

# UWB-NTIS Speaker Diarization System for the DIHARD II 2019 Challenge

Zbyněk Zajíc<sup>1</sup>, Marie Kunešová<sup>1,2</sup>, Marek Hruží<sup>1</sup>, Jan Vaněk<sup>1</sup>

University of West Bohemia  
Faculty of Applied Sciences

<sup>1</sup>NTIS - New Technologies for the Information Society and <sup>2</sup>Dept. of Cybernetics,  
Univerzitní 8, 306 14 Plzeň, Czech Republic

{zzajic, mkunes, mhruz, vanekej}@ntis.zcu.cz

## Abstract

In this paper, we present our system developed by the team from the New Technologies for the Information Society (NTIS) research center of the University of West Bohemia in Pilsen, for the Second DIHARD Speech Diarization Challenge. The base of our system follows the currently-standard approach of segmentation, i/x-vector extraction, clustering, and resegmentation. The hyperparameters for each of the subsystems were selected according to the domain classifier trained on the development set of DIHARD II. We compared our system with results from the Kaldi diarization (with i/x-vectors) and combined these systems. At the time of writing of this abstract, our best submission achieved a DER of 23.47% and a JER of 48.99% on the evaluation set (in Track 1 using reference SAD).

**Index Terms:** speaker diarization, i-vector, x-vector, agglomerative hierarchical clustering, neural network classifier, speaker change detection.

## 1. Introduction

In recent years, we have developed our Speaker Diarization (SD) system [1, 2, 3]. Last year, we also tailored an off-line system for the First DIHARD Speech Diarization Challenge [4], where we participated in Track 1 as well as Track 2 of the challenge.

The Second DIHARD Challenge [5] brought us an opportunity to extend our system and to try combining results from different sources. Besides i-vector [6, 7] extraction, we have also extracted x-vectors [8]. The main novelty in our system is an early-fusion of i-vectors and x-vectors into xi-vectors and modified segmentation. We also use the Kaldi system<sup>1</sup> with a recipe for diarization as an additional system for comparison and for a domain-specific combination of results. This year we have decided to participate only in the Track 1 part of the challenge, where a reference speech labeling is available and no speech activity detection (SAD) step is needed.

Our main advantage in the First DIHARD Challenge was the application of a Neural Network (NN) based domain classifier that allows the system to automatically identify the domain of each recording and to set the system's configuration accordingly. The same applies for the Second DIHARD Challenge.

## 2. Speaker Diarization System

Our system follows x-vector- and i-vector-based approaches [9, 10, 8]. A diagram of our diarization system is shown in Figure 1. The general structure is the same as in our previous system for the First DIHARD Challenge [4].

<sup>1</sup><http://kaldi-asr.org/>

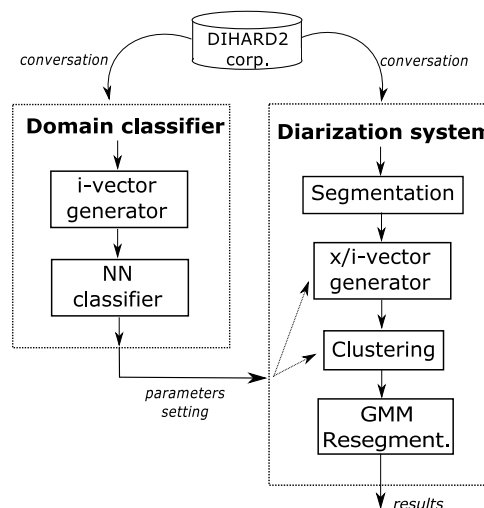


Figure 1: Diagram of the diarization process.

This section provides a description of the main steps of the diarization process. The domain classifier and related domain-dependent settings are described in section 3.

### 2.1. Segmentation

First, the entire conversation is split into multiple individual speech regions by breaking it on any non-speech events; these non-speech regions are excluded from subsequent processing. As a second step, the speech regions are further segmented according to the probability of speaker change given by our Speaker Change Detector, which is based on a Convolutional Neural Network (SCD-CNN) [2].

The SCD-CNN was trained as a regressor on spectrograms of the acoustic signal. The process is described in our previous paper for the First DIHARD Challenge [4]. The signal between two detected speaker changes is considered to be one segment.

