



The Third DIHARD Diarization Challenge

Neville Ryant¹, Prachi Singh², Venkat Krishnamohan², Rajat Varma², Kenneth Church³,
Christopher Cieri¹, Jun Du⁴, Sriram Ganapathy², Mark Liberman¹

¹Linguistic Data Consortium, University of Pennsylvania, Philadelphia, PA, USA

²LEAP Lab, Electrical Engineering, Indian Institute of Science, Bangalore, India

³Baidu Research, Sunnyvale, CA, USA

⁴University of Science and Technology of China, Hefei, China

nryant@ldc.upenn.edu

Abstract

DIHARD III was the third in a series of speaker diarization challenges intended to improve the robustness of diarization systems to variability in recording equipment, noise conditions, and conversational domain. Speaker diarization was evaluated under two speech activity conditions (diarization from a reference speech activity vs. diarization from scratch) and 11 diverse domains. The domains span a range of recording conditions and interaction types, including read audio-books, meeting speech, clinical interviews, web videos, and, for the first time, conversational telephone speech. A total of 30 organizations (forming 21 teams) from industry and academia submitted 499 valid system outputs. The evaluation results indicate that speaker diarization has improved markedly since DIHARD I, particularly for two-party interactions, but that for many domains (e.g., web video) the problem remains far from solved.

Index Terms: speaker diarization, speaker recognition, robust ASR, noise, conversational speech, DIHARD challenge

1. Introduction

Speaker diarization, often referred to as “who spoke when”, is the task of determining the number of speakers present in a conversation and their respective regions of activity in the given recording. In addition to being an interesting technical challenge, it forms an important part of the pre-processing pipeline for speech-to-text [1] and is essential for making objective measurements of turn-taking behavior. Early work in this area was driven by the NIST Rich Transcription (RT) evaluations [2], which ran from 2002 to 2009. In addition to driving substantial performance improvements, especially for meeting speech, the RT evaluations introduced diarization error rate (DER), which remains the principal evaluation metric in this area.

After the RT evaluation series ended in 2009, diarization systems continued to improve (e.g., i-vectors, x-vectors, PLDA scoring), though until quite recently there was no common benchmark for diarization, resulting in a fragmented research landscape where individual groups focused on different datasets or domains (e.g., conversational telephone speech [3, 4, 5, 6, 7], broadcast news [8, 9], or meeting [10, 11]), often with slightly differing evaluation methodologies. At best, this has made performance comparisons difficult, while at worst it may have engendered over-fitting to individual domains/datasets, resulting in systems that do not generalize.

Recently, there has been renewed interest in a common diarization task to facilitate systematic benchmarking. Whereas from 2009-2017 there were no major evaluations with a diarization component, there now is an annual diarization specific

evaluation – DIHARD – as well as numerous other challenges that include a diarization component; among others, the Fearless Steps series [12, 13], the Iberspeech-RTVE challenge [14], CHiME-6 [1], and VoxSRC-20 [15].

The first DIHARD challenge (DIHARD I) [16] ran in the spring of 2018 and evaluated diarization of single channel wide-band recordings drawn from a diverse range of domains. As expected, state-of-the-art systems performed poorly, with error rates of the best systems [17, 18] more than double the state-of-the-art for CALLHOME [19] at the time [5, 6]. This was followed by DIHARD II [20, 21] in 2019, which was even more successful, attracting 50 teams from 17 countries and 4 continents. While DIHARD II continued the single channel diarization tracks from DIHARD I, it also collaborated with the CHiME challenge series with the addition of two new tracks focusing on conversational speech from multiple far-field microphone arrays during a dinner party scenario. All tracks continued to be challenging for participants, with the most challenge tracks being the ones without reference speech segmentation and dinner party data. In the case of the latter, the CHiME-6 data, DER of the best performing system was over 45% when provided with an oracle speech segmentation and over 58% when required to produce its own segmentation.

