



# Speaker Diarization with Overlapping Community Detection Using Graph Attention Networks and Label Propagation Algorithm

Zhaoyang Li<sup>1</sup>, Jie Wang<sup>1</sup>, XiaoXiao Li<sup>2</sup>, Wangjie Li<sup>1</sup>, Longjie Luo<sup>1</sup>, Lin Li<sup>\*1</sup>, Qingyang Hong<sup>\*3</sup>

<sup>1</sup>School of Electronic Science and Engineering, Xiamen University, China

<sup>2</sup>School of Electronic Information, Beijing Jiaotong University, China

<sup>3</sup>School of Informatics, Xiamen University, China

23120231150269@stu.xmu.edu.cn, {lilin, qyhong}@xmu.edu.cn

## Abstract

In speaker diarization, traditional clustering-based methods remain widely used in real-world applications. However, these methods struggle with the complex distribution of speaker embeddings and overlapping speech segments. To address these limitations, we propose an Overlapping Community Detection method based on Graph Attention networks and the Label Propagation Algorithm (OCDGALP). The proposed framework comprises two key components: (1) a graph attention network that refines speaker embeddings and node connections by aggregating information from neighboring nodes, and (2) a label propagation algorithm that assigns multiple community labels to each node, enabling simultaneous clustering and overlapping community detection. Experimental results show that the proposed method significantly reduces the Diarization Error Rate (DER), achieving a state-of-the-art 15.94% DER on the DIHARD-III dataset without oracle Voice Activity Detection (VAD), and an impressive 11.07% with oracle VAD.

**Index Terms:** speaker diarization, graph attention network, label propagation, overlapping community detection

## 1. Introduction

Speaker diarization is designed to segment the conversational audio signals into distinct segments with labeled speakers' identities, effectively addressing the problem of "who spoke when" [1]. In recent years, speaker diarization technology has gained widespread application across diverse fields, such as meeting transcription, conversational analysis [2]. However, real-world scenarios still face issues like speech overlap and unknown speaker counts.

Current representative speaker diarization systems can be broadly categorized into two types: end-to-end systems and modular systems. End-to-end systems can directly process raw audio and produce diarization results [3, 4]. These systems demonstrate notable strengths in handling overlapping speech, but still face challenges such as ambiguous speaker label assignments. Additionally, methods like Target Speaker Voice Activity Detection (TS-VAD) [5, 6] are highly reliant on initial clustering results. In contrast, the modular framework, also known as Clustering-based Speaker Diarization (CSD), typically comprises four components: Voice Activity Detection (VAD), a speaker representation extractor, a clustering module, and a post-processing module. The clustering module often employs traditional unsupervised clustering methods such as Agglomerative Hierarchical Clustering (AHC) [7], or Spectral Clustering (SC) [8]. However, the complex distribution of speaker embeddings and the sensitivity of various methods to hyperparameters

\*Corresponding author.

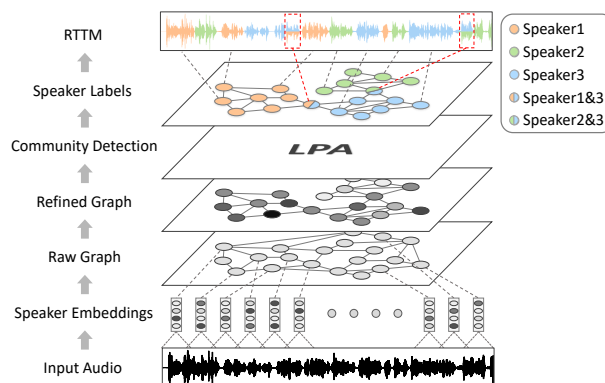


Figure 1: An illustration of the speaker diarization system pipeline with the OCDGALP clustering method.

can impact the classification performance of these systems.

For high-dimensional speaker embeddings such as x-vector [9], Graph Neural Network (GNN) can effectively leverage relational information between data points to further refine the speaker embeddings, thereby enhancing their representativeness. Wang et al. [10] initially used Graph Convolutional Network (GCN) to refine speaker embeddings, effectively separating different speakers within a conversation. Wang and Zheng [11, 12] formulated the clustering task as a community detection problem to achieve speaker classification. Despite significant progress in clustering, these methods still face challenges with overlapping audio, indicating that spatial relationships in speaker embeddings remain underutilized.