To ensure that each segment contains sufficient information about the speaker, we set the minimum duration of each segment to 0.5 s. Shorter segments are discarded from the clustering stage and the decision about the speaker is left for the resegmentation step (see section 2.5).

### 2.2. Feature Extraction

We use the same signal processing pipeline as for the First DIHARD challenge [4] – Linear Frequency Cepstral Coefficients (LFCCs). As a newly added step, Cepstral Mean Normalization

(CMN) is applied to compensate for channel variations.

### 2.3. Segment Description

Each segment is represented by a concatenation of the  $x$ -vector and  $i$ -vector for the same segment. We call this an  $xi$ -vector. The  $i$ - and  $x$ -vectors were obtained using a Kaldi recipe<sup>2</sup> [11]. To exploit the ability of  $x/i$ -vectors to represent small amounts of data and to minimize the presence of more than one speaker in a segment, longer segments are split into intervals of max. 2 s, with 1 s overlaps.

For generating  $i$ -vectors, we trained a UBM with 2048 components and a transformation matrix with a latent dimension of 128. A Time Delay Neural Network is used as an  $x$ -vector extractor, and  $x$ -vectors are extracted from the affine component of the second-to-last layer with dimension 128.

For whitening the  $xi$ -vectors, we subtract the mean of development set's  $xi$ -vectors. During the diarization process, we use a conversation-dependent Principal Component Analysis (PCA) [12] computed on the data in the current conversation to reduce the dimension of the  $xi$ -vectors into 3, 6, 9 or 12 (depending on the identified corpus – see Tab. 1).

The following corpora were used as training data: LibriSpeech<sup>3</sup>, VoxCeleb2 [13], TedLium3 [14], and the following ELRA corpora: Speecon database UK English (ELRA-S0215) and US English (ELRA-S0233). Additional data augmentation (additive noise, music, babble and reverberation) was used on the LibriSpeech and TedLium3 corpora.

### 2.4. Clustering

The general clustering approach remains mostly unchanged from what we used in the First DIHARD Challenge.

As in the previous year's challenge, the number of speakers in each recording is unknown in advance: the DIHARD corpus consists of several distinct domains where the number of speakers ranges from 1 to 10. Thus, we have chosen to primarily use the agglomerative hierarchical clustering (AHC) algorithm.

The clustering process is based on the average cosine distance between  $xi$ -vectors. We use a fixed distance threshold as the stopping condition. Additionally, we leverage our knowledge from the development set by also specifying a minimum and maximum number of clusters for each conversation, based on the number of speakers typically observed for the domain. So, we force the final number of clusters to be within this set range. The range and the clustering stopping condition – the distance threshold – were both established on a per-corpus basis using the development data (see Section 6).

In the Second DIHARD Challenge, there are three corpora in the development set with an overwhelming majority of two-speakers conversations. For these domains, we use a Probabilistic Linear Discriminant Analysis (PLDA) model [15] to evaluate the distance between the  $xi$ -vectors. Then, we apply  $k$ -medoids clustering into a constant number of clusters across the whole domain. For training the PLDA model, we used the same datasets as listed in Subsection 2.3. The between-class dimension is equal to the feature dimension.

### 2.5. Resegmentation

Finally, we refine the diarization results via resegmentation. The previous results are based on relatively long signal windows

<sup>2</sup>[https://github.com/kaldi-asr/kaldi/tree/master/egs/callhome\\_diarization/v1](https://github.com/kaldi-asr/kaldi/tree/master/egs/callhome_diarization/v1) and [/v2](https://github.com/kaldi-asr/kaldi/tree/master/egs/callhome_diarization/v2)

<sup>3</sup><http://www.openslr.org/12/>

and the boundaries between speaker segments are not precise. Therefore, we compute a GMM for each speaker cluster from all feature vectors assigned to the speaker. Then, likelihoods for all speaker GMMs are evaluated and filtered by a Gaussian window (length 75 ms with shift 50 ms) to smooth peaks. The number of GMM components ranges between 1 and 64 depending on the speaker data size. Then, the entire conversation is reclassified according to the GMM likelihoods on a frame-by-frame basis.

