



## DIHARD II is Still Hard: Experimental Results and Discussions from the DKU-LENOVO Team

Qingjian Lin<sup>1,2</sup>, Weicheng Cai<sup>1,2</sup>, Lin Yang<sup>4</sup>, Junjie Wang<sup>4</sup>, Jun Zhang<sup>2</sup>, Ming Li<sup>1,3</sup>

<sup>1</sup>Data Science Research Center, Duke Kunshan University, Kunshan, China

<sup>2</sup>School of Electronics and Information Technology, Sun Yat-sen University, Guangzhou, China

<sup>3</sup>School of Computer Science, Wuhan University, Wuhan, China

<sup>4</sup>AI Lab of Lenovo Research, Beijing, China

ming.li369@dukekunshan.edu.cn

### Abstract

In this paper, we present the submitted system for the second DIHARD Speech Diarization Challenge from the DKU-LENOVO team. Our diarization system includes multiple modules, namely voice activity detection (VAD), segmentation, speaker embedding extraction, similarity scoring, clustering, resegmentation and overlap detection. For each module, we explore different techniques to enhance performance. Our final submission employs the ResNet-LSTM based VAD, the Deep ResNet based speaker embedding, the LSTM based similarity scoring and spectral clustering. Variational Bayes (VB) diarization is applied in the resegmentation stage and overlap detection also brings slight improvement. Our proposed system achieves 18.84% DER in Track1 and 27.90% DER in Track2. Although our systems have reduced the DERs by 27.5% and 31.7% relatively against the official baselines, we believe that the diarization task is still challenging.

**Index Terms:** DIHARD, VAD, speaker embedding, similarity scoring, clustering, resegmentation, overlap detection

### 1. Introduction

Speaker diarization is the task of determining “who spoke when” in an audio file that usually contains an unknown number of speakers with variable speech duration [1, 2]. It has a wide range of applications such as telephone calls, meeting recordings and broadcast interviews. Diarization can also serve as the frontend of automatic speech recognition (ASR) to improve the transcription performance in multi-speaker conversations.

In general, a diarization system partitions multi-speaker audios into short segments and clusters them according to speaker identities. It consists of following modules (shown in Figure 1):

- **Voice activity detection (VAD):** VAD detects speech in the audio signals and removes the non-speech regions to reduce computation. Typical VAD systems include generative models like Hidden Markov Model (HMM) [3], and discriminative models like linear discriminant analysis (LDA) [4], support vector machine (SVM) [5] and deep neural network (DNN) methods [6, 7, 8].
- **Segmentation:** The segmentation step splits speech into multiple speaker-homogeneous segments. In general, a speaker changepoint detector (SCD) [9, 10] is employed to search speaker changepoints and split speech by the changepoints, but uniform segmentation with overlap [11] also works fine.

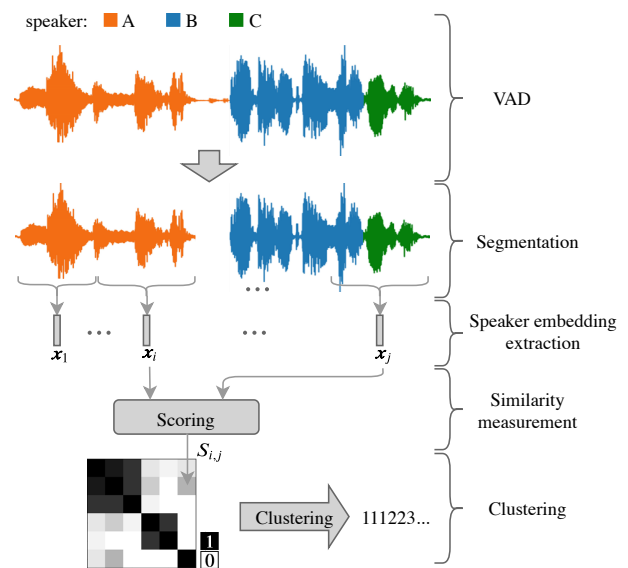


Figure 1: A standard speaker diarization framework with multiple modules.

- **Speaker embedding extraction:** After segmentation, short segments are mapped into the speaker subspace and generate fixed-dimensional speaker embeddings such as i-vector [12], x-vector [13] and Deep ResNet based speaker embeddings [14, 15].
- **Similarity measurement:** Similarity scores between any two speaker embeddings in the same audio are computed and later used in the clustering step. Popular techniques include cosine similarity, probabilistic linear discriminant analysis (PLDA) [16, 17] and DNN based measurement [18].
- **Clustering:** Clustering algorithms like K-means [19], agglomerative hierarchical clustering (AHC) [11, 20] and spectral clustering [18, 19] assign segments with high similarity scores to the same cluster. Note that the similarity measurement and clustering steps can be merged by a single online clustering backend, according to [21, 22].

