second), thereby hindering their applicability to settings requiring efficient realizations, typically when the solution needs to run at the edge. This leaves the need for robust **lightweight** online variants, especially variants exploiting smaller buffers.

Contributions: In this work, we present a novel online speaker diarization system based on end-to-end neural diarization with encoder-decoder attractors (EEND-EDA) [5] and leveraging online neural clustering to solve the speaker permutation problem between predictions of consecutive chunks. The system, illustrated in Fig. 1, is competitive with the state of the art on online speaker diarization in the two-speaker conversational telephone speech (CTS) domain, as measured by the diarization error rate (DER). The strength of this work relies on the improved trade-off between performance and computational complexity, as shown by the marginal degradation of the DER when constraining the computational budget. A particular contribution of this work is the introduction of a new centroid refinement decoder, used to improve the accuracy of the neural clustering. To show the importance of each component of the model, we conduct an ablation study under different settings of latency and computational budget. Finally, we analyze the performance of O-EENC-SD when neither data nor computation is shared between subsequent chunks, opening the path to future works focusing on real-time SD on edge devices. The code to reproduce the results of this work is freely available.¹

II. RELATED WORK

In speaker diarization, few models are online by nature [16], [17]. More commonly, online SD is performed with offline models adapted

¹<https://github.com/egruttadauria98/O-EENC-SD>




Fig. 2. **Left:** High-level view of the stitching process on non-overlapping chunks. The EEND-EDA model estimates \hat{Y}_n (I) and the relative attractors, matched to the centroids (II). Based on the optimal permutation between centroids and attractors, the local predictions are permuted accordingly and stitched to the previous chunks to form the clustering output (III). **Right:** Zoom on the matching process for the second chunk (enclosed in dashed box, $n = 2$). Two attractors need to be matched: the second is matched to the only centroid (green), while the first is matched to h^0 (grey), creating a new centroid.

to handle streaming data [10], [11], [13]–[15]. To do so, the data is divided into chunks to be processed “locally”, before combining the output of local (chunk-level) predictions through a stitching process.

In recent years, the end-to-end paradigm has gained a lot of traction in the landscape of offline SD. Currently, several alternative methods are available to adapt end-to-end diarization (EEND) models to operate online. Coria et al. [10] proposed using a stitching method based on unsupervised online clustering of the speaker embeddings associated with each predicted speech activity. The advantage of clustering is to enable handling more speakers in the whole recording than what is allowed by the EEND model in a single chunk. As an unsupervised method, clustering can be extremely flexible, but also hard to optimize, and a gap between training and inference is usually observed. Furthermore, clustering usually requires extensive domain-specific hyperparameter tuning to achieve high performance.

To optimize an EEND model to work online in an end-to-end fashion, a speaker-tracing buffer (STB) [13], [14] can be used. In its simplest form, using a first-in-first-out (FIFO) strategy, it is comparable to having an overlap between consecutive chunks with a hop size smaller than the chunk size. More advanced strategies involve heuristics for selecting past frames to summarize as much information into a limited buffer. In [15], using speaker-balanced sampling probabilities with a 100-s long frame-wise STB was shown to even improve the DER on CallHome compared to a fully causal model, from 15% to 14.93%. Unfortunately, using a 100-s buffer with a hop size of 1 s is not feasible for many applications with limited computational resources. As the buffer needs to preserve the information of all the speakers seen so far in the form of acoustic features to be further processed, the buffer size is critical for performance (see Table 3 of [15]). In contrast, neural clustering [9] stores speaker information in a much more compact format: a centroid for each seen speaker, evolving from step to step.

III. PROPOSED SYSTEM

The proposed system processes the data in chunks, as shown in Figs. 1 and 2. Online diarization is obtained by sequentially stitching the chunk-predictions together, thereby solving the speaker permutation problem. New speakers might also be found as data is processed. Roughly, the system is composed of two main components: an EEND-EDA model [5] to obtain local predictions from each chunk, and a recurrent neural network (RNN) [18] that performs online neural clustering for the stitching. To refine the latent representation of the speakers, two transformer decoders [19] are used: an *attractor refinement decoder* and a *centroid refinement decoder*. The latter is one of the novel contributions of this work.