## 3. Domain Classification

The DIHARD II corpus [16] consists of data taken from several different domains, with very diverse characteristics – including the number of speakers, the level of noise, and audio quality in general. To improve the results of the diarization system, we decided to use the supervised information about each domain given by the organizers (the possible number of speakers in the domain) and to tune specific settings on development data (mainly the threshold for AHC).

We have proposed the domain classifier as a hierarchical two-stage classifier. The first level is a special classifier to distinguish recordings with one speaker from multi-speaker data. The second stage classifier is applied when the first level class is considered to be the multiple speaker case. It evaluates the posterior probability that the input conversation belongs to one of the 11 corpora in the DIHARD II development set.

The same NN architecture is used for both classifiers. They differ only in the last layer, where the first level NN uses one neuron for the binary classification, and the second stage NN uses a softmax layer with 11 classes.

The NNs receive a single  $i$ -vector calculated over the entire conversation as the input. A special  $i$ -vector extractor was trained for the domain classifiers. The LFCC features are the same but also include frames marked as non-speech by SAD. The UBM has 512 diagonal components, and the final  $i$ -vector dimension is 100.

The NNs have one hidden layer with 2048 neurons and tanh activation function. Dropout with coefficient 0.9 was used during training. The network was implemented in TensorFlow, where the “adam” optimizer was used with 10 epochs and batch size of 32. The remaining hyper-parameters were left at default values.

Both NNs were trained on the development data + 10 randomly chosen recordings from the LibriSpeech corpus (cut to 10min length).

Because the evaluation dataset contains two unknown corpora with very different characteristics, we apply a threshold on the classifier result. The positive detection threshold was set to 0.6 for both stage classifiers. If there is no positive detection in the second-stage NN, the conversation is treated as “unknown domain”. The accuracy of the first-stage NN was 100% on a held-out part of the development data. The accuracy of the second-stage NN was 82%.

### 3.1. Domain-specific settings

Because of our domain classifier, we were able to use different system configuration for each of the 11 development set corpora and for unknown data. Here we describe the general approaches we selected for each domain. Specific experimentally-chosen parameters are listed in Table 1.

The descriptions of the individual DIHARD II corpora can be found in the challenge evaluation plan [5] as well as in the

main challenge paper [17]. As such, we do not replicate them here.

*LibriVox*: All recordings contain only 1 speaker. Thus, we did not need to perform diarization but simply used the information given by reference SAD. We rely on the first stage of the domain specific classifier.

*SEEDLingS, SCOTUS, RT-04S, SLX, VAST and YouthPoint*: For these corpora, we used the AHC approach with cosine distance, as described in section 2.4.

*ADOS, DCIEM and MIXER6*: These corpora have almost exclusively exactly 2 speakers in each conversation. For this reason, we could simply use k-medoids clustering into 2 clusters with PLDA scores.

*CIR*: For this corpus, our system gives the best results on development data for k-medoids clustering into 4 clusters with PLDA scores.

*Unknown*: For unrecognized evaluation data, we’ve chosen to use AHC with 2-6 target clusters.

Table 1: *Experimentally chosen parameters (Thr. = threshold, k-m = k-medoids) for each corpus.*

Corpus	Clustering	No. spk	Thr AHC	PCA dim
LibriVox	-	1	-	-
SEEDL.	AHC	2-3	0.62	6
CIR	k-m	4	-	-
ADOS	k-m	2	-	-
SCOTUS	AHC	5-10	0.46	12
DCIEM	k-m	2	-	-
RT-04S	AHC	3-10	0.46	6
SLX	AHC	2-6	0.762	6
MIXER6	k-m	2	-	-
VAST	AHC	1-9	0.58	3
YouthP.	AHC	3-5	0.54	9
other	AHC	2-6	0.1	-

## 4. Kaldi Diarization System

As an additional system, we have decided to use a Kaldi recipe for diarization [11]. The input features are the same LFCCs as in our system (details in Section 2.2).

The segmentation provides chunks of speech between important non-speech events (Kaldi SAD segmentation) and subsequently divides these segments into sub-segments with constant length 1.5 s and overlap 0.75 s (the minimum length of a segment is 0.5 s).

X-vectors or i-vectors are computed on segmented data and handled by a PLDA model to compute the similarity between these segments. X/i-vectors are whitened before the PLDA estimation by subtracting the mean and transforming by an LDA matrix.