Inspired by the aforementioned works, this paper proposes a novel speaker diarization method based on Graph Attention network (GAT) and label propagation, with the system flow illustrated in Figure 1. Our main contributions are as follows:

- **Enhanced Speaker Embedding Representation:** The speaker embedding graph is fed into the GAT, where high-dimensional embeddings are projected into a lower-dimensional space. This transformation provides a clearer representation of speaker relationships, thus facilitating subsequent clustering.
- **Integrated Clustering and Overlap Detection:** By combining clustering with overlap detection, label propagation leverages node connections predicted by the graph neural network. This enables iterative partitioning of overlapping communities, effectively addressing overlapping speech.
- **State-of-the-Art Performance:** Experimental results show that the proposed method significantly reduces Diarization Error Rate (DER). Specifically, it achieves a remarkable 15.94% on the DIHARD-III dataset [13] without oracle VAD and an impressive 11.07% with oracle VAD.

## 2. Related works

### 2.1. Graph attention network

In our work, a modified GAT model is used to update the speaker embedding graph. Assuming  $L$ -layer graph attention network, where  $l \in \{1, 2, \dots, L\}$  represents the layer index, we take the embedding matrix  $\mathbf{H} \in \mathbb{R}^{K \times D}$ , which consists of pre-extracted speaker embeddings, along with the adjacency matrix  $\mathbf{A} \in \mathbb{R}^{K \times K}$  as the network inputs, where  $K$  is the number of nodes in the graph and  $D$  is the dimension of the embeddings.

Each layer of the graph attention mechanism performs information aggregation for nodes with the following form:

$$z_i^{(l+1)} = \sigma\left(\sum_{j \in N_i} \alpha_{ij}^{(l)} \mathbf{W}^{(l)} z_j^{(l)}\right). \quad (1)$$

Here,  $z_i^{(l+1)}$  represents the output feature of node  $v_i$  in the embedding matrix  $\mathbf{H}$  at layer  $(l+1)$ . The set  $N_i$ , which includes the neighboring nodes connected to  $v_i$  (including itself), can be obtained from the adjacency matrix  $\mathbf{A}$ . The matrix  $\mathbf{W} \in \mathbb{R}^{D_1 \times D_2}$  represents the trainable weight matrix at layer  $l$ , with input and output dimensions of  $D_1$  and  $D_2$ , respectively. The attention coefficient  $\alpha_{ij}$  determines the importance of the neighboring node  $v_j$  to the node  $v_i$ . In addition,  $\sigma$  is a nonlinear activation function. According to the method proposed by Velickovic et al. [14], the attention coefficient  $\alpha$  is calculated as follows:

$$\begin{aligned} \alpha_{ij} &= \text{softmax}_j(\hat{\alpha}_{ij}) \\ &= \frac{\exp(\delta(\mathbf{a}^T [\mathbf{W}z_i \parallel \mathbf{W}z_j]))}{\sum_{b \in N_i} \exp(\delta(\mathbf{a}^T [\mathbf{W}z_i \parallel \mathbf{W}z_b]))}. \end{aligned} \quad (2)$$

where  $\mathbf{a} \in \mathbb{R}^{2 \cdot D_2}$  is a learnable weight vector,  $T$  denotes transposition,  $\parallel$  is the concatenation operation, and  $\delta(\cdot)$  is the LeakyReLU activation function.

