



# Speaker Diarization using Two-pass Leave-One-Out Gaussian PLDA Clustering of DNN Embeddings

Kiran Karra, Alan McCree

Human Language Technology Center of Excellence  
The Johns Hopkins University, Baltimore, MD, USA

kiran.karra@jhuapl.edu, alan.mccree@jhu.edu

## Abstract

Many modern systems for speaker diarization, such as the recently-developed VBx approach, rely on clustering of DNN speaker embeddings followed by resegmentation. Two problems with this approach are that the DNN is not directly optimized for this task, and the parameters need significant retuning for different applications. We have recently presented progress in this direction with a Leave-One-Out Gaussian PLDA (LGP) clustering algorithm and an approach to training the DNN such that embeddings directly optimize performance of this scoring method. This paper presents a new two-pass version of this system, where the second pass uses finer time resolution to significantly improve overall performance. For the Callhome corpus, we achieve the first published error rate below 4% without any task-dependent parameter tuning. We also show significant progress towards a robust single solution for multiple diarization tasks.

**Index Terms:** speaker diarization, x-vector, probabilistic linear discriminant analysis

## 1. Introduction

As highlighted in the recent DIHARD challenges (2018, 2019), there are many techniques for speaker diarization [1, 2, 3, 4, 5, 6, 7, 8]. Traditional approaches involve clustering of speaker embeddings (xvectors) followed by resegmentation, but we have concerns that the DNN is not directly optimized for this task and the parameters need significant retuning for different applications. The research community has recently explored end-to-end (E2E) approaches to overcome these drawbacks [9, 10, 11, 12, 13]. However, while E2E approaches are desirable from a philosophical perspective and can handle overlapping speech segments, they do not yet attain the best performance.

In this work, we begin with the Leave-One-Out Gaussian PLDA (LGP) approach, which clusters DNN embeddings using a Gaussian Mixture Model (GMM) [14]. In order to improve performance without task-dependent tuning, we introduce a two-pass LGP algorithm. Here, each pass examines the audio to be diarized at different segment lengths and overlaps, to fine-tune speaker assignments for speech segments. During each pass, the speaker assignment algorithm alternates between updating speaker models and generating speaker posteriors, while leaving out segments which were included in the model estimation to reduce bias.

The paper is organized as follows: we begin by reviewing the original LGP algorithm and DNN training process. We then motivate and explain the two-pass algorithm in detail. Modifications needed to the LGP algorithm to accommodate varying tasks are addressed. Finally, we discuss and compare the performance of this system against other published diarization sys-

tems on four separate datasets: 1) Callhome, 2) DIHARD2, 3) AMI Beamformed, and 4) AMI Mixed Headset.

## 2. LGP Diarization

In this section we review the LGP model as first presented in [14]. We use the PLDA generative model [15, 16] for clustering of DNN segment embeddings. To tackle the problem of joint estimation of speaker models and segment alignments, we use a leave-one-out (LOO) method to replace the VB approach [17, 8]. Besides giving a practical way to overcome the inherent bias of scoring models against segments which were included in the model estimation, this also allows us to improve performance by removing the independence assumption of PLDA.

### 2.1. Clustering

In more detail, the LGP algorithm works as follows. Given inputs of the length-normalized segment embeddings, initial segment posteriors over speakers, and PLDA parameters (within-class and across-class covariance), alternate between updating the speaker models and generating segment speaker posteriors.

#### 2.1.1. Model Update

Given the Gaussian PLDA model with known covariances  $\Sigma_{wc}$  and  $\Sigma_{ac}$  and a set of  $N$  enrollment segments  $\mathbf{z}_n$ , we update the posterior distribution of the speaker model  $S_i$  which is Gaussian [18] with mean:

$$\mathbf{m}_i = \Sigma_{ac} (\Sigma_{ac} + \Sigma_{ml})^{-1} \bar{\mathbf{z}}_{ml} \tag{1}$$

and covariance:

$$\Sigma_i = \Sigma_{ac} (\Sigma_{ac} + \Sigma_{ml})^{-1} \Sigma_{ml} \tag{2}$$

where  $\bar{\mathbf{z}}_{ml} = \frac{1}{N} \sum_{n=1}^N \mathbf{z}_n$  and  $\Sigma_{ml} = \frac{\Sigma_{wc}}{N}$  represent the maximum likelihood (ML) mean estimate and the covariance of this estimator. To update the speaker priors (weights), we follow [19, 20] and use a non-Bayesian maximum-likelihood approach, as it has been observed to have good properties of eliminating redundant speakers.