The vectors of segments are then clustered according to the AHC, with the stopping threshold set on development data. This threshold was found for the entire development set – this system does not treat different domains of the DIHARD II corpus differently.

Additionally, we create xi-vectors by concatenating the x-vector and i-vector for the same segment. The whitening transformation is also obtained by concatenating the means and PCA transformation matrices belonging to the x/i-vectors (this inde-

pendent treatment works better than computing the whitening transformation on the whole xi-vectors).

## 5. Late system combination

The Kaldi diarization system does not use the information from the domain classifier, and its setting is very general. Therefore, we have used the Kaldi system instead of our speaker diarization system (SD) in the cases where the domain classifier marked the conversation as “unknown domain”.

For the two most problematic corpora (Seedlings, VAST), we have also used the Kaldi system. On average, our system slightly outperforms the Kaldi system on the development data for these corpora. However, the DERs of individual conversations have a higher variance than the ones from Kaldi. We refer to this per-domain system selection as late system combination.

## 6. Experiments

This section describes our experiments on the development set of the Second DIHARD Challenge, as well as our final results on the evaluation set. The experiments mainly served for finding the optimal system configuration for each of the individual corpora. For details of the DIHARD II corpus [16, 18], see the evaluation plan [5].

### 6.1. Evaluation

The system performance was evaluated in terms of Diarization Error Rate (DER), as defined by NIST [19]. On the development set, we calculated this on a per-recording basis using NIST’s md-eval.pl script<sup>4</sup>.

DER and Jaccard Error Rate (JER) on the evaluation set were given by the official scoring system [17].

Unlike usual practice, DIHARD Challenge submissions were scored with no forgiveness collar around speaker boundaries, and overlapping speech was included in the evaluation.

### 6.2. Results

Table 2 presents a comparison between i/x/xi-vectors contribution with an earlier version of our system. Based on these preliminary results, the rest of our work was with xi-vectors only.

Table 3 shows results on the development set for each of the eleven corpora. Table 4 then presents the final results on the evaluation data for Track 1 – diarization using reference SAD.

Table 2: *Average DER [%] on DIHARD II development set for an earlier version of our system and for Kaldi with different segment descriptors (x/i/xi-vector).*

system	DER
SD i-vec	24.31
SD x-vec	23.81
SD xi-vec	<b>22.51</b>
Kaldi i-vec	25.83
Kaldi x-vec	25.32
Kaldi xi-vec	24.13

<sup>4</sup><https://github.com/usnistgov/SCTK/blob/master/src/md-eval>

Table 3: Average DER [%] on individual corpora of the DIHARD II development set, for our system (SD), Kaldi system, and the combination system, all using xi-vectors.

Corpus	SD	Kaldi	Comb.
LibriVox	0.00	14.52	0.0
SEEDLingS	31.32	33.90	33.90
CIR	45.83	52.25	45.83
ADOS	14.06	16.01	14.06
SCOTUS	6.92	18.03	6.92
DCIEM	8.88	9.65	8.88
RT-04S	33.14	36.30	33.14
SLX	17.56	16.90	17.56
MIXER6	9.42	9.72	9.42
VAST	38.00	39.65	39.65
YouthPoint	4.55	6.33	4.55
All	<b>20.78</b>	24.13	21.29

Table 4: Official results (DER [%] and JER [%]) on the DIHARD II evaluation data for our system (SD), Kaldi, and Comb.

	SD	Kaldi	Comb.
DER[%]	24.59	25.17	<b>23.47</b>
JER[%]	49.63	54.94	<b>48.99</b>

## 7. Discussion

For this challenge, we decided to extend our previous system for diarization with various enhancements; this chapter discusses their benefits. Table 5 presents the results for our SD system with comparison with several alterations: a system without SCD-CNN (only Kaldi SAD segmentation), with de-noised test data (using speech enhancement<sup>5</sup> [20], no de-noising for training data), and a system with reference information about overlapping speech. The latter information is used in the re-segmentation step – for such parts of the data, the second most likely speaker is also detected.

