



HAL
open science

Robust acoustic domain identification with its application to speaker diarization

A Kishore Kumar, Shefali Waldekar, Md Sahidullah, Goutam Saha

► **To cite this version:**

A Kishore Kumar, Shefali Waldekar, Md Sahidullah, Goutam Saha. Robust acoustic domain identification with its application to speaker diarization. *International Journal of Speech Technology*, 2022, 25 (December), pp.933-945. 10.1007/s10772-022-09990-9 . hal-03719697

HAL Id: hal-03719697

<https://inria.hal.science/hal-03719697v1>

Submitted on 11 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust Acoustic Domain Identification with its Application to Speaker Diarization

A Kishore Kumar¹ · Shefali Waldekar² ·
Md Sahidullah³ · Goutam Saha¹

the date of receipt and acceptance should be inserted later

Abstract With the rise in multimedia content over the years, more variety is observed in the recording environments of audio. An audio processing system might benefit when it has a module to identify the acoustic domain at its front-end. In this paper, we demonstrate the idea of *acoustic domain identification* (ADI) for *speaker diarization*. For this, we first present a detailed study of the various domains of the third DIHARD challenge highlighting the factors that differentiated them from each other. Our main contribution is to develop a simple and efficient solution for ADI. In the present work, we explore speaker embeddings for this task. Next, we integrate the ADI module with the speaker diarization framework of the DIHARD III challenge. The performance substantially improved over that of the baseline when the thresholds for agglomerative hierarchical clustering were optimized according to the respective domains. We achieved a relative improvement of more than 5% and 8% in DER for core and full conditions, respectively, on Track 1 of the DIHARD III evaluation set.

Index Terms: acoustic domain identification (ADI), agglomerative hierarchical clustering (AHC), diarization error rate (DER), Jaccard error rate (JER), speaker diarization.

¹ Department of Electronics & ECE, Indian Institute of Technology Kharagpur, India.
E-mail: kishore@iitkgp.ac.in, gsaha@ece.iitkgp.ac.in

² Department of Electrical, Electronics & Communication Engineering, GITAM School of Technology, GITAM (Deemed to be University), Bengaluru, India.
E-mail: swaldeka@gitam.edu

⁴ Université de Lorraine, CNRS, Inria, LORIA, F-54000, Nancy, France.
E-mail: md.sahidullah@inria.fr

1 Introduction

Spoken documents refer to audio recordings of multi-speaker events like lectures, meetings, news, etc. We use *acoustic domain* as a descriptor for the type of spoken document. Although it sounds similar to *acoustic scene*, it is not so because the latter concerns only the background environment of the audio which may or may not have speech (Barchiesi et al. 2015). In spoken documents, speech is the main source of information while other elements such as location and channel form the metadata. Examples of acoustic domains are broadcast news, telephone conversations, web video, movies, and so on, which may have different backgrounds across the data from same class, or even in the same document. Nowadays, easily accessible and simple to use audio recorders have led to an exponential increase in the amount of speech recordings which are available in myriad conditions in terms of environment, recording equipment, number of speakers, to name a few. An artificial intelligence system can deal with the multi-domain scenarios either by normalizing to the domains, i.e. *domain-invariant processing*, or by adapting to the domains i.e. *domain-dependent processing*. For diverse and challenging recording conditions, it is difficult to design a system following the first strategy. Therefore, in such cases, a system that identifies the domain and then proceeds with relevant information extraction accordingly is expected to perform better than a *one-size-fits-all* method.

For *rich transcription* (RT) of spoken documents, besides “what is being spoken” and “by whom”, it is also imperative to know “who spoke when”. The answers to the three questions are provided by automatic speech recognition, speaker recognition, and speaker diarization, respectively. *Speaker diarization* is the task of generating time-stamps with respect to the speaker labels (Anguera et al. 2012; Park et al. 2022). In earlier studies, RT researchers mainly considered recordings of broadcast news, telephone conversations and meetings/conferences (Anguera et al. 2012). The NIST RT challenges (RT02 - RT09)¹ were organized for metadata extraction from these acoustic domains. There, the type of data was known beforehand which implies that the diarization systems were designed to suit the application. The factors mainly contributing to poor diarization were number of speakers, speech activity detection, amount of overlap, and speaker modelling (Mirghafori and Wooters 2006; Huijbregts and Wooters 2007; Sinclair and King 2013).

In this paper, we propose a simple but efficient method for acoustic domain identification (ADI). The proposed ADI system uses speaker embeddings. Our study reveals that i-vector and OpenL3 embedding based method achieves considerably better performance than x-vector based approach in the third DIHARD challenge dataset. Of all the three, i-vector system’s results were found to be better. Next, we show that domain-dependent threshold for speaker clustering helps to improve the diarization performance only when *probabilistic linear discriminant analysis* (PLDA) adaptation uses audio-data

¹ <https://www.nist.gov/itl/iad/mig/rich-transcription-evaluation>

from all the domains. The proposed system stands apart from other submissions to the challenge because unlike most of the top-performers, it is not an ensemble or fusion-based system. Moreover, from the study of speaker diarization literature, we found that the domain-dependent processing approach is the first of its kind (Sahidullah et al. 2019).

In the upcoming sections, first, we talk about previous studies relevant to the topic, i.e. acoustic scene classification and speaker diarization in Section 2. Then, Section 3 presents a detailed description of the speaker diarization data used in our study. The major aspects of the contribution of this work, that is, ADI and classification of domains are discussed in Section 4 and Section 6 respectively. The experimental setup is outlined in Section 7 while the results of the experiments are presented and discussed in Section 8. We conclude the paper in Section 9.

2 Related works

In the classical audio processing applications, where speech is regarded as the primary information source, the domain was kept constant to get the best performance. However, in the present times of technological advancements, multi-domain data is unavoidable, whether it is for speech recognition, language recognition, speaker recognition, or speaker diarization. In this paper, we consider the speaker diarization (SD) application. The research in this area took a leap two decades ago with the NIST RT evaluations. Back then, broadcast news (BN), conversational telephone speech (CTS), and meetings/-conferences were the only three acoustic domains considered.

The use of *total variability* (TV) space was proposed in (Shum et al. 2011) to exploit intra-conversation variability. The same authors later applied Bayesian-Gaussian mixture model (GMM) to *principal component analysis* (PCA)-processed i-vectors as a probabilistic approach to speaker clustering (Shum et al. 2013). The authors of (Sell and Garcia-Romero 2014) incorporated PLDA for i-vector scoring and used unsupervised calibration of the PLDA scores to determine the clustering stopping criterion. They also showed improved diarization by denser sampling in the i-vector space with overlapping temporal segments. A *deep neural network* (DNN) architecture was introduced by (Garcia-Romero et al. 2017) that learned a fixed-dimensional embedding for variable length acoustic segments along with a scoring function for measuring the likelihood that the segments originated from the same or different speakers. Combination of long short-term memory based speaker-discriminative d-vector embeddings and non-parametric clustering was done by (Wang et al. 2018) while unbounded interleaved-state *recurrent neural network* (RNN) with d-vectors were employed by (Zhang et al. 2019) for SD.