#### 2.1.2. Speaker Assignment

The updated speaker models (means, covariances, and weights) form a Gaussian mixture model, so speaker assignment is done by computing LOO posteriors per class (responsibilities). LOO is implemented by leaving out the current sample,  $n$ , when computing GMM model parameter updates over all enrollment segments<sup>1</sup>. Posteriors are computed using the predictive distribu-

<sup>1</sup>See: <https://github.com/hltcoe/VBx> for more details

tion, which is again Gaussian:

$$\mathbf{z}_n | S_i \sim \mathcal{N}(\mathbf{m}_i, \Sigma_{wc} + \Sigma_i). \quad (3)$$

### 2.1.3. Dimension Reduction and Diagonalization

To reduce computation, this work uses diagonal PLDA covariance matrices based on the fact that two symmetric matrices can be simultaneously diagonalized with a linear transformation [21]. This process is similar to Linear Discriminant Analysis (LDA), and results in a transformed space where  $\Sigma_{wc} = I$  and  $\Sigma_{ac}$  is diagonal.

### 2.1.4. Removing the Independence Assumption

To better model the correlation between consecutive segments, [22] introduced non-independent enrollment update equations. For this model, the enrollment still follows Eq. 1 with the same ML mean, but the covariance of this ML mean estimator now decreases more slowly with increasing number of enrollment segments:

$$\Sigma_{ml} = \frac{\Sigma_{wc}}{N} \left( 1 + 2 \sum_{j=1}^{N-1} \frac{(N-j)}{N} r^j \right) \quad (4)$$

The parameter  $r$  represents the correlation between successive channel draws, and allows continuous variation between the two extremes of “by-the-book PLDA scoring” ( $r = 0$ ) and “average i-vector scoring” ( $r = 1$ ).

### 2.1.5. Selecting Number of Speakers

In all clustering applications, selecting the number of clusters is a challenging problem. While AHC with PLDA comparisons does work well for this task, it works best with a tuned, task-dependent stopping threshold. As in [20], we prefer to start with a maximum number of speakers and let the clustering algorithm automatically select the correct number. While LGP does produce overall likelihood estimates for any number of speakers, in practice the ML weight updates quickly reach zero for unnecessary speakers. We initialize the algorithm with k-means in the diagonalized embedding space with a fixed number of speakers, and let the algorithm eliminate speaker weights across iterations.

## 2.2. DNN Training

The baseline network architecture is a TDNN shown in Table 1. This represents an extension of the x-vector architecture [23], where the classification layers have been modified to match the diarization task. First, the traditional ReLU nonlinearity has been replaced by length normalization, since this is well-known to improve the performance of Gaussian PLDA [24]. To implement PLDA scoring in the DNN, we add a Gaussian quadratic layer. This layer maintains enrollment statistics for each class using a few (1-20) past values of embeddings for that class, and then uses the PLDA parameters  $\Sigma_{wc}$  and  $\Sigma_{ac}$  to perform Bayesian model enrollment using Eq. 1 and 2. Finally, log-likelihoods for each class are produced with the predictive distributions from Eq. 3.

The PLDA parameters are estimated in the following way. First, the within-class covariance  $\Sigma_{wc}$  is set to identity, as we again assume the DNN can force the embeddings to match this property. We do need a discriminatively-trained scale factor to compensate for the length-normalization constraint. For the across-class covariance  $\Sigma_{ac}$ , we use a generative update of this

Table 1: Baseline x-vector architecture. These experiments use a model with a layer size of 768 and an embedding dimension of 128.

Layer	Layer Type	Context	Size
1	TDNN-ReLU	t-2:t+2	L
2	Dense-ReLU	t	L
3	TDNN-ReLU	t-2, t, t+2	L
4	Dense-ReLU	t	L
5	TDNN-ReLU	t-3, t, t+3	L
6	Dense-ReLU	t	L
7	TDNN-ReLU	t-4, t, t+4	L
8	Dense-ReLU	t	L
9	Dense-ReLU	t	3*L
10	Pooling (mean+stddev)	Full-seq	6*L
11	Dense(Embedding)		D
12	Length-norm		D
13	Gauss quadratic-Softmax		Num. spks.