The aforementioned modules construct the general framework of speaker diarization. Additional modules like resegmentation [23, 24] and overlap detection [25] are not essential but

Table 1: Metadata analysis on the DIHARD II dev set.

domains	n_audios	n_speakers	average duration	speech percentage(%)	overlapped error(%)
restaurant	12	5~8	10min 6s	88.01	25.20
meeting	14	3~10	10min 27s	93.89	22.42
webvideo	32	1~9	3min 31s	75.12	21.70
child	23	2~5	9min 51s	59.78	11.73
socio_field	12	2~6	10min 2s	72.62	7.53
socio_lab	16	2	5min 59s	74.38	4.78
clinical	24	2~3	3min 26s	60.94	3.75
maptask	23	2	6min 33s	68.15	2.92
court	12	5~10	10min 24s	84.09	1.90
broadcast_interview	12	3~5	10min 17s	78.70	1.18
audiobooks	12	1	10min 3s	79.37	0.00
ALL	192	1~10	7min 25s	76.07	10.76

help further improve the system performance.

Researches have achieved state-of-the-art performance in some eval datasets like CallHome [19, 21]. However, the 2017 JSALT Summer Workshop at CMU found it hard to migrate the success to more challenging corpora including web videos, speech in the wild, child language recordings, etc [26]. To raise researchers’ attention, the DIHARD competition is therefore proposed as a new annual challenge focusing on the “hard” scenarios. The second DIHARD speech diarization challenge [27, 28, 29] contains four tracks. Track1 and Track2 share the same single-channel audios while Track3 and Track4 use the same multiple-channel audios. An oracle voice activity detector is provided in Track1 and Track3. As for the rest tracks, participants are required to distinguish between speech and nonspeech by themselves. Since we only participate in Track1 and Track2, experiments and discussions are limited in these two tracks.

The rest of this paper is organized as follows. Section 2 carries out metadata analysis on the DIHARD II development (dev) set. Section 3 describes models and algorithms we used in the competition. Experimental results and discussions are presented in Section 4, while conclusions are drawn in Section 5.

## 2. Metadata analysis

In this section, we carry out metadata analysis on the DIHARD II dev set to illustrate how “hard” the competition is. Several indicators are important: duration of audios, the number of speakers, speech percentage and overlapped error. It is discussed in [18] that performance of diarization systems degrades as audio duration increases. Besides, with more speakers involved in conversations, it becomes harder for systems to make correct predictions about each speaker. Speech percentage reflects proportion of speech in audios and helps us choose suitable datasets for VAD training. Overlapped error determines the minimum diarization error rate (DER) a diarization system is possible to achieve without handling overlapped speech<sup>1</sup>. Assume that  $D$  and  $N$  denote duration and the number of speakers respectively, and  $R_i$  denotes speech regions of specific speaker  $i$ . Then speech percentage  $R\%$  can be expressed as follows:

$$R\% = \frac{\text{dur}(R_1 \cup R_2 \cup \dots \cup R_N)}{D}. \quad (1)$$

<sup>1</sup>The DIHARD competition takes the strictest definition of DER: no collocation is tolerated around speech boundaries and overlapped speech accounts for DER.

The  $\text{dur}(\cdot)$  function gets duration of input regions. Besides, the overlapped error  $O\%$  is calculated by:

$$O\% = \frac{\sum_{n=1}^N \text{dur}(R_i) - \text{dur}(R_1 \cup R_2 \cup \dots \cup R_N)}{\sum_{n=1}^N \text{dur}(R_i)}. \quad (2)$$

Results are shown in Table 1. In comparison with existing datasets, the DIHARD II dev set is drawn from 11 different domains. In some challenging domains like *restaurant*, conversations arise under the background of strong noise, making it difficult for systems to detect speakers. The number of involved speakers ranges from 1 to 10, and most of the domains contain conversations with more than two speakers except *audiobooks*, *maptask* and *socio.lab*. The average audio duration is 7min 25s, and three quarters is speech. The overall overlapped error is 10.76%, but in specific domains such as *meeting*, *restaurant* and *webvideo*, the error is much higher. It indicates that audios drawn from these domains are much harder to cope with, even in the DIHARD competition. In summary, the competition is “hard” mainly because:

- 1) audios are drawn from a diverse set of challenging domains.
- 2) the number of speakers varies in a large range.
- 3) high overlapped error accounts for DER.

## 3. System description

### 3.1. Data

Multiple datasets are employed in our experiments. Voxceleb1&2 [30] contain 16k sampled short utterances with single speakers, suitable for speaker embedding training. Multi-speaker audios are drawn from databases in the meeting and telephone domains. The 16k sampled meeting data consists of AMI [31], ICSI [32], ISL (LDC2004S05), NIST (LDC2004S09) and SPINE1&2 (LDC2000S87, LDC2000S96, LDC2001S04, LDC2001S06, LDC2001S08). The 8k sampled telephone data covers six monolingual CallHome sets, including Arabic (LDC97S45), English (LDC97S42), German (LDC97S43), Japanese (LDC96S37), Mandarin (LDC96S34) and Spanish (LDC96S35). MUSAN [33] and RIRS [34] corpora are employed for data augmentation.