The current challenge (DIHARD III)¹ [22], which builds upon DIHARD I and II, addresses the problem of robust diarization which is resilient to variation in, among others, conversational domain, recording equipment, recording environment, reverberation, ambient noise, number of speakers, and speaker demographics. Like its predecessors, performance is evaluated under two SAD conditions: diarization from oracle reference SAD and diarization from scratch. There are no constraints on training data, with participants allowed to use any combination of open source noise sources, number of speakers, speaker demographics, and proprietary data for system development. Recordings are sampled from 11 diverse domains ranging from clean, near-field recordings of read audio-books to extremely noisy, highly interactive, far-field recordings of speech in restaurants to clinical interviews with children. Unlike DIHARD II, diarization from multi-channel audio is not evaluated; parties interested in this condition should instead consult the results from track 2 of CHiME-6 [1], which is essentially a rerun of the DIHARD II multichannel condition. A total of 30 organizations (forming 21 teams) from industry and academia submitted 499 valid system outputs (352 to track 1 and 147 to track 2).

In the remainder of this paper, we introduce the task, metrics, and data, as well as the baseline SAD and diarization sys-

¹<https://dihardchallenge.github.io/dihard3/>

Table 1: Overview of DIHARD III datasets. The **Part.** column indicates the partition (core or full), while the **% speech** and **% overlap** columns indicate, respectively, the percentage of speech/overlapped speech in the partition.

Set	Part.	# rec	# hours	% speech	% overlap
Dev	Core	181	23.94	78.43	10.04
	Full	254	34.15	79.81	10.70
Eval	Core	184	22.73	77.35	8.75
	Full	259	33.01	79.11	9.35

tems. Results of the evaluation for both tracks are reported in Section 5.2.

2. Task

The goal of the challenge is to automatically detect and label all speaker segments for each recording; that is: i) determine how many speakers are present; ii) for each speaker identify all corresponding speech segments. Because system performance is strongly influenced by the quality of the speech segmentation used, two different tracks are covered:

- **Track 1** – Diarization from reference SAD. Systems are provided with a reference speech segmentation that is generated by merging speaker turns in the reference diarization.
- **Track 2** – Diarization from scratch. Systems are provided with just the raw audio input for each recording session and are responsible for producing their own speech segmentation.

3. Performance Metrics

As in DIHARD I and II, the primary metric is DER [2], which is the sum of missed speech, false alarm speech, and speaker misclassification error rates. Because systems are provided with the reference speech segmentation for track 1, for this track DER exclusively measures speaker misclassification error. This is the metric used to rank systems on the leaderboards. For each system, we also compute a secondary metric, Jaccard error rate (JER), originally introduced for DIHARD II [22]. JER is based on the Jaccard similarity index [23, 24], a metric commonly used to evaluate the output of image segmentation systems, which is defined as the ratio between the sizes of the intersections and unions of two sets of segments. An optimal mapping between speakers in the reference diarization and speakers in the system diarization is determined and for each pair the Jaccard index of their segmentations is computed. JER is defined as 1 minus the average of these scores, expressed as a percentage.

All metrics are computed using version 1.0.1 of the *dscore* tool² without the use of forgiveness collars and with scoring of overlapped speech. For further details, please consult Section 4 of the DIHARD III evaluation plan [22] and the *dscore* repo.

4. Datasets

4.1. Overview

The development (DEV) and evaluation (EVAL) sets consist of selections of 5-10 minute duration samples drawn from 11 domains exhibiting wide variation in recording equipment,

²<https://github.com/nryant/dscore>

recording environment, ambient noise, number of speakers, and speaker demographics. These domains range in difficulty from the trivial, read audio-books recorded under clean conditions by a single speaker, to the extremely challenging, conversations between up to 6 diners recorded by a binaural microphone in restaurants with varying room acoustics and noise levels. Both adult and child speech (e.g., clinical interviews) are represented, as is speech from multiple languages (English and Chinese). For the first time, narrow-band recordings are included as well as wide-band recordings; in the narrow-band case, all recordings are drawn from the unreleased Phase II calls from the Fisher English collection conducted as part of the DARPA EARS project. All the audio recordings are distributed via LDC as 16 kHz, mono-channel files.