### 2.2. Label propagation algorithm

The Label Propagation Algorithm (LPA) [15] is a fast and efficient community detection method but struggles with detecting overlaps and maintaining stable results. To overcome these drawbacks, several improved LPA-based methods have been proposed, such as COPRA [16], SLPA [17], and DLPA [18]. In our work, we employ a variant of COPRA, known as LPANNI (Label Propagation Algorithm with Neighbor Node Influence) [19], whose algorithmic steps are as follows:

1. For a given network  $G = (V, \mathcal{E})$ , where  $V$  represents the set of nodes and  $\mathcal{E}$  represents the set of edges, and  $u \in V$ ,  $v \in Ng(u)$ , and  $Ng(u)$  represents the set of neighbors of  $u$ , we introduce additional definitions for nodes in the graph: Node Importance ( $NI(u)$ ), Node Similarity ( $Sim(u, v)$ ), and Neighbor Node Influence ( $NNI_v(u)$ ), which are defined in the same way as suggested by Lu et al. [19].
2. Nodes are sorted according to  $NI(u)$  for label updating. At this stage, nodes will receive multiple dominant labels from neighboring nodes, forming a label set:

$$L_{Ng} = \{l(c_1, b_1), l(c_2, b_2), \dots, l(c_v, b_v)\}, \quad (3)$$

where  $l(c_v, b_v)$  denotes the dominant label of neighbor node  $v$ , and  $b_v$  denotes the belonging coefficient of node  $v$  to community  $c_v$ . The dominant label is the one in the label set with the highest belonging coefficient.

3. Recalculate the new belonging coefficient  $b'(c, u)$  of node  $u$  to community  $c$  using  $NNI$  and  $L_{Ng}$ :

$$b'(c, u) = \frac{\sum_{l(c_v, b_v) \in L_{Ng}, v \in Ng(u), c_v = c} b(c_v, v) NNI_v(u)}{\sum_{l(c_v, b_v) \in L_{Ng}, v \in Ng(u)} b(c_v, v) NNI_v(u)} \quad (4)$$

Then, the updated label set  $L'$  of node  $u$  is generated:

$$L' = \{l(c_1, b'_1), l(c_2, b'_2), \dots, l(c_{|L'|}, b'_{|L'|})\}, \quad (5)$$

where  $|L'|$  denotes the number of the labels in  $L'$ . Next, labels satisfying  $b'(c, u) < 1/|L'|$  are adaptively removed, forming the refined label set  $L''$ .

4. Normalize the belonging coefficients of the labels in  $L''$  to obtain the final label set  $L_u$ .

The update process follows a historical label priority strategy to reduce propagation randomness. If multiple labels share the highest belonging coefficients, the dominant label is chosen from those present in the previous iteration; otherwise, it is selected randomly.

In summary, the LPANNI algorithm enhances the stability of label propagation through the proposed strategy and effectively assigns multiple labels to overlapping community nodes. These characteristics make LPANNI particularly well-suited for addressing overlapping speakers in speaker diarization tasks.

## 3. Proposed Method

The overall framework of OCDGALP is illustrated in Figure 2. The following sections provide a detailed introduction to each module of OCDGALP.

### 3.1. Graph Generation

In a graph  $G = (V, \mathcal{E}, A)$  composed of conversations, as illustrated in Figure 2(a), the nodes  $V = \{v_1, v_2, \dots, v_N\} \in \mathbb{R}^N$  are derived from speaker embeddings extracted from audio segments, while the edges  $\mathcal{E} \in \mathbb{R}^{N \times N}$  represent the connections between nodes. The affinity matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$  stores edge weights based on similarity scores computed using probabilistic linear discriminant analysis (PLDA) [20], normalized to the range  $[0, 1]$ . To reduce complexity, only edges with weights above a threshold  $\mu$ , treated as a hyperparameter, are retained.

### 3.2. Graph attention model design

As shown in Figure 2(b), the entire network consists of two components: a graph attention encoder and a scoring module. When the original graph passes through the GAT encoder, it learns information from surrounding nodes with different weights, generating a more refined latent speaker embedding graph. Each pair of nodes in this refined graph is then used as input to the scoring module.