A. EEND-EDA

The EEND-EDA model processes acoustic features $X \in \mathbb{R}^{D' \times T_n}$ into frame embeddings $E_n \in \mathbb{R}^{D \times T_n}$, where n is the chunk

index and T_n is the number of frames in a chunk. The EDA component estimates attractors $A_n = [a_1^{(n)}, \dots, a_{S_n}^{(n)}] \in \mathbb{R}^{D \times S_n}$ in an autoregressive fashion, where S_n is the number of speakers in the chunk. To let the network determine that there are no more speakers, existence probabilities $\hat{Z}_n = [\hat{z}_1^{(n)}, \dots, \hat{z}_{S_n+1}^{(n)}] \in (0, 1)^{S_n+1}$ are computed. The speaker probabilities are estimated for each frame as the dot product between the frame embeddings and the attractors, followed by a sigmoid activation: $\hat{Y}_n = \sigma(A_n^T E_n) \in (0, 1)^{S_n \times T_n}$, where $\hat{Y}_n = [\hat{y}_1^{(n)}, \dots, \hat{y}_{T_n}^{(n)}]$. To train the model, two losses are used: an existence loss $L_{\text{exist}}^{\text{chunk}}$ on the attractors and a diarization loss $L_{\text{diar}}^{\text{chunk}}$. Both losses are implemented using the binary cross-entropy (BCE):

$$L_{\text{exist}}^{\text{chunk}} = \frac{1}{N} \sum_{n=1}^N \text{BCE}(\hat{Z}_n, Z_n), \quad Z_n = \underbrace{[1, \dots, 1, 0]}_{S_n}, \quad (1)$$

$$L_{\text{diar}}^{\text{chunk}} = \frac{1}{N} \sum_{n=1}^N \frac{1}{T_n S_n} \min_{\phi \in \Phi_n} \sum_{t=1}^{T_n} \text{BCE}(\hat{y}_t^{(n)}, P_{\phi} y_t^{(n)}), \quad (2)$$

$$L_{\text{EEND-EDA}}^{\text{chunk}} = L_{\text{diar}}^{\text{chunk}} + L_{\text{exist}}^{\text{chunk}}, \quad (3)$$

where N is the total number of chunks, Φ_n is the set of all permutations ϕ of S_n speakers, and P_{ϕ} denotes the permutation matrix associated with ϕ . During training, the EEND-EDA model is also run on the entire audio and the losses are applied globally, similarly to [15]:

$$L_{\text{exist}}^{\text{global}} = \text{BCE}(\hat{Z}, Z), \quad Z = \underbrace{[1, \dots, 1, 0]}_S, \quad (4)$$

$$L_{\text{diar}}^{\text{global}} = \frac{1}{TS} \min_{\phi \in \Phi} \sum_{\tau=1}^T \text{BCE}(\hat{y}_{\tau}, P_{\phi} y_{\tau}), \quad (5)$$

$$L_{\text{EEND-EDA}}^{\text{global}} = L_{\text{diar}}^{\text{global}} + L_{\text{exist}}^{\text{global}}, \quad (6)$$

where T and S are the number of frames and the number of speakers in the entire recording, respectively.

B. Online neural clustering

Online diarization is performed by finding the correspondence of attractors across consecutive chunks. Neural clustering [9] leverages RNNs to model the evolution of speaker centroids, each modeled by a different RNN. The hidden state of the RNN represents the centroid of the cluster. All RNNs are initialized with a common trainable embedding h^0 . Rather than using a fixed rule to update the centroids, RNNs can provide a more robust and flexible alternative. Specifically, we use gated recurrent units (GRUs) [20].

The online neural clustering is trained by classifying each attractor of the chunk to one of the known C centroids or to h^0 , meant to represent an average speaker. Using h^0 as an option also allows online clustering to be framed as a closed set and fully differentiable

problem. A cross-entropy loss is used to learn the correct assignment of chunk attractors to centroids:

$$L_{CE}^{\text{cluster}} = \frac{1}{N} \sum_{n=1}^N \left(\frac{1}{S_n} \sum_{i=1}^{S_n} \sum_{j=1}^C -r_{ij}^{(n)} \cdot \log(p_{ij}^{(n)}) \right), \quad (7)$$

with $(p_{ij}^{(n)})_{j=1,\dots,C} = \sigma_c(H_{n-1}^\top a_i^{(n)})$, where σ_c denotes the softmax over the centroids and $H_{n-1} = [h_1^{(n-1)}, \dots, h_C^{(n-1)}] \in \mathbb{R}^{D \times C}$ the hidden states of the previous step. For an attractor i whose speaker identity is represented by centroid j , $r_{ij}^{(n)}$ is 1, otherwise it is 0. Another option that we propose to supervise the learning of the neural clustering is to apply $L_{\text{diar}}^{\text{global}}$ to the output of the stitching process, depicted in Fig. 2. We will refer to this loss as $L_{\text{diar}}^{\text{cluster}}$.