For validation on the DIHARD II dev set, we regard it as a held-out dataset. It means neither threshold tuning nor model adaptation is performed. As for evaluation on the DIHARD II evaluation (eval) set, we take the dev set to finetune our models and improve the performance.




Figure 2: The structure of ResNet-LSTM VAD.

## 3.2. VAD

### 3.2.1. WebRTC

WebRTC [35] is the official VAD baseline for Track2. Raw audios are split into frames with 20ms duration. For each input frame, WebRTC generates output 1 or 0, where 1 denotes speech and 0 denotes nonspeech. An optional setting of WebRTC is the aggressive mode, an integer between 0 and 3. 0 is the least aggressive about filtering out nonspeech while 3 is the most aggressive.

### 3.2.2. ResNet-LSTM VAD

We propose a DNN based approach for the VAD task. The network structure, shown in Figure 2, consists of a ResNet module, multiple Bi-LSTM layers and linear layers. Our motivation is that the ResNet module generates representative feature mappings for speech and nonspeech, and the Bi-LSTMs capture sequential information. The input is a long sequence of frame-wise features. Every 8 frames in the sequence are packed and fed into the ResNet, generating multi-channel feature maps  $F \in \mathbb{R}^{C \times H \times W}$ . We apply the global average pooling (GAP) on each channel and get a  $C$ -dimensional vector. Next are Bi-LSTM layers to capture forward and backward sequence information. Finally, outputs from the Bi-LSTMs pass through linear layers connected with the sigmoid function, and generate the speech likelihood.

To be specific, we extract 32-dimensional log-mel-filterbank (fbank) features from audios by 25ms length and 10ms step. An input sequence contains  $200 \times 8 = 1600$  frames. ResNet18 is employed and its channel width is set to  $\{16, 32, 64, 128\}$ . Two Bi-LSTM layers are stacked, each with 64 units per direction and a dropout rate of 0.5. Followings are two linear layers with 64 units and 1 unit, respectively. Details about model parameters and output size are shown in Table 2. Training sets include 16k meeting data and 8k telephone data with augmentation, as described in Section 3.1. They are all down-sampled to 8k sample rate. The stochastic gradient descent (SGD) optimizer and the binary cross entropy (BCE) loss are employed. The learning rate is initialized as 0.1 and decreases by a factor of 1/10 every 20 epochs. The training process terminates after 60 epochs. For evaluation, we take the dev set to finetune the model for 10 more epochs.

## 3.3. Segmentation

We use the uniform segmentation with overlap rather than SCD to reduce computation. In condition that the sliding window of uniform segmentation is short enough, we can assume large

Table 2: Parameters and output size of the ResNet-LSTM VAD.

Layer	Structure	Output size
Inputs	-	$8 \times 32$
ResNet	conv $3 \times 3$ , 16	$16 \times 8 \times 32$
	conv $3 \times 3$ , 16 conv $3 \times 3$ , 16	$\times 2$ $16 \times 8 \times 32$
	conv $3 \times 3$ , 32 conv $3 \times 3$ , 32	$\times 2, /2$ $32 \times 4 \times 16$
	conv $3 \times 3$ , 64 conv $3 \times 3$ , 64	$\times 2, /2$ $64 \times 2 \times 8$
	conv $3 \times 3$ , 128 conv $3 \times 3$ , 128	$\times 2, /2$ $128 \times 1 \times 4$
Pooling	GAP	128
Bi-LSTM	64 units per direction, 2 layers, drop=0.5	128
Linear1	$128 \times 64$	64
Linear2	$64 \times 1$	1

portions of the split segments are speaker-homogeneous. It is also mentioned in [1] that uniform segmentation does not significantly degrade the overall diarization performance in comparison with SCD based segmentation. In our experiments, the sliding window is 1.5s long with 0.75s step. The ground-truth speaker of a segment is the one who talks most in the central 0.75s region.

## 3.4. Speaker embedding training

### 3.4.1. i-vector

I-vector is an unsupervised speaker embedding with generative models [12]. Essential components for an i-vector system include the universal background model (UBM) and the total variability space  $T$ . For a pre-trained system, we extract features from input audios and project on UBM to calculate the statistics supervectors. Then supervectors are mapped into the low-rank subspace  $T$  as i-vectors.

We follow the dihard2018/v1 recipe in kaldi [36] to train the i-vector extractor. 24-dimensional MFCCs are extracted from voxceleb1&2 with 25ms length and 10ms step. Cepstral mean normalization (CMN) is applied. The UBM includes 2048 components, and the dimension of i-vector is 400.