The scoring module consists of fully connected layers with nonlinear activation functions and is jointly trained with the GAT layers. During this process, it re-predicts the relationships between updated graph nodes to generate the reconstructed speaker embedding graph  $G'$  and the affinity matrix  $\hat{\mathbf{A}}$ . To preserve important structural information from the original embeddings during the encoding process, we propose fusing the reconstructed affinity matrix  $\hat{\mathbf{A}}$  with the initial affinity matrix  $\mathbf{A}$ . The fused affinity matrix  $\bar{\mathbf{A}}$  is computed as follows:

$$\bar{\mathbf{A}} = (1 - \varepsilon)\hat{\mathbf{A}} + \varepsilon\mathbf{A} \quad (6)$$

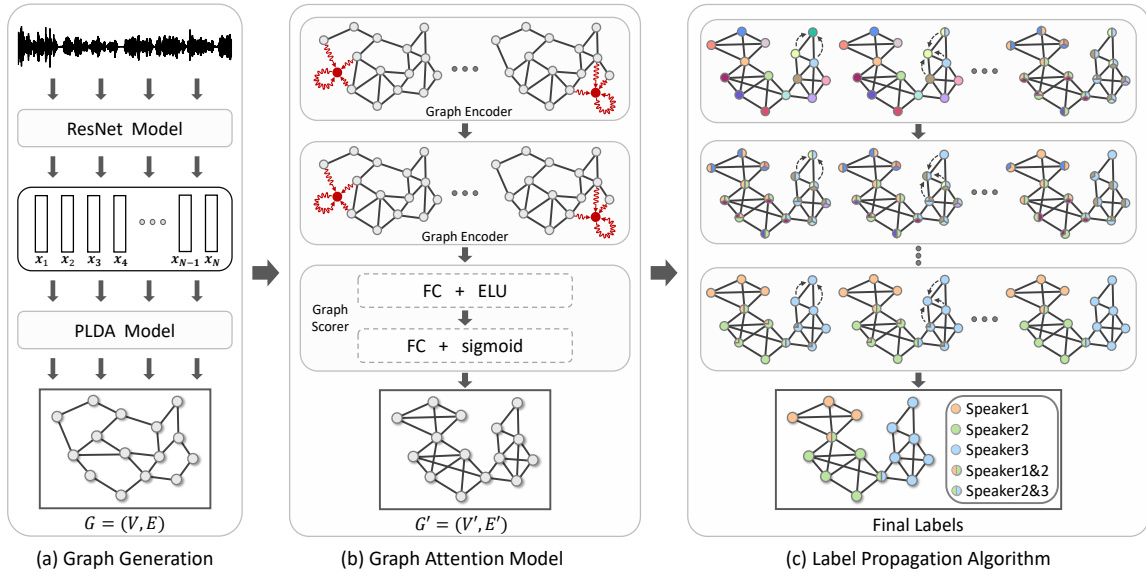


Figure 2: The architecture of the Overlapping Community Detection based on Graph Attention and Label Propagation Algorithm.

where  $\varepsilon \in [0, 1]$  is a balance coefficient that controls the influence of different affinity matrices. The loss function used for model training is defined as the difference between the fused affinity matrix and the ground truth adjacency matrix, calculated using Binary Cross-Entropy (BCE) loss.

### 3.3. Overlapping community detection

As shown in Figure 2(c), we integrate clustering with overlapping speaker detection and collectively refer to it as overlapping community detection.

The algorithm consists of two main stages. First, it calculates and sorts the  $NI$ , primarily to identify key speaker labels as propagation sources while also enhancing the algorithm’s stability. Then, we consider the order of  $NNI$  and  $Sim$  to be crucial for node updates. This is particularly important for overlapping nodes, which may be located at community boundaries and are more likely to possess multiple label assignments.

The convergence condition is met when the size of the label set and the dominant labels of all nodes remain stable or when the maximum number of iterations is reached. Upon completion of the iteration  $\tau$ , nodes with multiple speaker labels are considered overlapping nodes, and the overlapping community structure is detected. It is worth noting that when computing  $Sim$ , a path length threshold  $\beta$  is used to control computational complexity. Both  $\beta$  and the maximum iteration count  $\tau$  are regarded as hyperparameters.