The datasets are summarized in Table 1. For additional details about the domains, the reader is encouraged to consult the DIHARD III evaluation plan [22].

4.2. Scoring partitions

For DIHARD III, we define two partitions of the evaluation data:

- **core evaluation set** – a “balanced” evaluation set in which the total duration of each domain is approximately equal
- **full evaluation set** – a larger evaluation set that uses all available selections for each domain; it is a proper superset of the core evaluation set

The core evaluation set strives for balance across domains so that the evaluation metrics are not dominated by any single domain. It mimics the evaluation set composition from DIHARD I and II. The full evaluation set includes additional material from two domains (clinical interview and CTS), potentially resulting in more stable metrics at the expense of being unbalanced. All system submissions to all tracks are scored against both sets and the results reported on the leaderboards.

4.3. Annotation

Reference diarization was produced by segmenting the recordings into labeled speaker turns according to the following guidelines:

- split on pauses > 200 ms, where a pause by speaker “S” is defined as any segment of time during which “S” is not producing a vocalization of any kind³
- attempt to place boundaries within 10 ms of the true boundary, taking care not to truncate edges of words (e.g., utterance-final fricatives or utterance initial stops)
- where close-talking microphones exist for each speaker, perform the segmentation separately for each speaker using their individual microphone.

Reference SAD was then derived from these segmentations by merging overlapping speech segments and removing speaker identification.

During DIHARD II, it was found that manual annotation for these specifications needed highly skilled and experienced annotators using multiple spectrogram displays, making the annotation extremely slow and expensive. Many annotators were incapable of performing the task even after extensive training and the remainder found it extremely laborious with real time

³Vocalization is defined as any noise produced by the speaker by means of the vocal apparatus; e.g., speech (including yelled and whispered speech), backchannels, filled pauses, singing, speech errors and disfluencies, laughter, coughs, breaths, lipsmacks, and humming.

rates typically greater than 15X and sometimes exceeding 30X. Consequently, for DIHARD III we abandoned a commitment to entirely manual segmentation. Where a manual segmentation to these specs already existed (i.e., files annotated for DIHARD II), we used it. For all other data we instead produced a careful turn-level transcription, then established boundaries using a Kaldi-based forced aligner.

5. Baseline system

5.1. Speech activity detection

The baseline for track 2 uses a TDNN SAD model based on the Kaldi Aspire recipe (“`egs/aspire/s5`”). 40-D mel frequency cepstral coefficients (MFCCs) extracted every 30 ms using a 25 ms window are fed into a neural network consisting of 5 TDNN layers [25] followed by 2 statistics pooling layers [26]. The network context is set to approximately 1 second (left context: 0.8 sec; right context: 0.2 sec). The DNN was trained with two classes – speech and non-speech – on the DIHARD III DEV set. Training utilized the entire DEV set and was continued for 40 epochs. During inference, the posteriors of the model were converted to pseudo-likelihoods using the empirical speech/non-speech priors for the DEV set and Viterbi decoding was performed using an HMM with the following constraints: minimum speech duration: 240 ms, minimum non-speech duration: 30 ms. Miss rate, false alarm rate, and overall error (i.e., the actual frame-wise error rate) for the SAD system on the DEV and EVAL sets are depicted in Table 2.

5.2. Diarization

The diarization baseline is based on LEAP Lab’s submission to DIHARD II [27]. The system performs diarization by dividing each recording into short overlapping segments, extracting x-vectors [28, 29], scoring with probabilistic linear discriminant analysis (PLDA) [30], and clustering using agglomerative hierarchical clustering (AHC) [31]. The AHC output is then refined using variational Bayes hidden Markov model (VB-HMM) [32, 33] with posterior scaling [27]. The trained models and recipes for both tracks are distributed through GitHub⁴.