### 3.4.2. *x*-vector

*x*-vector is the DNN based speaker embedding [13]. In the training process, frame-wise features are fed into the time-delay neural network (TDNN) for supervised learning. Outputs from TDNN are pooled over the temporal dimension and transformed into a segment-level embedding, followed by multiple linear layers. For testing cases, embeddings from the second linear layer are taken as the *x*-vector.

Similarly, we follow the dihard2018/v2 kaldi recipe [20] to train our *x*-vector extractor. Voxceleb1&2 are augmented by MUSAN and RIRS corpora, and MFCCs are extracted with CMN. The dimension of *x*-vector is 512.

### 3.4.3. Deep ResNet based speaker embedding

The Deep ResNet structure [37] is an improved version of [14, 15]. It consists of three main components: a ResNet module, two parallelized pooling layers and two linear layers. First, the ResNet module transforms  $L$ -frame fbanks to feature maps  $\mathbf{F} \in \mathbb{R}^{C \times H \times W}$ . For each channel, the GAP layer and the global standard deviation pooling (GSP) layer are applied respectively, and outputs are concatenated together as a  $2C$ -dimensional vector  $\mathbf{v}$ :

$$v_{1c} = \text{GAP}(\mathbf{F}_c) = \frac{1}{H \times W} \sum_{h=1}^H \sum_{w=1}^W F_{c,h,w}, \quad (3)$$

$$v_{2c} = \text{GSP}(\mathbf{F}_c) = \frac{1}{H \times W} \sqrt{\sum_{h=1}^H \sum_{w=1}^W (F_{c,h,w} - v_{1c})^2}, \quad (4)$$

$$\mathbf{v} = [v_{11}, \dots, v_{1C}, v_{21}, \dots, v_{2C}]. \quad (5)$$

$\mathbf{v}$  is later fed into two linear layers and the softmax function sequentially, generating speaker likelihoods. Size of the last layer equals the number of speakers in training data.

In our experiments, the training sets include voxceleb1&2 and their augmentation, where there are 7323 speakers in total. We extract 64-dimensional fbank features from single-speaker utterances, with the number of frames  $L$  ranging from 200 to 400 dynamically. ResNet34 is employed and its channel width is set to  $\{32, 64, 128, 256\}$ . The two linear layers after pooling are 128-dimensional and 7323-dimensional, respectively. We list details of model parameters and output size in Table 3. The SGD optimizer and cross entropy loss are employed. The initial learning rate is 0.1, and reduces by a factor of 1/10 at the 25th and 40th epoch. Training process stops after 50 epochs. For testing cases, outputs from the first linear layer are taken as speaker embeddings.

## 3.5. Similarity measurement

Given speaker embeddings  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$  from the same audio, we compute similarity scores  $S_{i,j}$  between any two embeddings  $\mathbf{x}_i$  and  $\mathbf{x}_j$ , and construct similarity matrix  $\mathbf{S} \in \mathbb{R}^{n \times n}$  for the clustering step.

### 3.5.1. PLDA

The PLDA score between  $\mathbf{x}_i$  and  $\mathbf{x}_j$  can be expressed as:

$$S_{i,j} = \log\left(\frac{p(\mathbf{x}_i, \mathbf{x}_j | \mathcal{H}_0)}{p(\mathbf{x}_i | \mathcal{H}_1) \cdot p(\mathbf{x}_j | \mathcal{H}_1)}\right). \quad (6)$$

$\mathcal{H}_0$  hypothesizes  $\mathbf{x}_i$  and  $\mathbf{x}_j$  from the the same speaker while  $\mathcal{H}_1$  hypothesizes  $\mathbf{x}_i$  and  $\mathbf{x}_j$  from different speakers. In our ex-

Table 3: Parameters of the Deep ResNet based speaker verification system.

Layer	Structure	Output size
Inputs	-	$L \times 64$
ResNet	conv $3 \times 3, 32$	$32 \times L \times 64$
	conv $3 \times 3, 32$ conv $3 \times 3, 32$	$\times 3$ $32 \times L \times 64$
	conv $3 \times 3, 64$ conv $3 \times 3, 64$	$\times 4, /2$ $64 \times \frac{L}{2} \times 32$
	conv $3 \times 3, 128$ conv $3 \times 3, 128$	$\times 6, /2$ $128 \times \frac{L}{4} \times 16$
	conv $3 \times 3, 256$ conv $3 \times 3, 256$	$\times 3, /2$ $256 \times \frac{L}{8} \times 8$
Pooling	GAP + GSP	512
Linear1	$512 \times 128, \text{drop}=0.5$	128
Linear2	$128 \times 7323$	7323

periments, PLDA is trained by speaker embeddings from voxceleb1&2 without augmentation, and whitened by the dev set.