The annual DIHARD challenges focusing on *hard diarization* revamped SD research. The objective is to build SD systems for spoken documents from diverse and challenging domains where the existing state-of-the-art systems fail to deliver. The first one was held in 2018 (Ryant et al. 2018). The data

for this challenge was taken from nine domains. The top performing system (Sell et al. 2018) explored the i-vector and x-vector representations and found that the latter trained on wideband microphone data was better. They also showed that a *speech activity detection* (SAD) system specifically trained for SD with in-domain data helped in improving performance. For the second challenge, the data was drawn from 11 domains consisting of single-channel and multi-channel audio recordings (Ryant et al. 2019). This version was provided with baseline systems for speech enhancement, SAD and SD. The given SD baseline was based on the top-performing system of its predecessor. The first rank system performed *agglomerative hierarchical clustering* (AHC) over x-vectors with Bayesian hidden Markov model (Landini et al. 2020). In the third challenge, the multichannel audio condition and also the child speech data was not there, whereas CTS data was included for the first time (Ryant et al. 2020). The top performing system had speech enhancement and audio domain classification before SD (Wang et al. 2021). It combined multiple front-end techniques, such as speech separation, target-speaker based voice activity detection and iterative data purification. Its SD system was based on DIHARD II’s top-ranked system.

The challenges on detection and classification of acoustic scenes and events (DCASE) have rendered a boost to research in *acoustic scene classification* (ASC). Discussion about the state-of-the-art in this field is necessary before we go into the details of ADI because the two bear some degree of conceptual similarity. The baseline systems provided with the ASC task of the DCASE challenges have ranged from *Mel-frequency cepstral coefficients* (MFCC)-GMM based systems (Giannoulis et al. 2013; Mesaros et al. 2016), log mel-band energy with *multilayer perceptron* (MLP) (Mesaros et al. 2018a), with convolutional neural network (CNN) (Mesaros et al. 2018b, Oct. 2019) to OpenL3 embeddings with two fully-connected feed-forward layers (Heittola et al. 2020). On the other hand, the top-performing systems have shown a variety of approaches to extract as much as possible differences between the given acoustic scenes. These approaches include the use of spectrograms, *recurrence quantification analysis* (RQA) of MFCC, i-vectors, log-mel band energies, and constant-Q transform as features, with support vector machine, MLP, DNN, CNN, RNN, and trident residual neural network as classifiers (Roma et al. 2013; Eghbal-Zadeh et al. 2016; Mun et al. 2017; Sakashita and Aono 2018; Chen et al. 2019; Suh et al. 2020). A few systems have also used data augmentation techniques like generative adversarial network, graph neural network, temporal cropping and mixup (Mun et al. 2017; Sakashita and Aono 2018; Chen et al. 2019; Suh et al. 2020). The multi-classifier systems have employed majority vote, average vote, or random forest for decision making (Mun et al. 2017; Sakashita and Aono 2018; Chen et al. 2019; Suh et al. 2020). Some other noteworthy experiments have been done with various spectral and time-frequency features, classification techniques along with fusion approaches for ASC (Ren et al. 2018; Rakotomamonjy and Gasso 2015; Waldekar and Saha 2018, 2020). Matrix factorization was applied in (Bisot et al. 2017). i-vector

and x-vector embeddings were used with CNN in (Dorfer et al. 2018) and (Zeinali et al. 2018) respectively.

3 DIHARD III: A large-scale realistic speech corpora with multiple domains

In the third edition of the DIHARD challenge series, the DIHARD III data set is taken from eleven diverse domains. This brings a lot of variability in data conditions like background noise, language, source, number of speakers, amount of overlapped speech, and speaker demographics. Table 1 enlists the acoustic domains along with their corresponding source corpus, type of recording, and a brief description. The data for ‘Meeting’ and ‘Sociolinguistic (Field)’ are taken from two different speech corpora. The table also shows how the domains differ in terms of the number of speakers per sample and total speakers. The database is divided into development and evaluation sets. Both have 5–10 minute duration samples, except in the case of ‘Web videos’ where the sample duration ranges from less than 1 min to more than 10 min. Evaluation has to be done on two partitions of the evaluation data: Core evaluation set — a balanced set where the total duration of each domain is approximately equal, and Full evaluation set — all the samples are considered for each domain. The development set composition is the same as that of the evaluation set. The metadata provided for each data sample in the development set is its domain, source, language, and whether or not it was selected for the core set. This information was not provided for the evaluation data.

4 Audio domain analysis of the corpora

The heterogeneity in the recording conditions of the audio samples under consideration can be observed from the last two columns of Table 1. To gain better insights into their domain-wise differences, we analyzed the DIHARD III development data’s audio samples using three methods: *signal-to-noise ratio* (SNR) estimation, *long-term average spectrum* (LTAS) analysis, and speech-to-non-speech ratio computation.

Table 1 Description of different source domains for the DIHARD III dataset.

Domain	Source Corpora	Speakers per Recording	Total Speakers	Conversation Type	Description
Audiobooks (Ab)	LIBRIVOX	1	12	Read	Amateur readers reading aloud passages from public-domain English texts.
Broadcast Interview (Bi)	YOUTHPOINT	3-5	46	Radio Interview	Student-lead radio interviews conducted during the 1970s with then-popular figures. The audio files are recorded in a studio on open reel tapes, and these were digitized and transcribed later at LDC.
Clinical (Cl)	ADOS	2	96	Interview	Semi-structured interviews recorded with ceiling-mounted mic to identify whether the 2-16 yrs children fit the clinical autism diagnosis
Courtroom (Ct)	SCOTUS	5-10	83	Argument	Oral arguments from the U.S. Supreme Court (2001 term). The table-mounted speaker-controlled mics' outputs were summed and recorded on a single-channel reel-to-reel analog tape recorder.
CTS	FISHER	2	122	Telephonic	Ten min conversations between two native English speakers.
Map Task (Mt)	DCIEM	2	46	Formal	Pairs of speakers sat opposite each other. One communicated a route marked on a map so that the other may precisely reproduce it on his map. Speech recorded on separate mics was mixed later.
Meeting (Mg)	RT04/ROAR	3-10	75	Formal	Highly interactive recordings consisting of large amounts of spontaneous speech with frequent interruptions and overlaps, each recorded with a centrally located distant mic with small gain.
Restaurant (Rt)	CIR	5-8	86	Interactive Informal	Mix of the two channels of binaural mic recordings. Highly frequent overlap speech and interruptions along with speech from nearby tables and restaurant-typical background noise.
Sociolinguistic (Field) (Sf)	SLX/DASS	2-6	42	Interview	An interviewer tried to elicit vernacular speech from an informant during informal conversation. Most recordings are at home, while some are at public places (park or cafe).
Sociolinguistic (Lab) (Sl)	MIXER6	2	32	Interview	Controlled environment interviews recorded under quiet conditions.
Web Video (Wv)	Video-sharing sites	1-9	127	Variable	Collection of amateur videos on diverse topics and recording conditions in English and Mandarin.