## 4. Datasets and experimental setup

### 4.1. Evaluation and training data preparation

We evaluated our speaker diarization system on the DIHARD-III corpus, which includes a 34.15-hour development (DEV) set and a 33.01-hour evaluation (EVAL) set. The training sets for different modules in our system are described as follows:

- **VAD:** We fine-tuned a suitable VAD model on the DIHARD-III DEV set using the pre-trained models provided by pyannote.audio v3.1[21]. In addition, we also conducted experiments using the oracle VAD for comparison.
- **Embedding extractor and PLDA training:** We trained the embedding extractor and PLDA model using the VoxCeleb2

dataset, which contains over one million utterances.

- **GAT:** We generated a training dataset for the GAT model using VoxCeleb1, VoxCeleb2, and LibriSpeech. Following the methodology from [3], we simulated 6,000 mixed audio recordings featuring 2 to 9 speakers, with durations ranging from 20 to 600 seconds, totaling nearly 1,200 hours of data.

### 4.2. Experimental setup

In our diarization systems, we split the audio into 1.5s length segments with 0.75s window shift, and extract 256-dimensional embeddings of segments with the ResNet-34-SE model from ASV-Subtools [22]. The model consists of two GAT layers, followed by two fully connected layers with activation functions. The dimensionality settings are [d-128-64-64-1], with ELU and sigmoid as activation functions. The graph construction threshold  $\mu$  is set to 0.3, and the balance coefficient  $\varepsilon$  is set to 0.5.

During the subsequent overlapping community partitioning, we set the maximum number of iterations  $\tau$  to 80 and the path length  $\beta$  to 2. To mitigate the impact of algorithm instability on the results, each experimental outcome is obtained by averaging over 10 repeated trials.

## 5. Experimental results

### 5.1. Choice of hyperparameters

We investigate the effect of the balance coefficient  $\varepsilon$  and the path length threshold  $\beta$  on the OCDGALP system. As shown in Figure 3, the lowest DER is obtained when  $\varepsilon = 0.5$  and  $\beta = 3$ . This result demonstrates the effectiveness of using a fused affinity matrix and suggests that an optimal path length is crucial for efficient label propagation, as excessively long or short paths may hinder performance.

### 5.2. Speaker clustering results

In this section, we examine the impact of different modules on the performance of the speaker diarization system. We define three types of systems for comparison: S1 to S3 serve as baseline systems, each employing a different clustering method—AHC, SC, and LPA, respectively. In contrast, S4 to

Table 1: We evaluate the performance of different diarization systems in terms of DER (%) on the DIHARD III dataset under a 0ms collar condition, which includes the assessment of overlapping speech. The evaluation is conducted on the full dataset, comparing systems both with and without oracle VAD.

ID	System	Methods	wo/ Oracle VAD		w/ Oracle VAD	
			Dev	Eval	Dev	Eval
S1	Baselines	ResNet+PLDA+AHC	23.5	23.0	19.94	18.9
S2		ResNet+PLDA+SC	22.76	22.94	17.89	16.81
S3		Raw-LPA	23.68	23.73	20.45	19.32
S4	GCN-based	ResNet+GCN+AHC	22.36	22.45	18.04	17.55
S5		ResNet+GCN+SC	19.89	19.94	14.03	13.12
S6		ResNet+GCN+LPA-OCD	16.25	16.29	11.79	11.83
S7	GAT-based	ResNet+GAT+AHC	21.23	21.31	17.51	17.12
S8		ResNet+GAT+SC	19.56	19.63	13.88	13.01
S9		ResNet+GAT+LPA-OCD	<b>15.87</b>	<b>15.94</b>	<b>11.03</b>	<b>11.07</b>




Figure 3: The comparison of DER(%) with different  $\epsilon$  and  $\beta$  on the DIHARD III dataset.

S6 integrate GCN modules, while S7 to S9 incorporate GAT modules into the baseline systems. Notably, S9 represents our proposed OCDGALP system.

As illustrated in Table 1, the GAT-based system consistently outperforms both the baseline and GCN-based systems across all clustering methods. This highlights the superior effectiveness of the proposed GAT module. The attention mechanism in the GAT module, which assigns weighted importance to relationships between different nodes, is more effective than the uniform aggregation of node information in GCN. Notably, although Raw-LPA performs worse than AHC and SC when used independently, it significantly outperforms other clustering methods when integrated into the updated graph structure with GCN and GAT. This improvement is primarily attributed to its ability to detect overlapping structures, which enhances the overall performance of the proposed system.