At the end of each step, the GRUs of the active speakers are updated with the matched attractor using a teacher-forcing strategy [21]. The hidden states corresponding to the GRU of non-active speakers are not updated. Thus,

$$h_j^{(n)} = \begin{cases} h_j^{(n-1)} & \text{if } \forall i, r_{ij}^{(n)} = 0 \\ \text{GRU}(a_i^{(n)}, h_j^{(n-1)}) & \text{if } \exists i, r_{ij}^{(n)} = 1 \end{cases}, \quad (8)$$

where $1 \leq i \leq S_n$ and $1 \leq j \leq C$.

C. Refinement decoders

In the EEND-EDA model, the frame features are encoded by a stack of transformer layers [19]. Attractors are then estimated in an autoregressive fashion using an LSTM [22] encoder-decoder system. As attractors are optimized to minimize the diarization error within a chunk, [15] suggested using a transformer decoder to refine the attractors for inter-chunk stitching. Rather than using all past frame embeddings, we use only the ones of the current chunk as input to the decoder to make computation lightweight:

$$A'_n = \text{TransformerDecoder}(A_n, E_n, E_n) \in \mathbb{R}^{D \times S_n}. \quad (9)$$

A novel contribution of this work is the centroid refinement decoder. In online clustering, the update of the centroids usually happens after the new data points have been assigned to their corresponding clusters. The idea of the centroid refinement decoder is to perform a contextual update given the current attractors to be matched. To do so, the decoder uses a cross-attention mechanism between the centroids H_{n-1} and the concatenation of a trainable embedding $g_{\text{spk}} \in \mathbb{R}^D$, referred to as *ghost speaker*, with the attractors in the new chunk A_n :

$$H'_{n-1} = \text{TransformerDecoder}(H_{n-1}, A_n^+, A_n^+) \in \mathbb{R}^{D \times S_n}, \quad (10)$$

where $A_n^+ = \text{Concat}(g_{\text{spk}}, A'_n) \in \mathbb{R}^{D \times (S_n+1)}$.

The rationale behind the use of a ghost speaker is to add a degree of freedom to the refinement of the centroids of non-active speakers. In the latent space of the attention mechanism, queries are updated as a convex combination of keys and values. For active speakers, the decoder can push the centroid towards the correct attractor to facilitate the stitching step. For non-active speakers, the ghost speaker represents an extra option to attend to rather than the chunk's attractors.

IV. EXPERIMENTAL SETUP

A. Data

As we focus on conversational telephone speech (CTS), we evaluate O-EENC-SD on the CallHome corpus [23], using a 0.25-s tolerance collar to compare our results to previous works. For training, we distinguish between pretraining on simulated data and

fine-tuning on real data. For the simulated data, we rely on the licence-free CallHome and CallFriend corpora available on TalkBank [24] as input data, which are provided in stereo format and with no annotation. The data simulation pipeline is as follows: the stereo channels are analyzed individually with a pretrained diarization model to filter out the ones with more than one speaker. The remaining channels are combined randomly into simulated phone conversations. To improve the realism, speakers are sampled from the same language group. For the fine-tuning part, we use the development set of the CallHome corpus.

B. Architecture and Training Details

The configuration of the EEND-EDA model, and of the input features to the model, is the same as described in [5]. For the pretraining phase of O-EENC-SD, the EEND-EDA model is initialized as the average of the checkpoints of the last 10 epochs after being trained on simulated conversations (option 2), as described in [25]. The neural clustering component is initialized randomly. The learning rate is set to 10^{-4} for the pre-training phase and 10^{-5} for the fine-tuning phase. Before pre-training with the desired buffer size, all models are first trained on 10-s audio segments, split into 10 non-overlapping 1-s chunks. Our experiments show that this step leads to performance gains in the fine-tuned model.