Fig. 1 SNR distribution of DIHARD III development data.

4.1 Signal-to-noise ratio (SNR) analysis

SNR is the most popular measure to quantify the presence of noise in signals. In SD, noisier audio samples are expected to be difficult to diarize. *Waveform amplitude distribution analysis* (WADA) SNR estimation is a well-known technique for SNR estimation. It assumes that speech is totally independent of the background noise and clean speech always follows a gamma distribution with a fixed shaping parameter (between 0.4 and 0.5), whereas background noise follows Gaussian distribution (Kim and Stern 2008). The result of this estimation for the data in hand is shown in Fig. 1. ‘CTS’ and ‘Map task’ show considerable variation in values and are on the higher side in the box-plot. This could be attributed to having only two speakers speaking on separate channels. ‘Meeting’ and ‘Restaurant’ have spontaneous speech from multiple speakers in noisy environment, recorded with one mic contributing to the low SNR.

4.2 Long-term average spectrum (LTAS) analysis

With LTAS, we are trying to gain knowledge of the spectral distribution of the signals over a period of time (Löfqvist 1986). LTAS tends to average out segmental variations (Assmann and Summerfield 2004). This information is displayed in Fig. 2 for DIHARD III development data in domain-wise manner. The fact that speech signals generally have no energy below 300 Hz but are low-frequency signals is discernible in all the plots. The peculiar nature of ‘CTS’ after 4 kHz is because telephonic audio is sampled at 8 kHz. Controlled and quiet recording conditions like that of ‘Map task’ and ‘Sociolinguistic (Lab)’ show lower LTAS values. Note that the latter had low SNR too. On the other hand, multi-speaker and noisy environments of ‘Courtroom’, ‘Restaurant’, and ‘Web videos’ exhibit higher magnitudes.




Fig. 2 Comparison of long-term average spectrum (LTAS) of speech files from each domain.




Fig. 3 Domain-wise percentage of Speech to non-speech average duration for DIHARD III development data.

4.3 Speech-to-nonspeech ratio

The final goal of this work is to assign time-stamps to spoken documents. So, the information regarding the amount of speech in the audio samples would be worthwhile. The speech-to-nonspeech ratio offers an acoustic representation of the language in daily conversations. It helps obtain additional information on the spectral energy distribution of a speech signal in a longer speech sample. As can be seen in Fig. 3, around 20% of the recorded signals from most of the domains are non-speech. The reason for exceptional behaviour of ‘Clinical’ signals could be the unusual language use by children with autism, non-autistic language disorders, and ADHD. For ‘Courtroom’ the high nonspeech value could be attributed to the participants’ control on the table-top microphone and the overall noisy environment.




Fig. 4 Percentage of overlapped speech against non-overlapped speech averaged across all the samples of each domain for DIHARD III development data.

4.4 Percentage of overlapped speech

The presence of overlapped speech in a spoken document is a serious factor affecting the performance of a diarization system. The SD systems are generally designed to assign one speaker label to a segment. Thus, when more than one speaker is present, the other speaker’s or speakers’ speech contributes to the *missed detection rate* (MDR) component of the DER. A system having the capacity to assign multiple labels would also find it challenging to characterize particular speakers when their speech is superimposed. As can be seen in Fig. 4, overlapped speech amount varies with the domain. ‘Restaurant’ has the highest value because many speakers are involved in informal conversation. In ‘Meeting’, although the conversation is formal, the speakers can still interrupt each other resulting in high percentage of overlapped speech. In spite of multiple speakers having argument in ‘Courtroom’, the reason for less overlapped speech is the same as that of high speech-to-nonspeech ratio.

5 Baseline Diarization System

A general SD system consists of an SAD, a segmentation, and a clustering module. In this work, we are concerned with only identifying the various domains of spoken documents and hence have only considered the Task 1 of DIHARD III (Ryant et al. 2020) where the reference SAD was given. The diarization baseline provided with this challenge (Ryant et al. 2021), which was based on one of the submissions of the predecessor challenge (Singh et al. 2019), was used to benchmark our proposed SD system.

The baseline system extracts 512-dimensional x-vectors using 30-dimensional MFCCs from 25 ms frames with overlap of 15 ms. Three second sliding windows are used for mean normalization. VoxCeleb 1 (Nagrani et al. 2017) and VoxCeleb 2 (Chung et al. 2018) are combined and augmented with additive

noise from MUSAN (Snyder et al. 2015) and reverberation from RIR datasets (Ko et al. 2017) for training. Segments of duration 1.5 sec with shift of 0.25 sec are used for embeddings extraction during testing. Statistics estimated from DIHARD III Dev and Eval sets are used for centering and whitening of the embeddings, post which they are length normalized.

Before scoring with Gaussian PLDA model (Prince and Elder 2007) trained from the x-vectors extracted from VoxCeleb, conversation-dependent PCA (Zhu and Pelecanos 2016) preserving 30% of the total variability is used for dimensionality reduction. Clustering is done using AHC with minimum DER for Dev set as the stopping criteria (Han et al. 2008). The output is then refined by using *variational Bayes hidden Markov model* (VB-HMM) re-segmentation (Diez et al. 2019) which is initialized separately for each recording and run for one iteration. For this, 24-dimensional MFCCs are used without mean or variance normalization nor delta coefficients, extracted from 15 ms frames with 5 ms overlap. A 1024-diagonal covariance components GMM-*universal background model* (UBM) and 400 eigenvoices TV matrix, both trained with the x-vector extractor data, are used. The zeroth-order statistics are boosted before VB-HMM likelihood computation for posterior scaling to reduce frequent speaker transitions (Singh et al. 2019).

6 Domain classification

The framework of proposed ADI system is shown in Fig. 5. It is based on the speaker embeddings as recording-level feature and nearest neighbor classifier. Though the speaker embeddings are principally developed for speaker characterization, they also capture information related to acoustic scene (Zeinali et al. 2018), recording session (Raj et al. 2019), and channel (Wang et al. 2017). In this work, we studied two frequently used speaker embeddings: discriminatively trained x-vectors and generative i-vectors. In an earlier study, these speaker embeddings were investigated for the second DIHARD dataset on related tasks (Sahidullah et al. 2019; Fennir et al. 2020). Owing to the similarity of ADI with ASC, we additionally investigated L3-net embeddings, which were a part of the baseline system provided with the ASC task of the DCASE 2020 challenge (Heittola et al. 2020).