### 5.3. Performance comparison

As shown in Table 2, we compare our proposed OCDGALP-based system with the latest systems in the literature on the DIHARD III dataset. To ensure a fair evaluation, we adhere to the same conditions specified in the original challenge: no collar tolerance is applied, and overlapping speech is considered.

The results clearly demonstrate that our system, when equipped with oracle VAD, not only outperforms existing single-clustering systems but also matched the performance

Table 2: DER(%) on the DIHARD III: A comparison with other systems in the literature

System	Method	DER(%)
<b>With Oracle VAD</b>		
SA-EEND[3]	EEND	16.19
EEND-EDA[23]	EEND	14.42
Hitachi /- fusion[24]	Fusion/cluster	11.58 / 15.74
USTC System[25]	Fusion	11.3
BUT system/- fusion[26]	Fusion/cluster	14.56 / 15.65
ANSD-MA-MSE[5]	TS-VAD	11.12
Seq2Seq-TSVAD[6]	TS-VAD	<b>10.77</b>
OCDGALP(ours)	Cluster	<u>11.07</u>
<b>Without Oracle VAD</b>		
Pyannote.audio v3.1[21]	Cluster	21.3
EEND-EDA[23]	EEND	20.69
DiaPer[27]	EEND	20.3
EEND-GLA[28]	EEND	19.49
VBx + Reseg[29]	Cluster	19.3
USTC System[25]	Fusion	16.78
ANSD-MA-MSE[5]	TS-VAD	16.76
EEND-M2F[30]	EEND	<u>16.07</u>
OCDGALP(ours)	Cluster	<b>15.94</b>

of the best system reported to date. Furthermore, in scenarios without oracle VAD, our system achieves the best publicly available performance with a DER of 15.94%, further demonstrating the effectiveness of our OCDGALP model.

## 6. Conclusion

In this paper, we propose the OCDGALP model for speaker diarization. Our primary goal is to leverage the topological structure of speaker embeddings and fully exploit graph-based information to achieve both clustering and overlap detection. To this end, we incorporate a GAT model to optimize the graph structure formed by speaker embeddings. Additionally, we employ LPA to determine the optimal partition of overlapping communities. Experimental results demonstrate that the OCDGALP-based speaker diarization system outperforms the baseline and several recent speaker diarization systems on the DIHARD-III corpus when oracle VAD is used. Moreover, it achieves state-of-the-art performance even without oracle VAD.

## 7. Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grants 62371407 and 62276220, and the Innovation of Policing Science and Technology, Fujian province (Grant number: 2024Y0068)