Table 6 shows the segmentation coverage and purity [21] achieved on the development set for SCD-CNN and Kaldi SAD segmentation. Our SCD-CNN method outperforms the Kaldi SAD segmentation in both coverage and purity, except two cases. This result is expected, as SAD segmentation does not take into account the speaker changes in one speech activity segment. The purity measure is more important in the diarization task, as it tells us how pure the segments are – that they contain only one speaker.

Based on these results, we chose to apply the SCD-CNN in our final system, and we omitted the de-noising. Unfortunately, despite the clear benefits of detecting overlapping speech, we were not able to train a real overlap detector with reasonable accuracy on the DIHARD II data.

## 8. Conclusion

In this paper, we presented a new version of our diarization system and its results for the Second DIHARD Diarization Challenge. Compared to our previous system, we applied xi-vectors and modified the SCD-based segmentation step to take advantage of x/i-vectors’ ability to represent short segments. Using

<sup>5</sup>[https://github.com/staplesinLA/denoising\\_DIHARD18](https://github.com/staplesinLA/denoising_DIHARD18)

Table 5: Average DER [%] on individual corpora of the DIHARD II development set, for our system (SD) with different setting – with Kaldi SAD segmentation instead of SCD-CNN, test data de-noise and with reference overlap labels.

Corpus	SD	without SCD	with de-noise	with ref. overlap
LibriVox	0.00	0.0	0.0	0.0
SEEDLingS	31.32	31.22	32.30	24.56
CIR	45.83	47.88	46.70	37.71
ADOS	14.06	13.26	14.25	10.73
SCOTUS	6.92	10.67	8.01	5.99
DCIEM	8.88	8.66	8.74	6.24
RT-04S	33.14	36.38	34.53	25.69
SLX	17.56	19.14	17.36	13.64
MIXER6	9.42	9.29	9.93	5.02
VAST	38.00	38.91	38.61	30.09
YouthPoint	4.55	5.26	5.49	3.89
All	20.78	21.52	21.31	16.16

Table 6: Overall Coverage and Purity for segments provided by Kaldi (SAD) and by SCD-CNN

Corpus	(a) SAD segments		(b) SCD-CNN	
	Cov	Pur	Cov	Pur
VAST	<b>0.945</b>	0.454	0.926	<b>0.820</b>
YP	0.950	0.445	<b>0.991</b>	<b>0.971</b>
LIBRIVOX	0.948	0.466	<b>0.987</b>	<b>0.991</b>
SEEDLINGS	0.917	0.600	<b>0.937</b>	<b>0.868</b>
CIR	0.963	0.345	<b>0.979</b>	<b>0.588</b>
ADOS	0.896	0.669	<b>0.992</b>	<b>0.875</b>
SCOTUS	0.980	0.375	<b>0.990</b>	<b>0.936</b>
DCIEM	0.901	0.563	<b>0.987</b>	<b>0.862</b>
RT04	<b>0.969</b>	0.421	0.889	<b>0.670</b>
SLX	0.916	0.507	<b>0.988</b>	<b>0.871</b>
MIXER6	0.933	0.440	<b>0.958</b>	<b>0.868</b>

a domain classifier as in the previous challenge, we were able to use a different system configuration for each subset of data. For comparison, we applied a Kaldi recipe for diarization and combined the results into a single system. Additionally, we have investigated the potential gains of detecting overlapping speech, de-noising, and various segmentation methods. Our best Track 1 submission achieved a DER of 23.47% and JER of 48.99%.

## 9. Acknowledgements

This research was supported by the Ministry of Education, Youth and Sports of the Czech Republic project, No. LO1506, by the Technology Agency of the Czech Republic, project No. TN01000024 and by the grant of the University of West Bohemia project, No. SGS-2019-027. Access to computing and storage facilities provided by the project CESNET LM2015042 is greatly appreciated.

## 10. References

- [1] Z. Zajíč, M. Kunešová, and V. Radová, “Investigation of Segmentation in i-Vector Based Speaker Diarization of Telephone Speech,” in *Speccom*. Budapest: Springer, 2016, pp. 411–418.