Similarly to [15], we use both a global and a chunk-level EEND-EDA loss. In total, four loss terms are used: two for the EEND-EDA model ($L_{\text{EEND-EDA}}^{\text{global}}$ and $L_{\text{EEND-EDA}}^{\text{chunk}}$) and two for the neural clustering (L_{CE} and L_{diar}). The total loss is computed as:

$$L = L_{\text{EEND-EDA}}^{\text{global}} + 10L_{\text{EEND-EDA}}^{\text{chunk}} + L_{CE}^{\text{cluster}} + L_{\text{diar}}^{\text{cluster}}, \quad (11)$$

where 10 is the number of chunks into which the train data is usually split. During pre-training, we have found beneficial to only use the classification loss for the neural clustering. Models are trained to work at different latencies by modifying the attention mask of the EEND-EDA encoder. During inference, we only consider first-in-first-out (FIFO) [13], [14] as the buffering strategy.

V. RESULTS

In this work, we refer to the **entire** context used to process a chunk (including past observations and innovation) as *buffer*. As such, the buffer size is always greater than or equal to the latency, equivalent to the online processing unit in [15].

Table I collects the results of O-EENC-SD on CallHome to compare our proposal to state-of-the-art systems. All shown variants are trained with a 50-s buffer size, except for models in Table II, where the buffer size is equal to the latency. Unless otherwise specified, all models are trained with the same latency as used at test time.

A. Comparison to the state of the art

Our system demonstrates competitive performance with previous proposals, particularly when higher latency is allowed. Using a 100-s FIFO buffer, O-EENC-SD achieves a DER of 9.53% and 9.50% at 5-s and 10-s latency, respectively. The best model achieves a DER of 9.33% at 5-s latency when all past frames are included in the buffer. Our experiments reveal that O-EENC-SD achieves better performance when trained with higher latency values. Indeed, when performing inference with a 1-s latency, using a 100-s buffer size, the model trained with 5-s latency achieves 11.96% DER, better than the 12.47% of the 1-s latency model. Even with a significantly smaller 25-s buffer size, the 5-s latency model reaches a better 12.14% DER when tested at 1-s latency than the 1-s latency model using a 100-s buffer size.

TABLE I
 DERs (%) ON CALLHOME 2-SPKS TEST SET WITH 0.25S COLLAR
 TOLERANCE.

| Model | Buffer (s) | Latency (s) | DER (%) |
|-----------------------------------|------------|-------------|-------------|
| BW-EDA-EEND [16] | 10 | 10 | 11.82 |
| EEND-EDA+FW-STB [14] | 101 | 1 | 12.70 |
| EEND-EDA+FW-STB [†] [15] | 101 | 1 | 9.08 |
| EEND-GLA-Small+BW-STB [15] | 100 | 1 | 9.01 |
| EEND-GLA-Large+BW-STB [15] | 100 | 1 | 9.20 |
| <hr/> | | | |
| O-EENC-SD | ∞ | 10 | 9.50 |
| O-EENC-SD | 100 | 10 | 9.50 |
| O-EENC-SD | 10 | 10 | 10.75 |
| <hr/> | | | |
| O-EENC-SD | ∞ | 5 | 9.33 |
| O-EENC-SD | 100 | 5 | 9.53 |
| O-EENC-SD | 5 | 5 | 13.20 |
| <hr/> | | | |
| O-EENC-SD* | ∞ | 1 | 11.45 |
| O-EENC-SD* | 100 | 1 | 11.96 |
| O-EENC-SD* | 50 | 1 | 12.54 |
| O-EENC-SD* | 25 | 1 | 12.14 |
| O-EENC-SD* | 10 | 1 | 14.50 |
| O-EENC-SD* | 5 | 1 | 19.99 |
| <hr/> | | | |
| O-EENC-SD Base | 100 | 1 | 15.38 |
| + Attractors decoder | 100 | 1 | 15.07 |
| + Centroid decoder | 100 | 1 | 12.69 |
| O-EENC-SD | 100 | 1 | 12.47 |

[†] Reproduced results with improved training methodology.

* Model trained with a 5-s latency and a 50-s buffer size.