6.1 i-vector embeddings

Speaker diarization involves splitting multi-speaker spoken documents according to speaker changes and segment clustering. Since it closely resembles speaker recognition, the latter’s algorithms are applied to the former with good results. i-vectors are one such example (Dehak et al. 2010). Speaker recognition systems work in two steps, enrolment and recognition. During the speaker enrolment process, a UBM (Reynolds et al. 2000) is generated using the data collected from non-target utterances. A target model is generated by

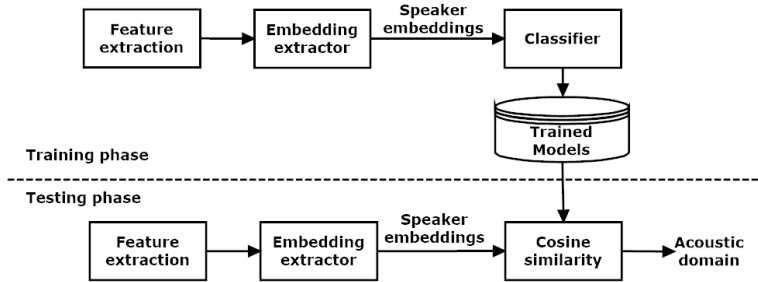


Fig. 5 Block diagram of the proposed acoustic domain identification system.

adapting the UBM according to the target data. Whenever any test utterance is given as input to the system, features are extracted from it, and pattern matching algorithms are applied using one or more kinds of target models. i-vectors is designed to improve the joint factor analysis (Kenny et al. 2007) by combining the inter and intra-domain variability and modeling it in the same low dimensional total variability space. This space \mathbf{T} accounts for both speaker and channel variability.

The speaker and channel dependent supervector \mathbf{M} is given by

$$\mathbf{M} = \mathbf{m} + \mathbf{T}\mathbf{x} \quad (1)$$

If C is the number of components in UBM and F is the dimension of acoustic feature vectors then for a given utterance, we concatenate the F -dimensional GMM mean vectors to get the speaker and channel independent UBM supervectors \mathbf{m} of dimension CF . The normally distributed random vectors \mathbf{x} are known as identity vectors or i-vectors. It is assumed that \mathbf{M} is distributed normally with $\mathbf{T}\mathbf{T}^\top$ as its covariance matrix and \mathbf{m} as its mean vector. For training of the matrix \mathbf{T} , it is assumed that the utterances of the same speaker are produced by different speakers.

6.2 x-vector embeddings

x-vectors represent a recent but popular approach from speaker recognition research (Snyder et al. 2018; Bai and Zhang 2021). Here, long-term speaker characteristics are captured in a feed-forward deep neural network by a temporal pooling layer that aggregates over the input speech (Snyder et al. 2017). The network has layers that operate on speech frames, a statistics pooling layer that aggregates over the frame-level representations, additional layers that operate at the segment-level, and a softmax output layer at the end. The nonlinearities are *rectified linear units* (ReLUs).

Suppose an input segment has K frames. The first five layers operate on speech frames, with a small temporal context centered at the current frame. The statistics pooling layer aggregates all K frame-level outputs from fifth

frame-layer and computes its mean and standard deviation. This process aggregates information across the time dimension so that subsequent layers operate on the entire segment. The mean and standard deviation are concatenated together and propagated through segment-level layers, and finally, the softmax output layer.

The goal of training the network is to produce embeddings that generalize well to speakers that have not been seen in the training data. The embeddings should capture speaker characteristics over the entire utterance, rather than at the frame-level. After training, the embeddings extracted from the affine component of first segment-layer are referred to as x-vectors. The pre-softmax affine layer is not considered because of its large size and dependence on the number of speakers.

6.3 L3-net embeddings

The look, listen and learn (L3)-network embeddings were proposed for learning *audio-visual correspondence* (AVC) in videos in a self-supervised manner (Arandjelovic and Zisserman 2017). The goal of AVC learning was to *learn* audio and visual semantic information simultaneously by simply *watching* and *listening* to many unlabelled videos. Without needing annotated data and a relatively simple convolutional architecture, L3-net successfully produced powerful embeddings that led to state-of-the-art sound classification performance. Thus, it can be applicable to a wide variety of machine learning tasks with scarce labeled data. However, in (Arandjelovic and Zisserman 2017), non-trivial parameter choices that may affect the performance and computation cost of the model were not elaborated. This issue was addressed in (Cramer et al. 2019) in the context of sound event classification. The authors worked with three publicly available environmental audio datasets and found that mel-spectrogram gave better audio representation than linear spectrogram, which was used in (Arandjelovic and Zisserman 2017). Their results also indicated that for a downstream classification task, using the audio that makes the embeddings most discriminative, independently of the downstream domain, is more important. Regarding the training data, it was suggested that at least 40M samples should be used to train the L3-Net embedding.

We first computed the average of the embeddings for each domain and stored it as trained models. We calculated the cosine similarity during testing and assigned the class label with the highest similarity.

7 Experimental setup

The ADI experiments were performed on the development set of 254 speech utterances from 11 domains. We randomly selected 200 utterances for training and used the remaining 54 for test. For similarity measurement between training and test data, *cosine similarity* was employed. The experiments were repeated 1000 times.

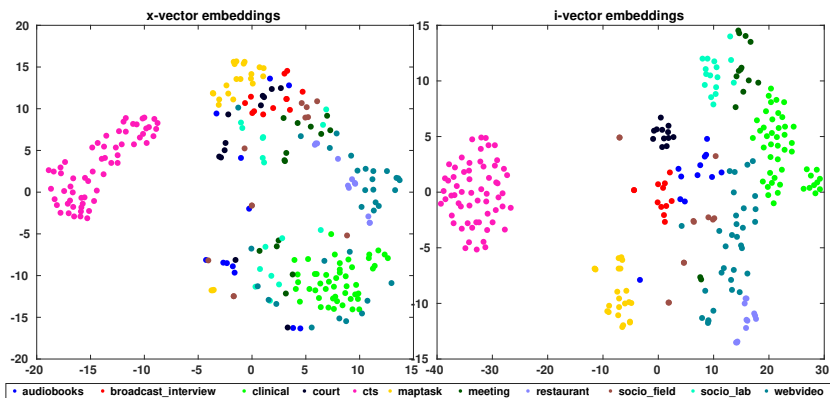


Fig. 6 Domain-wise scatter plot of x-vector and i-vector features.

To extract utterance-level embeddings for the ADI task, we used pre-trained x-vector and i-vector models trained on VoxCeleb audio-data². Pre-trained versions of the L3-Net variants studied in (Cramer et al. 2019) are made freely available online by the name OpenL3³. Open L3 embeddings were used as feature representation for the baseline system provided with the DCASE 2020’s ASC task (Heittola et al. 2020). The system had a one-second analysis window, with a 100 ms hop, 256 mel filters for input representation, content type music, and a 512-dimensional embedding. We used the same for our ADI task.