### 3.5.2. LSTM

Noticing that PLDA handles speaker embeddings in a pairwise and independent manner which ignores sequential information, we proposed LSTM-based scoring model in [18] to capture the forward and backward messages. In comparison with PLDA, scores are calculated between vector and sequence rather than vector and vector. Given speaker embeddings  $\mathbf{x}_i$  and  $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$ , we repeatedly concatenate  $\mathbf{x}_i$  with each embedding in the sequence and feed the concatenated sequence into the LSTM network. Outputs are scores of input concatenated vectors:

$$[S_{i,1}, S_{i,2}, \dots, S_{i,n}] = f_{\text{LSTM}}\left(\left[\begin{array}{c} \mathbf{x}_i \\ \mathbf{x}_1 \end{array}\right], \left[\begin{array}{c} \mathbf{x}_i \\ \mathbf{x}_2 \end{array}\right], \dots, \left[\begin{array}{c} \mathbf{x}_i \\ \mathbf{x}_n \end{array}\right]\right). \quad (7)$$

For large score matrix  $\mathbf{S}$  where  $n \geq 400$ , we partition it into small blocks and then perform LSTM scoring respectively, as is described in [18].

The architecture of LSTM network includes two Bi-LSTM layers followed by two linear layers. Both Bi-LSTM layers have 512 units (256 units per direction), and a dropout rate of 0.5 is added. The first linear layer is 64-dimensional and the second layer is 1-dimensional, connected with the sigmoid function to generate a similarity score between 0 and 1. Training features are speaker embeddings extracted from 16k meeting data described in Section 3.1. The SGD optimizer and the BCE loss are employed. The learning rate is initialized as 0.01 and decays by a factor of 1/10 every 40 epochs. The training process terminates after 100 epochs and the model is validated on the dev set. For evaluation, we further take the dev set to finetune our model for 30 more epochs.

## 3.6. Clustering

### 3.6.1. Agglomerative hierarchical clustering

Agglomerative hierarchical clustering (AHC) is performed as the binary-tree building process [38]. Segments are initialized as singleton clusters, and clusters with the highest pairwise similarity score are merged. The process iterates until scores between any two clusters is below a given threshold. For evalu-

ation, we tune the threshold on the dev set to achieve the best performance.

### 3.6.2. Spectral clustering

Spectral clustering (SC) is a graph-based clustering algorithm [39]. Given the similarity matrix  $\mathbf{S}$ , it considers  $S_{i,j}$  as the weight of the edge between nodes  $i$  and  $j$  in an undirected graph. By removing weak edges with small weights, spectral clustering divides the original graph into subgraphs. As described in [39], spectral clustering consists of the following steps:

- Construct  $\mathbf{S}$  and set diagonal elements to 0.
- Compute Laplacian matrix  $\mathbf{L}$  and perform normalization:

$$\mathbf{L} = \mathbf{D} - \mathbf{S}, \quad (8)$$

$$\mathbf{L}_{\text{norm}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{L} \mathbf{D}^{-\frac{1}{2}}, \quad (9)$$

where  $\mathbf{D}$  is a diagonal matrix and  $D_i = \sum_{j=1}^n S_{i,j}$ .

- Compute eigenvalues and eigenvectors of  $\mathbf{L}_{\text{norm}}$ .
- Compute the number of clusters  $k$ . One property of  $\mathbf{L}_{\text{norm}}$  indicates that the number of clusters in the graph equals algebraic multiplicity of the 0 eigenvalue. In our implementation, we count the number of eigenvalues below the a given threshold as  $k$ . The threshold is tuned by the dev set.
- Take the  $k$  smallest eigenvalues  $\lambda_1, \lambda_2, \dots, \lambda_k$  and corresponding eigenvectors  $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_k$  of  $\mathbf{L}_{\text{norm}}$  to construct matrix  $\mathbf{P} \in \mathbb{R}^{n \times k}$  using  $\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_k$  as columns.
- Cluster row vectors  $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n$  of  $\mathbf{P}$  to  $k$  classes by the K-means algorithm.  $\mathbf{y}_i \in \text{class } j$  indicates that segment  $i$  belongs to speaker  $j$ .

### 3.7. Resegmentation

Since results from clustering modules are segment-wise while the metric of DER computes error in a frame-wise manner, it unavoidably imposes loss of precision. To deal with the problem, resegmentation modules estimate distributions of clustering results and refine them by frames.

#### 3.7.1. GMM resegmentation

Given diarization outputs, the Gaussian Mixture Model (GMM) resegmentation constructs speaker-specific GMMs for each speaker by their speech. Then for each frame in the audio, we compute its posterior probabilities of belonging to different GMMs and reassign it to the GMM with the highest probability. The process iterates until converge.

We extract 24-dimensional MFCCs from raw audios to construct speaker-specific GMMs, each with 8 components. The resegmentation process iterates for a maximum number of 5 turns.