A possible explanation is that higher latency provides a stronger training effect on the EEND-EDA model to make better predictions at the chunk level. On the contrary, at low latency, the stitching is easier as the overlap portion between chunks is greater, leading to a smaller training effect. More specifically, the average EEND-EDA performance on a chunk is 11.84% DER when trained at 1-s latency and 8.89% at 5-s latency.

Training strategy improvements can lower the DER considerably in EEND-like models. In the top section of Table I, it is shown that the performance of EEND-EDA+FW-STB can be improved from 12.70% DER to 9.08% with the use of variable chunk-size training (VCT) [15], paired with improvements in the sampling probabilities of the STB. The profound effect of VCT stresses how much EEND-EDA performance on small chunks is essential for the overall online DER. From our experiments, high-latency training appears to be a good strategy to improve chunk performance, but it presents a domain shift with low-latency inference. We believe future works can bridge this gap and make online neural clustering competitive even at low latency requirements.

B. Architecture Ablation

The bottom of Table I also includes an ablation study on the model architecture at 1-s latency. We find that both decoders do improve performance over the base model. When the decoders are used jointly, O-EENC-SD exhibits the best results, increasing the clustering accuracy from 90% to 93% over the base model. Also, the performance of the EEND-EDA model in each chunk improves, decreasing the DER from 12.68% to 11.84%. Overall, the online clustering performance improves from 15.38% to 12.47% DER.

C. Efficiency

The proposed O-EENC-SD system particularly shines when tested at low buffer sizes and on non-overlapping chunks, making the system

TABLE II
 LOW-COMPUTATION RESULTS ON CALLHOME 2-SPKS TEST SET.
 L DENOTES THE NUMBER OF CONTEXT BLOCKS IN [16].

| Model | Complexity | Buffer (s) | DER (%) | Accuracy (%) |
|--------------------------------|------------------|------------|--------------|--------------|
| BW-EDA-EEND, $L = \infty$ [16] | $\mathcal{O}(T)$ | 10 | 11.82 | N/A |
| BW-EDA-EEND, $L = 1$ [16] | $\mathcal{O}(1)$ | 10 | 16.18 | N/A |
| <hr/> | | | | |
| O-EENC-SD Base | $\mathcal{O}(1)$ | 5 | 13.94 | 93.14 |
| O-EENC-SD | $\mathcal{O}(1)$ | 5 | 13.20 | 96.16 |
| <hr/> | | | | |
| O-EENC-SD Base | $\mathcal{O}(1)$ | 10 | 12.93 | 94.24 |
| O-EENC-SD | $\mathcal{O}(1)$ | 10 | 10.75 | 97.17 |

Latency is set to the buffer size for all models.

Accuracy refers to the ratio of correctly classified attractors to the total.

efficient from a computational perspective. By contrast, the paradigm of the previous methods [13]–[15] requires a long buffer to work properly. The only other EEND-based online model that does not use a buffer is BW-EDA-EEND [16], which implements a version of the EEND-EDA encoder motivated by Transformer-XL [26]: for every chunk, the hidden states of all layers are cached after being computed, to be used in the computation of the next L chunks. If $L = \infty$, all previously cached hidden states are used in the current chunk and the time complexity of the attention mechanism is linear with respect to the sequence length T . If $L = 1$, only the hidden states from the previous chunk are used and the time complexity does not depend on the sequence length. For O-EENC-SD, when the latency is set equal to the buffer size as in Table II, each chunk is considered independently. As the RNN update does not depend on T , the computation for a single chunk has constant time complexity. Table II shows that O-EENC-SD with a 10-s buffer size outperforms BW-EDA-EEND with $L = \infty$ by a significant margin, lowering the DER from 11.82% to 10.75%. Noticeably, all the model versions in Table II achieve better performance than BW-EDA-EEND with $L = 1$.

To conclude, we also want to remark on the impressive clustering accuracy of O-EENC-SD when trained with small buffer sizes, even surpassing 97% for the 10-s buffer size model. For models trained with a 50-s buffer size, the average accuracy is instead around 93%. As for the high-latency training, it seems that the online neural clustering prefers harder training strategies, which will be leveraged in future work to make the neural clustering paradigm stronger at low-latency requirements.