## 8. 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," *Computer Speech & Language*, vol. 72, p. 101317, 2022.
- [2] S. H. Yella, A. Stolcke, and M. Slaney, "Artificial neural network features for speaker diarization," in *2014 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2014, pp. 402–406.
- [3] 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.
- [4] 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.
- [5] M.-K. He, J. Du, Q.-F. Liu, and C.-H. Lee, "Ansd-ma-mse: Adaptive neural speaker diarization using memory-aware multi-speaker embedding," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 1561–1573, 2023.
- [6] M. Cheng, W. Wang, Y. Zhang, X. Qin, and M. Li, "Target-speaker voice activity detection via sequence-to-sequence prediction," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [7] G. Sell, A. McCree, and D. Garcia-Romero, "Priors for speaker counting and diarization with ahc," in *InterSpeech*, 2016, pp. 2194–2198.
- [8] W. Wang, Q. Lin, D. Cai, and M. Li, "Similarity measurement of segment-level speaker embeddings in speaker diarization," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 2645–2658, 2022.
- [9] 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.
- [10] J. Wang, X. Xiao, J. Wu, R. Ramamurthy, F. Rudzicz, and M. Brudno, "Speaker diarization with session-level speaker embedding refinement using graph neural networks," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7109–7113.
- [11] J. Wang, Z. Chen, H. Zhou, L. Li, and Q. Hong, "Community detection graph convolutional network for overlap-aware speaker diarization," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [12] S. Zheng and H. Suo, "Reformulating speaker diarization as community detection with emphasis on topological structure," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 8097–8101.
- [13] N. Ryant, P. Singh, V. Krishnamohan, R. Varma, K. Church, C. Cieri, J. Du, S. Ganapathy, and M. Liberman, "The third dihard diarization challenge," *arXiv preprint arXiv:2012.01477*, 2020.
- [14] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio *et al.*, "Graph attention networks," *stat*, vol. 1050, no. 20, pp. 10–48 550, 2017.
- [15] U. N. Raghavan, R. Albert, and S. Kumara, "Near linear time algorithm to detect community structures in large-scale networks," *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, vol. 76, no. 3, p. 036106, 2007.
- [16] S. Gregory, "Finding overlapping communities in networks by label propagation," *New journal of Physics*, vol. 12, no. 10, p. 103018, 2010.
- [17] J. Xie, B. K. Szymanski, and X. Liu, "Slpa: Uncovering overlapping communities in social networks via a speaker-listener interaction dynamic process," in *2011 IEEE 11th international conference on data mining workshops*. IEEE, 2011, pp. 344–349.
- [18] H.-L. Sun, J.-B. Huang, Y.-Q. Tian, Q.-B. Song, and H.-L. Liu, "Detecting overlapping communities in networks via dominant label propagation," *Chinese Physics B*, vol. 24, no. 1, p. 018703, 2015.
- [19] M. Lu, Z. Zhang, Z. Qu, and Y. Kang, "Lpanni: Overlapping community detection using label propagation in large-scale complex networks," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 9, pp. 1736–1749, 2018.
- [20] S. Ioffe, "Probabilistic linear discriminant analysis," in *Computer Vision—ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006, Proceedings, Part IV 9*. Springer, 2006, pp. 531–542.
- [21] A. Plaquet and H. Bredin, "Powerset multi-class cross entropy loss for neural speaker diarization," *arXiv preprint arXiv:2310.13025*, 2023.
- [22] F. Tong, M. Zhao, J. Zhou, H. Lu, Z. Li, L. Li, and Q. Hong, "Asv-subtools: Open source toolkit for automatic speaker verification," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6184–6188.
- [23] 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.
- [24] S. Horiguchi, N. Yalta, P. Garcia, Y. Takashima, Y. Xue, D. Raj, Z. Huang, Y. Fujita, S. Watanabe, and S. Khudanpur, "The hitachi-ju dihard iii system: Competitive end-to-end neural diarization and x-vector clustering systems combined by dover-lap," *arXiv preprint arXiv:2102.01363*, 2021.
- [25] Y. Wang, M. He, S. Niu, L. Sun, T. Gao, X. Fang, J. Pan, J. Du, and C.-H. Lee, "Ustc-nelslip system description for dihard-iii challenge," *arXiv preprint arXiv:2103.10661*, 2021.
- [26] F. Landini, A. Lozano-Diez, L. Burget, M. Diez, A. Silnova, K. Zmohková, O. Glembek, P. Matejka, T. Stafylakis, and N. Brümmer, "But system description for the third dihard speech diarization challenge," in *Proc. 3rd DIHARD Speech Diarization Challenge Workshop*, 2021.
- [27] F. Landini, T. Stafylakis, L. Burget *et al.*, "Diaper: End-to-end neural diarization with perceiver-based attractors," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- [28] S. Horiguchi, S. Watanabe, P. García, Y. Takashima, and Y. Kawaguchi, "Online neural diarization of unlimited numbers of speakers using global and local attractors," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 706–720, 2022.
- [29] H. Bredin and A. Laurent, "End-to-end speaker segmentation for overlap-aware resegmentation," *arXiv preprint arXiv:2104.04045*, 2021.
- [30] M. Härkönen, S. J. Broughton, and L. Samarakoon, "Eend-m2f: Masked-attention mask transformers for speaker diarization," *arXiv preprint arXiv:2401.12600*, 2024.