#### 3.7.2. VB diarization

The VB diarization method [23] constructs a HMM with eigen-voice priors. In comparison with GMM resegmentation, it takes speaker transition probabilities into consideration and avoids frequent speaker turns. Besides, a UBM is pretrained and speaker-specific GMMs are adapted on the UBM. Speaker-specific GMMs share the same component weights  $w_c^{ubm}$  and covariance matrices  $\Sigma_c^{ubm}$ , and thus only mean vectors require adaptation. Let  $\mathbf{u}_s$  be the mean supervector concatenated by

GMM component means for speaker  $s$ , and  $\mathbf{u}_{ubm}$  be the mean supervector of UBM. The algorithm assumes

$$\boldsymbol{\mu}_s = \boldsymbol{\mu}_{ubm} + \mathbf{T}\mathbf{z}, \quad (10)$$

where  $\mathbf{z}$  is the low-dimensional latent vector and subject to normal distribution  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ .  $\mathbf{T}$  is the total variability subspace pretrained from i-vector systems. From e.q. 10 we can infer the distribution of  $\boldsymbol{\mu}_s$ :

$$\boldsymbol{\mu}_s \sim \mathcal{N}(\boldsymbol{\mu}_{ubm}, \mathbf{T}\mathbf{T}^\top). \quad (11)$$

It further imposes priors for speaker-specific GMMs.

In our experiments, we employ the tool released by [23]. UBM with 1024 components are trained on 24-dimensional MFCCs and  $\mathbf{z}$  is 400-dimensional. Other parameters for the VB algorithm are: maxIters=1, downsample=3, loopProb=0.99, statScale=0.3.

### 3.8. Speech overlap detection

In Section 2, we point out high overlapped errors for DER metric through metadata analysis. To cope with the problem, we attempt to carry out overlap detection experiments. The model structure, data and training configurations are all the same as those in ResNet-LSTM VAD. The only difference is that we remove the nonspeech regions and mark overlapped speech and non-overlapped speech as 1 and 0, respectively. For testing cases, when a segment is discriminated as overlapped speech in audios, we extend its boundary by  $\pm 20$  frames and take all speakers appearing in the extended segment as labels of the original segment.

## 4. Results and discussions

### 4.1. VAD

To have a direct view at VAD performance, we carry out independent evaluation on WebRTC and our ResNet-LSTM based VAD. Ground-truth VAD labels of the eval set are available in Track1 but should be held-out in Track2, so the experiments are performed after the competition. Since the VAD model works as a submodule for the whole diarization system, we focus on the accuracy rate as the evaluation metric. Results are shown in Table 4.

Without adaptation, the ResNet-LSTM VAD model only gains slightly higher accuracy than the WebRTC baseline. When we take the dev set to finetune the model, performance on the eval set significantly improves to 91.4%. It mainly results from the fact that our training data is drawn from meeting and telephone domains while the DIHARD sets covers 11 domains. Domain mismatch leads to poor performance and adaptation brings significant improvement.

Table 4: Accuracy rate of different VAD modules.

Methods	Dev(%)	Eval(%)	Eval(adapt)(%)
WebRTC, mode=0	80.9	79.6	-
WebRTC, mode=1	81.5	79.7	-
WebRTC, mode=2	81.9	78.4	-
WebRTC, mode=3	78.3	75.4	-
ResNet-LSTM	82.2	80.7	<b>91.4</b>

Table 5: Evaluation on different combinations of speaker embeddings, scoring, and clustering methods.

System ID	Speaker embedding	Scoring & Clustering	Dev(%)	Eval(adapt)(%)
1	i-vector	PLDA + AHC	28.96	25.84
2		LSTM + SC	25.98	24.72
3	x-vector	PLDA + AHC	27.55	25.26
4		LSTM + SC	23.22	22.03
5	Deep ResNet	PLDA + AHC	23.51	21.7
6		LSTM + SC	21.95	<b>20.87</b>
7	fusion (2 + 4 + 6)		21.36	<b>20.24</b>

#### 4.2. Speaker embedding, similarity scoring and clustering

In Table 5, we compare different combinations of speaker embeddings, similarity scoring and clustering methods in Track1. It is observed that the Deep ResNet based speaker embedding outperforms i-vector and x-vector in all combinations. Besides, the backend of LSTM based scoring followed by spectral clustering behaves better in comparison with PLDA and AHC. Our best single system is System 6, which achieves a DER of 20.87%. When we fuse System 2, 4 and 6 by averaging their score matrices, the DER further reduces to 20.24%.

#### 4.3. Resegmentation

Resegmentation is carried out on our best single system (System 6) and the fusion system (System 7). Results are shown in Figure 6. In our expectation, the VB algorithm should outperform the GMM method, and resegmentation modules should bring similar improvement for both systems. To our surprise, for the fusion system, diarization predictions after resegmentation does not become more accurate. The most significant improvement is gained in System 6 with VB diarization, reducing the DER by 1.65% absolutely. The phenomenon is quite confusing, and we will keep investigating the reasons in our future works.

Table 6: Resegmentation on System 6 and the fusion system.