The x-vector extractor configuration is identical to that of [17, 29] with two exceptions: i) 30-D MFCCs are used instead of a mel filterbank; ii) the embedding layer uses 512 dimensions. MFCCs are extracted every 10 ms using a 25 ms window and mean-normalized using a 3-second sliding window. For training, we use a combination of VoxCeleb 1 and VoxCeleb 2 [34, 35] augmented with additive noise and reverberation according to the recipe from [28]. Segments under 4-second duration are discarded, resulting in a training set with 7,323 speakers. Reverberation is added by convolution with room responses from the RIR dataset [36], while additive noises are drawn from MUSAN [37]. At test time, x-vectors are extracted from 1.5 sec segments with a 0.25 sec shift. X-vectors are centered and whitened using statistics estimated from the DIHARD III DEV and EVAL sets, followed by length normalization [38].

The x-vectors are then clustered using AHC and a similarity matrix produced by scoring with a Gaussian PLDA model [30]. The PLDA model was trained using centered, whitened, and length normalized x-vectors extracted from VoxCeleb segments with duration ≥ 3 sec. Prior to PLDA scoring, dimensionality reduction was performed using conversation-dependent

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

Table 2: Baseline SAD results for the core/full DEV and EVAL sets. The **Part.** column indicates whether scoring was performed using the full or core DEV/EVAL set.

Set	Part.	Miss (%)	FA (%)	Overall error (%)
Dev	Core	1.84	3.98	2.30
	Full	1.88	4.55	2.42
Eval	Core	4.97	15.07	7.26
	Full	4.35	14.65	6.51

Table 3: Baseline track 1 diarization results for the core/full DEV and EVAL sets with/without VB-HMM resegmentation.

Part.	VB-HMM reseg.	DER (%)		JER (%)	
		Dev	Eval	Dev	Eval
Core	No	21.05	21.66	46.34	48.10
Core	Yes	20.25	20.65	46.02	47.74
Full	No	20.71	20.75	42.44	43.31
Full	Yes	19.41	19.25	41.66	42.45

PCA [4] preserving 30% of the total variability. For each track, the stopping criteria for AHC was tuned to minimize DER on the DEV set.

We then refine the AHC output using frame-level VB-HMM resegmentation as described by [32, 33]. 24-D MFCCs are extracted every 10 ms using a 15 ms window; neither mean nor variance normalization are applied, nor do we use delta coefficients. We use a universal background model (UBM-GMM) with 1,024 diagonal-covariance components and a total variability (\mathbf{V}) matrix containing 400 eigenvoices. Both the UBM-GMM and \mathbf{V} were trained using the same data as the x-vector extractor. Following [27], posterior scaling was applied to discourage frequent speaker transitions by the VB-HMM. This scaling was accomplished by boosting the zeroth order, but not first or second order, statistics prior to VB-HMM likelihood computation. The VB-HMM is initialized separately for each recording from the result of AHC and run for one iteration. Parameters were set to the following values by tuning on the DIHARD III DEV set: scaling factor $\beta = 10$, loop probability $P_{loop} = 0.45$, downsampling factor $downSamp = 25$.

DER and JER of the baseline diarization recipe for tracks 1 and 2 are reported in Tables 3 and 4, respectively. Mirroring the findings of [18, 39], VB-HMM resegmentation reliably improves DER and JER for both tracks, though the gains are more pronounced for track 2. Possibly, the effects of VB-HMM resegmentation could be enhanced by using a UBM-GMM and variability matrix trained on or adapted to domain materials, though we did not explore this possibility.

Table 4: Baseline track 2 diarization results for the core/full DEV and EVAL sets with/without VB-HMM resegmentation.

Part.	VB-HMM reseg.	DER (%)		JER (%)	
		Dev	Eval	Dev	Eval
Core	No	24.06	29.51	49.17	53.82
Core	Yes	22.28	27.34	47.75	51.91
Full	No	24.08	28.00	45.61	49.35
Full	Yes	21.71	25.36	43.66	46.95




Figure 1: Track 1 (red) and track 2 (blue) DER of primary submissions by domain on the core EVAL set.