We performed domain-dependent speaker diarization on the evaluation set of DIHARD III data by first predicting the domain for every utterance using the ADI system and grouping the utterances according to the predicted domains. i-vector embeddings manifest superior domain-clustering abilities (refer Fig. 6 and Fig. 7 in Section 8.1), so we resort to them for ADI. We used domain-specific thresholds for speaker clustering in the diarization pipeline. Towards this, we computed the domain-specific speaker diarization performance for different thresholds of the development set. We selected the threshold value that shows the best SD performance for a specific subset.

For diarization, our experimental setup is based on the baseline system created by the organizers (Ryant et al. 2021). We have used the toolkit⁴ with the same frame-level acoustic features, embedding extractor, scoring method, etc. The system is trained with all 254 sentences for identifying the domains of evaluation data samples.

The primary evaluation metric of the challenge is *diarization error rate* (DER), which is the sum of all the errors associated with the diarization task. *Jaccard error rate* (JER) is the secondary evaluation metric, and it is based

² <https://kaldi-asr.org/models/m7>

³ <https://github.com/marl/openl3>

⁴ https://github.com/dihardchallenge/dihard3_baseline




Fig. 7 Acoustic domain identification performance in terms of accuracy, F1-score, unweighted average recall (UAR), Matthews correlation coefficient (MCC) using i-vector, x-vector, and OpenL3 embeddings.

on the Jaccard similarity index. The details of the challenge with rules are available in the challenge evaluation plan (Ryant et al. 2020).

8 Results

8.1 ADI experiments

First, we present in Fig. 6 the domain-wise clustering ability of the two speaker embeddings: i-vectors and x-vectors. The scatter-plots show that ‘CTS’ samples are well-separated from the rest of the data because this subset consists of narrowband telephone speech, whereas the others are from wideband speech. We also observe that i-vector embeddings have done better for the other domains than the x-vector embeddings. However, for ‘Sociolinguistic (Field)’, the clusters are not well represented by both the embeddings, likely due to the within-class variability in the subset (refer Table 1, Row 9).

We carried out ADI with the three embeddings described in Section 6. Figure 7 shows the corresponding performance comparison. Besides the commonly used classification accuracy, we also reported *unweighted average recall* (UAR), *F1-score*, and *Mathews correlation coefficient* (MCC), because of the unbalanced nature of the data. The figure shows that the i-vector and OpenL3 systems are substantially better than the x-vector system for ADI on the DI-HARD III dataset. With the present experimental setup, i-vectors surpass the other two embeddings on all four metrics. For instance, the average domain classification accuracy over 1000 repetitions was 91.11%, 72.98%, and 89.64% for i-vector, x-vector, and OpenL3 systems, respectively. The poor performance of x-vectors could be because they are originally designed to capture better

Table 2 Performance of ADI system on DIHARD III evaluation for i-vector, x-vector, and OpenL3 based systems.

	Accuracy (%)	F1-Score (%)	UAR (%)	MCC (%)
i-vector	86.49	80.26	82.76	80.24
x-vector	69.88	56.61	60.14	55.85
OpenL3	75.29	72.84	73.43	71.66

speaker characteristics over the entire utterance. However, here we are trying to classify domains where one utterance may have multiple speakers, and the same domain data may not have common speakers.

The ADI results over evaluation data of the DIHARD III challenge for the three embeddings are displayed in Table 2. Here, we notice that the four metrics follow the same pattern as in Fig. 7, but the difference in the performance of the embeddings is more prominently visible. Based on these values, we chose i-vectors for ADI and proceeded to the speaker diarization stage.

8.2 Speaker diarization experiments: Development set

In our first SD experiment with the baseline system, we evaluated the development set of the DIHARD III dataset and analyzed the diarization performance for each domain separately. The top row for each domain listed in Table 3 shows the domain-wise baseline results in terms of DER and JER. As expected, we notice that the SD performance varied with domains. The DER values range from less than 4% for ‘Broadcast Interview’ domain to greater than 46% corresponding to ‘Restaurant’ domain.

Next, we performed the diarization experiments with domain-dependent processing. We selected a domain-dependent threshold for speaker clustering. In addition, for PLDA adaptation during the scoring process, we considered audio data from specific domains only. This system is referred to as **P1** in Table 3. To our surprise, this degraded the performance compared to the baseline for most of the domains. We hypothesize that the limited speaker variability in domain-wise audio data (refer to Table 1) could be a factor as we used this data as in-domain target data for PLDA adaptation. The PLDA adaption with domain-specific data was helpful for ‘Audiobooks’ and ‘CTS’ because the two domains are substantially different from the others. For example, ‘Audiobooks’ consists of high-quality recordings with a single speaker, and ‘CTS’ consists of narrow-band telephone speech with two speakers.

In the next set of SD experiments, we performed PLDA adaptation with audio data from all the eleven subsets while applying domain-specific thresholds for speaker clustering. We refer to this method as **P2** in Table 3. The noticeable improvement in ‘Courtroom’ audio diarization could be because of the following reason. As pointed out earlier in Section 4, despite having multiple speakers in argument, ‘Courtroom’ has less overlapped speech and a low

Table 3 The impact of domain-dependent processing on speaker diarization performance (DER in %/JER in %) for different acoustic domains of the development set of the third DIHARD challenge. (P1: Domain-dependent threshold and PLDA adaptation, P2: Domain-dependent threshold but PLDA adaptation with full data.)