System ID	Reseg	Eval(%)
6	-	20.87
	GMM	20.52
	VB	<b>19.22</b>
7	-	20.24
	GMM	20.18
	VB	20.12

#### 4.4. Overlap detection

The last module in our diarization system is overlap detection. Since the overlapped error is as high as 10.76% in the dev set, it reasonable for us to assume there is around 10% of the same type error in the eval set. Experiments are carried out on System 6 with VB diarization. Results are shown in Table 7. Our attempt to detect overlapped speech only slightly reduces DERs by 0.38% in Track1 and 0.69% in Track2. It is still very challenging because we recalled less than 1/10 of the overlapped speech.

Our final submitted pipeline includes the ResNet-LSTM based VAD, the Deep ResNet based speaker embedding, the LSTM based scoring, spectral clustering, VB diarization and overlap detection. We achieve DERs of 18.84% for Track1 and

Table 7: Overlap detection on System 6 with VB diarization.

System ID	Reseg	Overlap	Track1(%)	Track2(%)
6	VB	no	19.22	28.59
		yes	<b>18.84</b>	<b>27.90</b>

27.90% for Track2, both ranking the 2nd on the leaderboard.

#### 4.5. Discussions

To understand how our system performs in each specific domain, we group DERs of the dev set on System 6 by domains. Results shown in Figure 3 are highly diverse between different domains. Our system performs worst in these four scenes: *restaurant*, *child*, *webvideo* and *meeting*, three of which are discussed in Section 2 due to high overlapped errors. The *child* domain, despite the low overlapped error, raises a high DER of 37.38%. This is probably because the audios are drawn from Seedlings corpus where children are 6-18 months old, which is a mismatch compared with the adult speakers in our training databases. As a result, our system still performs poorly in these challenging domains.




Figure 3: The performance of the proposed system in each specific domain on the DIHARD dev set.

## 5. Conclusion

In this paper, we provide detailed introduction and compare combinations of different modules in our diarization system, and report their performance on the DIHARD II. Our final submissions achieve DERs of 18.84% in Track1 and 27.9% in Track2. We further discuss DERs for specific domains, and point out that diarization is still challenging for unknown domains and domains with high overlapped errors.

## 6. Acknowledgements

This research is funded in part by the National Natural Science Foundation of China (61773413), Key Research and Development Program of Jiangsu Province (BE2019054), Six talent peaks project in Jiangsu Province (JY-074), Guangzhou Municipal People's Livelihood Science and Technology Plan (201903010040), Science and Technology Program of Kunshan City.

## 7. References

- [1] Sue E Tranter and Douglas A Reynolds, "An overview of automatic speaker diarization systems," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 5, pp. 1557–1565, 2006.
- [2] Xavier Anguera, Simon Bozonnet, Nicholas Evans, Corinne Fredouille, Gerald Friedland, and Oriol Vinyals, "Speaker diarization: A review of recent research," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no. 2, pp. 356–370, 2012.
- [3] Chuck Wooters, James Fung, Barbara Peskin, and Xavier Anguera, "Towards robust speaker segmentation: The icsi-sri fall 2004 diarization system," in *RT-04F Workshop*, 2004.
- [4] Elias Rentzperis, Andreas Stergiou, Christos Boukis, Aristodemos Pnevmatikakis, and Lazaros C Polymenakos, "The 2006 athens information technology speech activity detection and speaker diarization systems," in *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 2006, pp. 385–395.
- [5] Andrey Temko, Dusan Macho, and Climent Nadeu, "Enhanced svm training for robust speech activity detection," in *2007 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2007, pp. 1025–1028.
- [6] Rubén Zazo Candil, Tara N Sainath, Gabor Simko, and Carolina Parada, "Feature learning with raw-waveform cldnns for voice activity detection," in *Interspeech*, 2016, pp. 3668–3672.
- [7] Florian Eyben, Felix Weninger, Stefano Squartini, and Björn Schuller, "Real-life voice activity detection with lstm recurrent neural networks and an application to hollywood movies," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 483–487.
- [8] Shuo-Yiin Chang, Bo Li, Gabor Simko, Tara N Sainath, Anshuman Tripathi, Aäron van den Oord, and Oriol Vinyals, "Temporal modeling using dilated convolution and gating for voice-activity-detection," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 5549–5553.
- [9] Marek Hruš and Zbyněk Zajíc, "Convolutional neural network for speaker change detection in telephone speaker diarization system," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 4945–4949.
- [10] Ruiqing Yin, Hervé Bredin, and Claude Barras, "Speaker change detection in broadcast tv using bidirectional long short-term memory networks," in *Interspeech*, 2017, pp. 3827–3831.
- [11] Gregory Sell and Daniel Garcia-Romero, "Speaker diarization with plda i-vector scoring and unsupervised calibration," in *2014 IEEE Spoken Language Technology Workshop*, 2014, pp. 413–417.
- [12] Stephen Shum, Najim Dehak, Reda Dehak, and James R Glass, "Unsupervised speaker adaptation based on the cosine similarity for text-independent speaker verification.," in *Odyssey 2010 The Speaker and Language Recognition Workshop*, 2010.
- [13] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 5329–5333.
- [14] Weicheng Cai, Jinkun Chen, and Ming Li, "Exploring the encoding layer and loss function in end-to-end speaker and language recognition system," in *Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 74–81.
- [15] Weicheng Cai, Jinkun Chen, and Ming Li, "Analysis of length normalization in end-to-end speaker verification system," in *Interspeech*, 2018, pp. 3618–3622.
- [16] Simon JD Prince and James H Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *2007 IEEE 11th International Conference on Computer Vision*, 2007, pp. 1–8.
- [17] Patrick Kenny, Themis Stafylakis, Pierre Ouellet, Md Jahangir Alam, and Pierre Dumouchel, "Plda for speaker verification with utterances of arbitrary duration," in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 7649–7653.
- [18] Qingjian Lin, Ruiqing Yin, Ming Li, Hervé Bredin, and Claude Barras, "Lstm based similarity measurement with spectral clustering for speaker diarization," in *Interspeech*, 2019, pp. 366–370.
- [19] Quan Wang, Carlton Downey, Li Wan, Philip Andrew Mansfield, and Ignacio Lopez Moreno, "Speaker diarization with lstm," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 5239–5243.
- [20] Gregory Sell, David Snyder, Alan McCree, Daniel Garcia-Romero, Jesús Villalba, Matthew Maciejewski, Vimal Manohar, Najim Dehak, Daniel Povey, Shinji 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.
- [21] Aonan Zhang, Quan Wang, Zhenyao Zhu, John Paisley, and Chong Wang, "Fully supervised speaker diarization," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 6301–6305.
- [22] Qiujia Li, Florian L Kreyszig, Chao Zhang, and Philip C Woodland, "Discriminative neural clustering for speaker diarisation," in *arXiv preprint arXiv:1910.09703*, 2019.
- [23] Mireia Diez, Lukás Burget, and Pavel Matejka, "Speaker diarization based on bayesian hmm with eigenvoice priors.," in *Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 147–154.
- [24] Gregory Sell and Daniel Garcia-Romero, "Diarization resegmentation in the factor analysis subspace," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2015, pp. 4794–4798.