- [2] M. Hružík and Z. Zajíc, “Convolutional Neural Network for Speaker Change Detection in Telephone Speaker Diarization System,” in *ICASSP*. New Orleans: IEEE, 2017, pp. 4945–4949.
- [3] Z. Zajíc, M. Hružík, and L. Müller, “Speaker diarization using convolutional neural network for statistics accumulation refinement,” in *Interspeech*, Stockholm, 2017, pp. 3562–3566.
- [4] Z. Zajíc, M. Kunešová, J. Zelinka, and M. Hružík, “ZCU-NTIS Speaker Diarization System for the DIHARD 2018 Challenge,” in *Interspeech*, Hyderabad, 2018, pp. 2788–2792.
- [5] N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, and M. Liberman, “Second DIHARD Challenge Evaluation Plan,” Linguistic Data Consortium, Tech. Rep., 2019.
- [6] P. Kenny, P. Ouellet, N. Dehak, V. Gupta, and P. Dumouchel, “A Study of Interspeaker Variability in Speaker Verification,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 5, pp. 980–988, 2008.
- [7] L. Machlica and Z. Zajíc, “Factor Analysis and Nuisance Attribute Projection Revisited,” in *Interspeech*, vol. 2, Portland, 2012, pp. 1570–1573.
- [8] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, “X-Vectors: Robust DNN Embeddings for Speaker Recognition,” in *ICASSP*, 2018, pp. 5329–5333.
- [9] G. Sell and D. Garcia-Romero, “Speaker Diarization with PLDA I-vector Scoring and Unsupervised Calibration,” in *IEEE Spoken Language Technology Workshop*, South Lake Tahoe, 2014, pp. 413–417.
- [10] M. Senoussaoui, P. Kenny, T. Stafylakis, and P. Dumouchel, “A Study of the Cosine Distance-Based Mean Shift for Telephone Speech Diarization,” *Audio, Speech and Language Processing*, vol. 22, no. 1, pp. 217–227, 2014.
- [11] G. Sell, D. Snyder, A. Mccree, D. Garcia-romero, V. Manohar, N. Dehak, D. Povey, S. Watanabe, and S. Khudanpur, “Diarization is Hard : Some Experiences and Lessons Learned for the JHU Team in the Inaugural DIHARD Challenge,” in *Interspeech*, Hyderabad, 2018, pp. 2808–2812.
- [12] S. Shum, N. Dehak, E. Chuangsuwanich, D. Reynolds, and J. Glass, “Exploiting Intra-Conversation Variability for Speaker Diarization,” in *Interspeech*, Florence, 2011, pp. 945–948.
- [13] J. S. Chung, A. Nagrani, and A. Zisserman, “VoxCeleb2: Deep Speaker Recognition,” in *Interspeech*, Hyderabad, 2018, pp. 1086—1090.
- [14] F. Hernandez, V. Nguyen, S. Ghannay, N. Tomashenko, and Y. Estève, “TED-LIUM 3: Twice as Much Data and Corpus Repartition for Experiments on Speaker Adaptation,” in *Speccom*. Leipzig: Springer, Cham, sep 2018, pp. 198–208.
- [15] S. Ioffe, “Probabilistic Linear Discriminant Analysis,” *Lecture Notes in Computer Science*, vol. 3954 LNCS, pp. 531–542, 2006.
- [16] N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, and M. Liberman, “DIHARD Corpus,” LDC, Tech. Rep., 2019.
- [17] —, “The Second DIHARD Diarization Challenge: Dataset, task, and baselines,” in *Interspeech*, Gratz, 2019.
- [18] E. Bergelson, “Bergelson Seedlings HomeBank Corpus,” Linguistic Data Consortium, Tech. Rep., 2018.
- [19] J. G. Fiscus, N. Radde, J. S. Garofolo, A. Le, J. Ajot, and C. Laprun, “The Rich Transcription 2006 Spring Meeting Recognition Evaluation,” *Machine Learning for Multimodal Interaction*, vol. 4299, pp. 309–322, 2006.
- [20] L. Sun, J. Du, C. Jiang, X. Zhang, S. He, B. Yin, and C.-h. Lee, “Speaker Diarization with Enhancing Speech for the First DIHARD Challenge,” in *Interspeech*, 2018, pp. 2793–2797.
- [21] H. Bredin, “TristouNet: Triplet Loss for Speaker Turn Embedding,” in *ICASSP*. New Orleans: IEEE, 2017, pp. 5430–5434.