Domain	Method	First Pass		Re-segmentation with VB-HMM	
		Full	Core	Full	Core
Audiobooks	Baseline	4.95 / 4.81	4.95 / 4.81	4.55 / 4.45	4.55 / 4.45
	P1	0.00 / 0.00	0.00 / 0.00	0.40 / 0.40	0.40 / 0.40
	P2	0.00 / 0.00	0.00 / 0.00	0.40 / 0.40	0.40 / 0.40
Broadcast interview	Baseline	3.75 / 25.62	3.75 / 25.62	3.56 / 23.83	3.56 / 23.83
	P1	6.47 / 39.38	6.47 / 39.38	6.10 / 38.98	6.10 / 38.98
	P2	3.51 / 24.05	3.51 / 24.05	3.29 / 22.27	3.29 / 22.27
Clinical	Baseline	17.55 / 28.88	16.08 / 26.38	17.69 / 28.46	16.72 / 27.21
	P1	20.06 / 29.92	18.88 / 28.11	16.71 / 25.50	15.21 / 23.66
	P2	15.81 / 23.69	14.61 / 22.37	14.69 / 22.07	13.78 / 21.59
Courtroom	Baseline	10.81 / 38.75	10.81 / 38.75	10.17 / 37.63	10.17 / 37.63
	P1	12.19 / 43.99	12.19 / 43.99	9.03 / 37.97	9.03 / 37.97
	P2	5.82 / 23.91	5.82 / 23.91	4.77 / 22.22	4.77 / 22.22
CTS	Baseline	22.25 / 27.85	23.48 / 28.41	18.38 / 22.32	18.84 / 22.37
	P1	20.22 / 27.73	20.28 / 27.18	17.31 / 22.31	16.89 / 21.07
	P2	19.23 / 25.93	19.81 / 25.96	16.55 / 20.94	17.10 / 21.20
Map Task	Baseline	9.92 / 18.90	9.92 / 18.90	8.75 / 16.66	8.75 / 16.66
	P1	12.19 / 25.75	12.19 / 25.75	9.23 / 20.25	9.23 / 20.25
	P2	9.21 / 18.20	9.21 / 18.20	8.50 / 16.50	8.50 / 16.50
Meeting	Baseline	29.65 / 51.57	29.65 / 51.57	28.80 / 51.01	28.80 / 51.01
	P1	31.31 / 52.71	31.31 / 52.71	29.35 / 50.85	29.35 / 50.85
	P2	29.80 / 50.72	29.80 / 50.72	28.62 / 50.09	28.62 / 50.09
Restaurant	Baseline	46.99 / 73.80	46.99 / 73.80	46.60 / 74.83	46.60 / 74.83
	P1	48.68 / 74.11	48.68 / 74.11	46.40 / 73.83	46.40 / 73.83
	P2	45.34 / 70.49	45.34 / 70.49	45.34 / 72.98	45.34 / 72.98
Sociolinguistic (Field)	Baseline	17.67 / 49.48	17.67 / 49.48	16.24 / 47.68	16.24 / 47.68
	P1	23.06 / 60.61	23.06 / 60.61	21.53 / 58.31	21.53 / 58.31
	P2	17.46 / 51.19	17.49 / 51.19	16.16 / 49.16	16.16 / 49.16
Sociolinguistic (Lab)	Baseline	11.97 / 17.65	11.97 / 17.65	10.33 / 14.94	10.33 / 14.94
	P1	13.60 / 23.84	13.60 / 23.84	11.41 / 18.70	11.41 / 18.70
	P2	10.38 / 15.92	10.38 / 15.92	9.21 / 13.87	9.21 / 13.87
Web Video	Baseline	40.99 / 76.97	40.99 / 76.97	40.06 / 76.77	40.06 / 76.77
	P1	41.60 / 79.93	41.60 / 79.93	40.97 / 79.66	40.97 / 79.66
	P2	40.14 / 79.29	40.14 / 79.29	40.25 / 79.38	40.25 / 79.38

value of the speech-to-nonspeech ratio. Another interesting characteristic of this domain’s data is its high spectral energy, even at high frequencies, along with only a slight variation in the SNR. The minor improvement in the DER and worsened JER for ‘Sociolinguistic (Field)’ reiterates the poor clustering of this domain’s samples (refer Fig. 6). A similar argument applies to the trends of results in ‘Meeting’ and ‘Web videos.’

After the first pass for the ‘Full’ data condition, we observe that ‘Audiobooks,’ ‘Courtroom,’ and ‘CTS’ exhibit a major relative improvement of 100%, 46.16%, and 13.57% in DER, respectively, to the baseline system. However, following re-segmentation (second pass), the highest improvement in DER is ex-

Table 4 Results showing the speaker diarization performance using baseline and proposed method on development and evaluation set of third DIHARD challenge. For this phase, we consider the second pass with re-segmentation as it consistently gives superior SD performance over the first pass.

Set	Method	Full		Core	
		DER	JER	DER	JER
Development	Baseline	19.59	43.01	20.17	47.28
	Proposed	17.97	40.33	18.73	44.77
Evaluation	Baseline	19.19	43.28	20.39	48.61
	Proposed	17.56	38.60	19.23	43.74

hibited by ‘Audiobooks’: 91.21%, ‘Clinical’: 16.96%, and ‘Courtroom’: 53.09%, respectively. These domains have less speaker overlap and lower speech-to-nonspeech ratios except the ‘CTS’ domain. On the other hand, a poor relative improvement of -0.50%, 1.19%, and 2.07% in DER is registered for ‘Meeting,’ ‘Sociolinguistic (Field),’ and ‘Web video’ respectively after the first pass of diarization. The same domains show low relative DER improvements of 0.63%, 0.49%, and -0.47%, respectively, after the second pass of diarization. This can be explained by the fact that these domains have considerable overlapping speech and poor SNR. For most domains, the relative improvement is higher after the first pass than after the second because the clustering was done in the first pass at an ideal threshold, leaving little room for improvement during re-segmentation.

Our detailed analysis of the development set results indicates that the **P2** method shows reduced DER/JER values for eight out of eleven domains. This confirms our hypothesis that domain-dependent threshold for speaker clustering helps. We subsequently used **P2** method for diarization of the evaluation set.

8.3 Speaker diarization experiments: Evaluation set

We performed experiments on the evaluation set by predicting the domain for each of its recordings which is followed by selecting the corresponding clustering threshold. For domain prediction, we used the i-vector-based embeddings as features and the nearest neighbor classifier as this gave promising results for ADI (See details in Section 8.1).

Table 4 summarizes the results for the entire development and evaluation sets. We observe substantial overall improvements in development and evaluation data. For the ‘Full’ condition of the evaluation set consisting of all the audio files, we obtained a relative improvement of 8.49% and 10.81% in terms of DER and JER, respectively. Whereas for the ‘Core’ condition with a balanced amount of audio data per domain, we achieved a relative improvement of 5.69% and 10.02%. This reflects the results of the detailed analysis of individual domains in Section 8.2 where we established that the relative improvements with the proposed method varied for different domains.

9 Conclusion

This work advances the state-of-the-art DNN-based speaker diarization performance with the help of domain-specific processing. We applied speech embeddings computed from the entire audio recording for the supervised audio domain identification task. After experimenting with three different embeddings, we employed i-vector embeddings and the nearest neighbor classifier for the ADI task. We integrated the ADI system with the x-vector based state-of-the-art diarization system and performed the experiments on the DIHARD III challenge dataset consisting of multiple domains. We obtained up to 8% relative improvement on the ‘Full’ condition of the evaluation set. The proposed method substantially improved the diarization performance on most of the subsets of the DIHARD III dataset.

Our work involved the application of domain-specific thresholds for speaker clustering. This study can be extended by incorporating other domain-specific operations such as speech enhancement and feature engineering. We used a basic *time delay neural network* (TDNN) based x-vector embedding for audio segment representation. The future work may also include investigation of more advanced embeddings (e.g., *emphasized channel attention, propagation, and aggregation in TDNN* or ECAPA-TDNN), which have shown promising results in related tasks.

The present study is limited to a supervised domain classification problem where the training phase requires class-annotated data. However, in practice, the audio data may come from myriad unknown sources, and our domain classification method may not be applicable in its current form. Nevertheless, an investigation of unsupervised domain clustering to exploit the advantage of domain-dependent processing in such a scenario can be carried out.