- [25] Ondrej Novotný, Karel Veselý, Ondrej Glembek, Oldrich Plchot, Ladislav Mošner, and Pavel Matejka, “But system for dihard speech diarization challenge 2018,” in *Interspeech*, 2018, pp. 2798–2802.
- [26] Neville Ryanta, Erika Bergelson, Kenneth Church, Alejandrina Cristia, Jun Du, Sriram Ganapathy, Sanjeev Khudanpur, Diana Kowalski, Mahesh Krishnamoorthy, Rajat Kulshreshta, et al., “Enhancement and analysis of conversational speech: Jsalt 2017,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 5154–5158.
- [27] Neville Ryant, Kenneth Church, Christopher Cieri, Alejandrina Cristia, Jun Du, Sriram Ganapathy, and Mark Liberman, “The second dihard diarization challenge: Dataset, task, and baselines,” in *Interspeech*, 2019, pp. 978–982.
- [28] E Bergelson, “Bergelson seedlings homebank corpus,” 2016.
- [29] Ryant et al, “Dihard corpus,” in *Linguistic Data Consortium*, 2019.
- [30] Arsha Nagrani, Joon Son Chung, and Andrew Zisserman, “Voxceleb: a large-scale speaker identification dataset,” in *Interspeech*, 2017, pp. 2616–2620.
- [31] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al., “The ami meeting corpus: A pre-announcement,” in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [32] Adam Janin, Don Baron, Jane Edwards, Dan Ellis, David Gelbart, Nelson Morgan, Barbara Peskin, Thilo Pfau, Elizabeth Shriberg, Andreas Stolcke, et al., “The icisi meeting corpus,” in *2003 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2003, pp. I–I.
- [33] David Snyder, Guoguo Chen, and Daniel Povey, “Musan: A music, speech, and noise corpus,” in *arXiv preprint arXiv:1510.08484*, 2015.
- [34] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 5220–5224.
- [35] “The webrtc project,” 2011, [Online]. Available: <https://webrtc.org/>.
- [36] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al., “The kaldic speech recognition toolkit,” in *IEEE 2011 workshop on automatic speech recognition and understanding*, 2011.
- [37] Weicheng Cai, Danwei Cai, Jun Zhang, and Ming Li, “Global temporal-frequential statistics pooling for text-independent speaker verification,” submitted to *IEEE Signal Processing Letter*.
- [38] K Chidananda Gowda and G Krishna, “Agglomerative clustering using the concept of mutual nearest neighbourhood,” *Pattern recognition*, vol. 10, no. 2, pp. 105–112, 1978.
- [39] Ulrike Von Luxburg, “A tutorial on spectral clustering,” *Statistics and computing*, vol. 17, no. 4, pp. 395–416, 2007.