6. Results and Discussion

Figure 2 presents boxplots of DER and JER on the core/full EVAL sets for each team’s final submissions. The majority of submissions outperformed the baseline with the mean improvement in DER for these submissions being nearly 4% absolute for track 1 and 5% for track 2. Restricting our attention to just the top submissions, we see that the improvement is even more extreme with the winning submission [40] outperforming the baseline by 7% DER absolute on the core EVAL set for track 1 and 8% on track 2. As was the case in previous DIHARD challenges, track 2 was substantially more difficult than track 1 (on the order of 5% for both DER and JER), indicating SAD remains a challenging problem for some of the covered domains. Though the full EVAL set was more difficult than the core set, the same trends and rankings are observed for both partitions; for sake of exposition, we will report only results on the core partition throughout the rest of this discussion.

As can be seen from Figure 1, system performance varies greatly across domains. For track 1, the median DER is below 10% for 6 domains. With the exception of the courtroom data, these domains all consist of 1 or 2 party interactions, from which we conclude that systems are able to reliably handle clean data from a small number of speakers when a high-quality speech segmentation is available. However, in the absence of an accurate speech segmentation, for all of these domains save audiobooks DER increases substantially for all systems and catastrophically for some; particularly, for the sociolinguistic field recordings. For two of the remaining domains, clinical interviews and conversational telephone speech, DER generally ranges from 10% to 20% for track 1 and 15% to 25% for track 2, indicating that even in the two-party case there remain substantial challenges, particularly in the presence of speakers with unusual characteristics (i.e., children in the clinical interviews). For the final three domains – meeting speech, web videos, and restaurant – performance is awful with median track 1 DER ranging from 35% to 45%.

Compared to DIHARD I and II, there is a notable improvement in performance. Median DER on track 1 has fallen from over 30% in DIHARD I, to 25% in DIHARD II, to under 20% in the current iteration. The improvement is even more striking for track 2, where median DER has fallen from 40% in DIHARD I to under 25% in DIHARD III, a 38% relative reduction. Considering only performance of the best single system in each evaluation, DER has decreased by 43% for track 1 (from 23.73% to 13.45%) and 46% for track 2 (35.51% to 19.37%). These improvements are observed across all domain, including those which have not changed since DIHARD I (e.g., courtroom

and meeting), indicating that these are real improvements and not an artifact of changes to the composition and annotation of the evaluation set over time. From a survey of the submitted system descriptions, these improvements appear to be due to a combination of improvements in handling of speaker overlap, use of VB-HMM in place of traditional frame-level clustering [32, 33], supervised neural speech enhancement models, and target-speaker based voice activity detection (TS-VAD) [41], and system combination methods such as DOVER-Lap [42]. While not utilized by the winning submission, end-to-end approaches [43] were prominent in DIHARD III and exhibited strong performance on both tracks.




Figure 2: DER and JER of primary submissions for each track on the core (red) and full (blue) partitions of the EVAL set.

7. Summary

We present a summary of DIHARD III, whose objective was to evaluate the current state of robust speaker diarization. Results from this present challenge indicate great progress has been made since DIHARD I, although performance of even the best system was poor for half of the domains considered, particularly when a reference SAD was not supplied. While encouraging, these results demonstrate that further work is needed to achieve the desired goal of truly robust diarization that gracefully handles the wide variety of interaction types and recording conditions observed in real world data.