Acknowledgements

We thank the anonymous reviewers for their careful reading of our manuscript and their insightful comments and suggestions. This work is funded by the Jagadish Chandra Bose Centre of Advanced Technology (JCBCAT), Govt. of India. Experiments presented in this paper were partially carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

References

- X. Anguera, S. Bozonnet, N. Evans, C. Fredouille, G. Friedland, O. Vinyals, Speaker diarization: A review of recent research. *IEEE Transactions on Audio, Speech, and Language Processing* **20**(2), 356–370 (2012). IEEE
- R. Arandjelovic, A. Zisserman, Look, listen and learn, in *Proc. of the IEEE International Conference on Computer Vision*, 2017, pp. 609–617

- P. Assmann, Q. Summerfield, The perception of speech under adverse conditions. *Speech Processing in the Auditory System* **18**, 231–308 (2004)
- Z. Bai, X.-L. Zhang, Speaker recognition based on deep learning: An overview. *Neural Networks* **140**, 65–99 (2021)
- D. Barchiesi, D. Giannoulis, D. Stowell, M.D. Plumbley, Acoustic scene classification: Classifying environments from the sounds they produce. *IEEE Signal Processing Magazine* **32**(3), 16–34 (2015). IEEE
- V. Bisot, R. Serizel, S. Essid, G. Richard, Feature learning with matrix factorization applied to acoustic scene classification. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **25**(6), 1216–1229 (2017)
- H. Chen, Z. Liu, Z. Liu, P. Zhang, Y. Yan, Integrating the data augmentation scheme with various classifiers for acoustic scene modeling, Technical report, DCASE2019 Challenge, 2019
- J.S. Chung, A. Nagrani, A. Zisserman, VoxCeleb2: Deep speaker recognition. *Proc. INTERSPEECH*, 1086–1090 (2018). ISCA
- J. Cramer, H.-H. Wu, J. Salamon, J.P. Bello, Look, listen, and learn more: Design choices for deep audio embeddings, in *Proc. ICASSP*, IEEE, 2019, pp. 3852–3856. IEEE
- N. Dehak, P.J. Kenny, R. Dehak, P. Dumouchel, P. Ouellet, Front-end factor analysis for speaker verification. *IEEE Transactions on Audio, Speech, and Language Processing* **19**(4), 788–798 (2010)
- M. Diez, L. Burget, F. Landini, J. Černocký, Analysis of speaker diarization based on bayesian HMM with eigenvoice priors. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* **28**, 355–368 (2019)
- M. Dorfer, B. Lehner, H. Eghbal-zadeh, H. Christop, P. Fabian, W. Gerhard, Acoustic scene classification with fully convolutional neural networks and i-vectors, Technical report, DCASE2018 Challenge, 2018
- H. Eghbal-Zadeh, B. Lehner, M. Dorfer, G. Widmer, CP-JKU submissions for DCASE-2016: a hybrid approach using binaural i-vectors and deep convolutional neural networks, Technical report, DCASE2016 Challenge, 2016
- T. Fennir, F. Habib, C. Macaire, Acoustic scene classification for speaker diarization, Technical report, Université de Lorraine, 2020
- D. Garcia-Romero, D. Snyder, G. Sell, D. Povey, A. McCree, Speaker diarization using deep neural network embeddings, in *Proc. ICASSP*, IEEE, 2017, pp. 4930–4934. IEEE
- D. Giannoulis, D. Stowell, E. Benetos, M. Rossignol, M. Lagrange, M.D. Plumbley, A database and challenge for acoustic scene classification and event detection, in *Proc. EUSIPCO*, IEEE, 2013, pp. 1–5. IEEE
- K.J. Han, S. Kim, S.S. Narayanan, Strategies to improve the robustness of agglomerative hierarchical clustering under data source variation for speaker diarization. *IEEE Transactions on Audio, Speech, and Language Processing* **16**(8), 1590–1601 (2008)
- T. Heittola, A. Mesaros, T. Virtanen, Acoustic Scene Classification in DCASE 2020 Challenge: Generalization Across Devices and Low Complexity Solutions, in *Proc. of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop*, Tokyo, Japan, 2020, pp. 56–60
- M. Huijbregts, C. Wooters, The blame game: performance analysis of speaker diarization system components, in *Proc. INTERSPEECH*, ISCA, 2007, pp. 1857–1860. ISCA
- P. Kenny, G. Boulianne, P. Ouellet, P. Dumouchel, Joint factor analysis versus eigenchannels in speaker recognition. *IEEE Transactions on Audio, Speech, and Language Processing* **15**(4), 1435–1447 (2007)
- C. Kim, R.M. Stern, Robust signal-to-noise ratio estimation based on waveform amplitude distribution analysis, in *Proc. INTERSPEECH*, ISCA, 2008, pp. 2598–2601. ISCA
- T. Ko, V. Peddinti, D. Povey, M.L. Seltzer, S. Khudanpur, A study on data augmentation of reverberant speech for robust speech recognition, in *Proc. ICASSP*, IEEE, 2017, pp. 5220–5224. IEEE
- F. Landini, S. Wang, M. Diez, L. Burget, P. Matějka, K. Žmolíková, L. Mošner, A. Silnova, O. Plchot, O. Novotný, H. Zeinali, J. Rohdin, BUT System for the Second Diarization Challenge, in *Proc. ICASSP 2020*, IEEE, 2020, pp. 6529–6533. IEEE
- A. Löfqvist, The long-time-average spectrum as a tool in voice research. *Journal of Phonetics* **14**(3-4), 471–475 (1986)