VI. CONCLUSIONS

In this work, we have proposed O-EENC-SD, a novel online speaker diarization method providing a great trade-off between DER and latency/complexity. Our results show that online neural clustering is competitive with previous systems, thanks in particular to the proposed centroid refinement decoder and the downstream diarization loss. Future work will focus on leveraging the strong potential of high-latency training for low-latency inference. To conclude, we believe online neural clustering has the potential to become state-of-the-art in online speaker diarization, combining efficiency and performance into a single paradigm.

ACKNOWLEDGMENTS

We would like to thank Chenyu Yang for their implementation of the offline neural clustering method, and Federico Landini for the checkpoints of the pretrained EEND-EDA model.

REFERENCES

- [1] T. J. Park, N. Kanda, D. Dimitriadis, K. J. Han, S. Watanabe, and S. Narayanan, "A Review of Speaker Diarization: Recent Advances with Deep Learning," 2021.
- [2] M. Sahidullah, J. Patino, S. Cornell, R. Yin, S. Sivasankaran, H. Bredin, P. Korshunov, A. Brutti, R. Serizel, E. Vincent *et al.*, "The Speed Submission to DIHARD II: Contributions & Lessons Learned," 2019.
- [3] F. Landini, O. Glembek, P. Matějka, J. Rohdin, L. Burget, M. Diez, and A. Silnova, "Analysis of the BUT Diarization System for VoxConverse Challenge," 2021.
- [4] Y. Fujita, N. Kanda, S. Horiguchi, Y. Xue, K. Nagamatsu, and S. Watanabe, "End-to-End Neural Speaker Diarization with Self-attention," 2019.
- [5] 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," 2020.
- [6] 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," 2021.
- [7] K. Kinoshita, M. Delcroix, and N. Tawara, "Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds," *Proc. ICASSP*, 2021.
- [8] —, "Advances in integration of end-to-end neural and clustering-based diarization for real conversational speech," *Proc. Interspeech*, 2021.
- [9] C. Yang and Y. Wang, "Robust End-to-end Speaker Diarization with Generic Neural Clustering," in *Proc. Interspeech*, 2022, pp. 1471–1475.
- [10] J. M. Coria, H. Bredin, S. Ghannay, and S. Rosset, "Overlap-Aware Low-Latency Online Speaker Diarization Based On End-To-End Local Segmentation," in *Proc. ASRU*, 2021.
- [11] E. Gruttadauria, M. Fontaine, and S. Essid, "Online speaker diarization of meetings guided by speech separation," 2024.
- [12] Y. Fujita, N. Kanda, S. Horiguchi, K. Nagamatsu, and S. Watanabe, "End-to-End Neural Speaker Diarization with Permutation-Free Objectives," 2019.
- [13] Y. Xue, S. Horiguchi, Y. Fujita, S. Watanabe, and K. Nagamatsu, "Online End-to-End Neural Diarization with Speaker-Tracing Buffer," 2021.
- [14] Y. Xue, S. Horiguchi, Y. Fujita, Y. Takashima, S. Watanabe, P. Garcia, and K. Nagamatsu, "Online Streaming End-to-End Neural Diarization Handling Overlapping Speech and Flexible Numbers of Speakers," 2021.
- [15] S. Horiguchi, S. Watanabe, P. Garcia, Y. Takashima, and Y. Kawaguchi, "Online Neural Diarization of Unlimited Numbers of Speakers," 2022.
- [16] E. Han, C. Lee, and A. Stolcke, "BW-EDA-EEND: streaming END-TO-END Neural Speaker Diarization for a Variable Number of Speakers," in *Proc. ICASSP*, 2021, pp. 7193–7197.
- [17] A. Zhang, Q. Wang, Z. Zhu, J. Paisley, and C. Wang, "Fully Supervised Speaker Diarization," 2019.
- [18] R. M. Schmidt, "Recurrent Neural Networks (RNNs): A gentle Introduction and Overview," 2019.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention Is All You Need," 2017.
- [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," 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] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, 1997.
- [23] Przybocki, Mark and Martin, Alvin, "2000 NIST Speaker Recognition Evaluation," 2001.
- [24] "TalkBank." [Online]. Available: <https://talkbank.org/>
- [25] F. Landini, A. Lozano-Diez, M. Diez, and L. Burget, "From Simulated Mixtures to Simulated Conversations as Training Data for End-to-End Neural Diarization," 2022.
- [26] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context," 2019.