8. References

- [1] S. Watanabe, M. Mandel, J. Barker *et al.*, “CHiME-6 challenge: Tackling multispeaker speech recognition for unsegmented recordings,” *arXiv preprint arXiv:2004.09249*, 2020.
- [2] J. G. Fiscus, J. Ajot, M. Michel *et al.*, “The Rich Transcription 2006 Spring Meeting Recognition Evaluation,” in *International Workshop on Machine Learning for Multimodal Interaction*. Springer, 2006, pp. 309–322.
- [3] G. Sell and D. Garcia-Romero, “Speaker diarization with PLDA i-vector scoring and unsupervised calibration,” in *Proc. SLT*, 2014, pp. 413–417.
- [4] W. Zhu and J. Pelecanos, “Online speaker diarization using adapted i-vector transforms,” in *Proc. ICASSP*, 2016.
- [5] D. Garcia-Romero, D. Snyder, G. Sell *et al.*, “Speaker diarization using deep neural network embeddings,” in *Proc. ICASSP*, 2017, pp. 4930–4934.
- [6] Q. Wang, C. Downey, L. Wan *et al.*, “Speaker diarization with LSTM,” in *Proc. ICASSP*, 2018, pp. 5239–5243.
- [7] A. Zhang, Q. Wang, Z. Zhu *et al.*, “Fully supervised speaker diarization,” *Proc. ICASSP*, 2019.
- [8] M. Rouvier, G. Dupuy, P. Gay *et al.*, “An open-source state-of-the-art toolbox for broadcast news diarization,” in *Proc. Interspeech*, 2013, pp. 1477–1481.
- [9] I. Viñals, A. Ortega, J. A. V. López *et al.*, “Domain adaptation of PLDA models in broadcast diarization by means of unsupervised speaker clustering,” in *Proc. Interspeech*, 2017, pp. 2829–2833.
- [10] S. H. Yella and H. Bourlard, “Improved overlap speech diarization of meeting recordings using long-term conversational features,” in *Proc. ICASSP*, 2013, pp. 7746–7750.
- [11] S. H. Yella, A. Stolcke, and M. Slaney, “Artificial neural network features for speaker diarization,” in *Proc. SLT*, 2014, pp. 402–406.
- [12] J. H. Hansen, A. Joglekar, M. C. Shekhar *et al.*, “The 2019 inaugural Fearless Steps challenge: A giant leap for naturalistic audio,” in *Proc. Interspeech*, 2019, pp. 1851–1855.
- [13] A. Joglekar, J. H. Hansen, M. C. Shekar *et al.*, “Fearless Steps challenge (fs-2): Supervised learning with massive naturalistic apollo data,” *arXiv preprint arXiv:2008.06764*, 2020.
- [14] E. Lleida, A. Ortega, A. Miguel *et al.*, “Albayzin 2018 evaluation: the iberspeech-rtve challenge on speech technologies for spanish broadcast media,” *Applied Sciences*, vol. 9, no. 24, p. 5412, 2019.
- [15] A. Nagrani, J. S. Chung, J. Huh *et al.*, “Voxsrc 2020: The second voxceleb speaker recognition challenge,” *arXiv preprint arXiv:2012.06867*, 2020.
- [16] N. Ryant, K. Church, C. Cieri *et al.*, “First DIHARD challenge evaluation plan,” Tech. Rep., 2018. [Online]. Available: <https://zenodo.org/record/1199638>
- [17] G. Sell, D. Snyder, A. McCree *et al.*, “Diarization is hard: Some experiences and lessons learned for the JHU team in the inaugural DIHARD Challenge,” in *Proc. Interspeech*, 2018, pp. 2808–2812.
- [18] M. Diez, F. Landini, L. Burget *et al.*, “BUT system for DIHARD Speech Diarization Challenge 2018,” in *Proc. Interspeech*, 2018, pp. 2798–2802.
- [19] C. Cieri, D. Miller, and K. Walker, “From Switchboard to Fisher: Telephone collection protocols, their uses and yields,” in *Proc. EUROSPEECH*, 2003.
- [20] N. Ryant, K. Church, C. Cieri *et al.*, “Second DIHARD challenge evaluation plan,” Tech. Rep., 2019. [Online]. Available: <https://coml.lscp.ens.fr/dihard/2019/second.dihard.eval.plan.v1.1.pdf>
- [21] —, “The Second DIHARD Diarization Challenge: Dataset, task, and baselines,” *arXiv preprint arXiv:1906.07839*, 2019.
- [22] —, “Third DIHARD Challenge evaluation plan,” *arXiv preprint arXiv:2006.05815*, 2020.