- A. Mesaros, T. Heittola, T. Virtanen, TUT database for acoustic scene classification and sound event detection, in *Proc. EUSIPCO*, IEEE, 2016, pp. 1128–1132. IEEE
- A. Mesaros, T. Heittola, T. Virtanen, Acoustic scene classification: an overview of DCASE 2017 challenge entries, in *Proc. 2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC)*, IEEE, 2018a, pp. 411–415. IEEE
- A. Mesaros, T. Heittola, T. Virtanen, A multi-device dataset for urban acoustic scene classification, in *Proc. of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop*, 2018b, pp. 9–13
- A. Mesaros, T. Heittola, T. Virtanen, Acoustic scene classification in DCASE 2019 challenge: Closed and open set classification and data mismatch setups, in *Proc. of the Detection and Classification of Acoustic Scenes and Events 2019 Workshop*, Oct. 2019, pp. 164–168
- N. Mirghafari, C. Wooters, Nuts and flakes: A study of data characteristics in speaker diarization, in *Proc. ICASSP*, vol. 1, IEEE, 2006, p. . IEEE
- S. Mun, S. Park, D. Han, H. Ko, Generative adversarial network based acoustic scene training set augmentation and selection using SVM hyper-plane, Technical report, DCASE2017 Challenge, 2017
- A. Nagrani, J.S. Chung, A. Zisserman, VoxCeleb: A large-scale speaker identification dataset, in *Proc. INTERSPEECH*, ISCA, 2017, pp. 2616–2620. ISCA
- T.J. Park, N. Kanda, D. Dimitriadis, K.J. Han, S. Watanabe, S. Narayanan, A review of speaker diarization: Recent advances with deep learning. *Computer Speech Language* **72**, 101317 (2022)
- S.J. Prince, J.H. Elder, Probabilistic linear discriminant analysis for inferences about identity, in *Proc. of the IEEE International Conference on Computer Vision*, 2007, pp. 1–8
- D. Raj, D. Snyder, D. Povey, S. Khudanpur, Probing the information encoded in x-vectors, in *Proc. IEEE ASRU*, 726-733, 2019. 726-733
- A. Rakotomamonjy, G. Gasso, Histogram of gradients of time-frequency representations for audio scene classification. *IEEE/ACM Transactions on Audio, Speech and Language Processing* **23**(1), 142–153 (2015)
- Z. Ren, K. Qian, Z. Zhang, V. Pandit, A. Baird, B. Schuller, Deep scalogram representations for acoustic scene classification. *IEEE/CAA Journal of Automatica Sinica* **5**(3), 662–669 (2018)
- D.A. Reynolds, T.F. Quatieri, R.B. Dunn, Speaker verification using adapted Gaussian mixture models. *Digital Signal Processing* **10**(1-3), 19–41 (2000)
- G. Roma, W. Nogueira, P. Herrera, Recurrence quantification analysis features for auditory scene classification, Technical report, DCASE2013 Challenge, 2013
- N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, M. Liberman, First DIHARD challenge evaluation plan, Technical report, DIHARD challenge 2018, 2018
- N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, M. Liberman, The Second DIHARD Diarization Challenge: Dataset, Task, and Baselines, in *Proc. INTERSPEECH 2019*, ISCA, 2019, pp. 978–982. ISCA
- N. Ryant, K. Church, C. Cieri, J. Du, S. Ganapathy, M. Liberman, Third DIHARD challenge evaluation plan, Technical report, Linguistic Data Consortium, University of Pennsylvania, Philadelphia, PA, USA, 2020
- N. Ryant, P. Singh, V. Krishnamohan, R. Varma, K. Church, C. Cieri, J. Du, S. Ganapathy, M. Liberman, The third DIHARD diarization challenge, in *Proc. INTERSPEECH*, ISCA, 2021, pp. 3570–3574. ISCA
- M. Sahidullah, J. Patino, S. Cornell, R. Yin, S. Sivasankaran, H. Bredin, P. Korshunov, A. Brutti, R. Serizel, E. Vincent, N. Evans, S. Marcel, S. Squartini, C. Barras, The Speed submission to DIHARD II: Contributions & lessons learned. arXiv preprint arXiv:1911.02388 (2019)
- Y. Sakashita, M. Aono, Acoustic scene classification by ensemble of spectrograms based on adaptive temporal divisions, Technical report, DCASE2018 Challenge, 2018
- G. Sell, D. Garcia-Romero, Speaker diarization with PLDA i-vector scoring and unsupervised calibration, in *Proc. 2014 IEEE Spoken Language Technology Workshop (SLT)*, IEEE, 2014, pp. 413–417. IEEE
- G. Sell, D. Snyder, A. McCree, D. Garcia-Romero, J. Villalba, M. Maciejewski, V. Manohar,

- N. Dehak, D. Povey, S. Watanabe, S. Khudanpur, Diarization is hard: some experiences and lessons learned for the JHU team in the inaugural DIHARD challenge., in *Proc. INTERSPEECH*, ISCA, 2018, pp. 2808–2812. ISCA
- S.H. Shum, N. Dehak, R. Dehak, J.R. Glass, Unsupervised methods for speaker diarization: An integrated and iterative approach. *IEEE Transactions on Audio, Speech, and Language Processing* **21**(10), 2015–2028 (2013)
- S. Shum, N. Dehak, E. Chuangsuwanich, D. Reynolds, J. Glass, Exploiting intra-conversation variability for speaker diarization, in *Proc. INTERSPEECH*, ISCA, 2011, pp. 945–948. ISCA
- M. Sinclair, S. King, Where are the challenges in speaker diarization?, in *Proc. ICASSP*, IEEE, 2013, pp. 7741–7745. IEEE
- P. Singh, H. Vardhan, S. Ganapathy, A. Kanagasundaram, LEAP diarization system for the second DIHARD challenge, in *Proc. INTERSPEECH*, ISCA, 2019, pp. 983–987. ISCA
- D. Snyder, G. Chen, D. Povey, MUSAN: A music, speech, and noise corpus. arXiv preprint arXiv:1510.08484 (2015)
- D. Snyder, D. Garcia-Romero, D. Povey, S. Khudanpur, Deep neural network embeddings for text-independent speaker verification, in *Proc. INTERSPEECH*, ISCA, 2017, pp. 999–1003. ISCA
- D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, S. Khudanpur, X-vectors: Robust DNN embeddings for speaker recognition, in *Proc. ICASSP*, IEEE, 2018, pp. 5329–5333. IEEE
- S. Suh, S. Park, Y. Jeong, T. Lee, Designing acoustic scene classification models with CNN variants, Technical report, DCASE2020 Challenge, 2020
- S. Waldekar, G. Saha, Classification of audio scenes with novel features in a fused system framework. *Digital Signal Processing* **75**, 71–82 (2018)
- S. Waldekar, G. Saha, Analysis and classification of acoustic scenes with wavelet transform-based mel-scaled features. *Multimedia Tools and Applications* **79**(11), 7911–7926 (2020)
- Q. Wang, C. Downey, L. Wan, P.A. Mansfield, I.L. Moreno, Speaker diarization with LSTM, in *Proc. ICASSP*, IEEE, 2018, pp. 5239–5243. IEEE
- S. Wang, Y. Qian, K. Yu, What does the speaker embedding encode?, in *Proc. INTERSPEECH*, ISCA, 2017, pp. 1497–1501. ISCA
- Y. Wang, M. He, S. Niu, L. Sun, T. Gao, X. Fang, J. Pan, J. Du, C.-H. Lee, USTC-NELSLIP system description for DIHARD-III challenge. arXiv preprint arXiv:2103.10661 (2021)
- H. Zeinali, L. Burget, J.H. Cernocky, Convolutional neural networks and x-vector embedding for DCASE2018 Acoustic Scene Classification challenge, in *Proc. of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop*, 2018, pp. 202–206
- A. Zhang, Q. Wang, Z. Zhu, J. Paisley, C. Wang, Fully supervised speaker diarization, in *Proc. ICASSP*, IEEE, 2019, pp. 6301–6305. IEEE
- W. Zhu, J. Pelecanos, Online speaker diarization using adapted i-vector transforms, in *Proc. ICASSP*, IEEE, 2016, pp. 5045–5049. IEEE