- [23] L. Hamers *et al.*, “Similarity measures in scientometric research: The Jaccard index versus Salton’s cosine formula,” *Information Processing and Management*, vol. 25, no. 3, pp. 315–18, 1989.
- [24] R. Real and J. M. Vargas, “The probabilistic basis of Jaccard’s index of similarity,” *Systematic Biology*, vol. 45, no. 3, pp. 380–385, 1996.
- [25] V. Peddinti, D. Povey, and S. Khudanpur, “A time delay neural network architecture for efficient modeling of long temporal contexts,” in *Proc. Interspeech*, 2015.
- [26] P. Ghahremani, V. Manohar, D. Povey *et al.*, “Acoustic modelling from the signal domain using CNNs,” in *Proc. Interspeech*, 2016, pp. 3434–3438.
- [27] P. Singh, M. Harsha Vardhan, S. Ganapathy *et al.*, “LEAP diarization system for the second DIHARD challenge,” 2019, pp. 983–987.
- [28] D. Snyder, P. Ghahremani, D. Povey *et al.*, “Deep neural network-based speaker embeddings for end-to-end speaker verification,” in *Proc. SLT*, 2016, pp. 165–170.
- [29] D. Snyder, D. Garcia-Romero, G. Sell *et al.*, “X-vectors: Robust DNN embeddings for speaker recognition,” in *Proc. ICASSP*, 2018, pp. 5329–5333.
- [30] S. J. Prince and J. H. Elder, “Probabilistic linear discriminant analysis for inferences about identity,” in *2007 IEEE 11th International Conference on Computer Vision*, 2007, pp. 1–8.
- [31] K. J. Han, S. Kim, and S. S. Narayanan, “Strategies to improve the robustness of agglomerative hierarchical clustering under data source variation for speaker diarization,” *IEEE Trans. Audio, Speech, Language Process*, vol. 16, no. 8, pp. 1590–1601, 2008.
- [32] M. Diez, L. Burget, and P. Matejka, “Speaker diarization based on Bayesian HMM with eigenvoice priors,” in *Proc. Odyssey*, 2018, pp. 147–154.
- [33] M. Diez, L. Burget, F. Landini *et al.*, “Analysis of speaker diarization based on Bayesian HMM with eigenvoice priors,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 355–368, 2019.
- [34] A. Nagrani, J. S. Chung, and A. Zisserman, “VoxCeleb: a large-scale speaker identification dataset,” *arXiv preprint arXiv:1706.08612*, 2017.
- [35] J. S. Chung, A. Nagrani, and A. Zisserman, “VoxCeleb2: Deep speaker recognition,” *Proc. Interspeech*, pp. 1086–1090, 2018.
- [36] T. Ko, V. Peddinti, D. Povey *et al.*, “A study on data augmentation of reverberant speech for robust speech recognition,” in *Proc. ICASSP*, 2017, pp. 5220–5224.
- [37] D. Snyder, G. Chen, and D. Povey, “MUSAN: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [38] D. Garcia-Romero and C. Y. Espy-Wilson, “Analysis of i-vector length normalization in speaker recognition systems,” in *Proc. Interspeech*, 2011, pp. 249–252.
- [39] F. Landini, S. Wang, M. Diez *et al.*, “BUT system for the Second DIHARD Speech Diarization Challenge,” in *Proc. ICASSP*. IEEE, 2020, pp. 6529–6533.
- [40] Y. Wang, M. He, S. Niu *et al.*, “USTC-NELSLIP system description for DIHARD III challenge,” *arXiv preprint arXiv:2103.10661*, 2021.
- [41] I. Medennikov, M. Korenevsky, T. Prisyach *et al.*, “Target-Speaker Voice Activity Detection: a novel approach for multi-speaker diarization in a dinner party scenario,” *arXiv preprint arXiv:2005.07272*, 2020.
- [42] D. Raj, L. P. Garcia-Perera, Z. Huang *et al.*, “DOVER-Lap: A method for combining overlap-aware diarization outputs,” in *Proc. SLT*, 2021, pp. 881–888.
- [43] S. Horiguchi, N. Yalta, P. Garcia *et al.*, “The Hitachi-JHU DIHARD III system: Competitive end-to-end neural diarization and x-vector clustering systems combined by DOVER-Lap,” *arXiv preprint arXiv:2102.01363*, 2021.