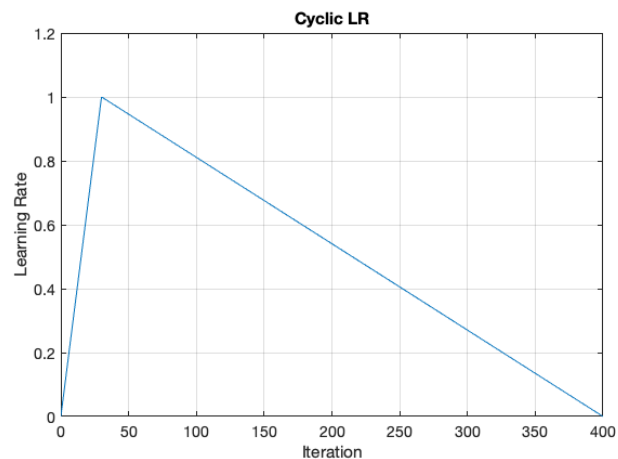


Figure 1: Learning rate schedule used for training xvector DNN.

matrix. The PLDA layer, which estimates the global PLDA parameters, maintains its own separate copy of enrollment statistics over previous embeddings for each class, and  $\Sigma_{ac}$  is approximated by the covariance of these ML model means.

The training cost is normalized multiclass cross-entropy. To compensate for possible instability of generative parameter updates which are not exposed to gradient optimization, the learning rate uses a linear ramp-up from zero for the initial few training epochs. The learning rate schedule is shown in Fig. 1.

## 3. Algorithm Improvements

We have made a number of improvements to the LPG system, including replacing the TDNN with a ResNet, introducing a second pass for time refinement, and improving the duration modeling function.

### 3.1. ResNet

For wideband speech signals, we have experimented with a modified ResNet architecture shown in Table 2 [25]. The Angular Margin Softmax described in [25] is replaced with a Gauss quadratic-Softmax to allow PLDA scoring within the DNN, as

Table 2: Modified ResNet-34 architecture with 15.4 million parameters. Batch-norm and ReLU layers are not shown. The  $1 \times 1$  convolutions are used to match the dimensions for the residual connections. The dimensions are (Channels  $\times$  Frequency  $\times$  Time). The input comprises 80 Melfilter bank energies from speech segments. During training we use a fixed segment length of  $T = 400$

Layer Name	Structure	Output ( $C \times F \times T$ )
Input	—	$1 \times 80 \times T$
Conv2D	$3 \times 3$ , stride=1	$128 \times 80 \times T$
ResBlock-1	$\begin{matrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{matrix} \times 3$ , stride=1	$128 \times 80 \times T$
ResBlock-2a	$\begin{matrix} 3 \times 3, 128 \\ 3 \times 3, 128 \\ 1 \times 1, 128 \end{matrix} \times 1$ , stride=2	$128 \times 40 \times T/2$
ResBlock-2b	$\begin{matrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{matrix} \times 3$ , stride=1	$128 \times 40 \times T/2$
ResBlock-3a	$\begin{matrix} 3 \times 3, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 256 \end{matrix} \times 1$ , stride=2	$256 \times 20 \times T/4$
ResBlock-3b	$\begin{matrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{matrix} \times 5$ , stride=1	$256 \times 20 \times T/4$
ResBlock-4a	$\begin{matrix} 3 \times 3, 256 \\ 3 \times 3, 256 \\ 1 \times 1, 256 \end{matrix} \times 1$ , stride=2	$256 \times 10 \times T/8$
ResBlock-4b	$\begin{matrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{matrix} \times 2$ , stride=1	$256 \times 10 \times T/8$
Flatten ( $C, F$ )	—	$2560 \times T/8$
StatsPooling	—	5120
Dense (Emb.)	—	128
Length-norm	—	128
Gauss quadratic-Softmax	—	Num. spks.

described in Section 2.2.

### 3.2. Two Pass Approach

We motivate the need for a two-pass enhancement to the LPG algorithm by noting that operating the LPG algorithm does not work well if the segment length is small. Long enough segments are needed to: 1) extract enough information from the speech segment to train an effective embedding, and 2) have enough stable SAD information to process the correct speech segments. However, longer segments translate to coarser speaker assignments in the diarization problem. To enable fine temporal assignment of speakers to segments, we improve the LGP algorithm by adding a second pass.

The two pass modification works by running the LGP algorithm twice, with different segment lengths and overlaps. The first pass is unchanged from our baseline approach, with a larger segment length (2 seconds) and no overlap. This configuration provides good performance overall. The second pass refines these diarization marks with a smaller segment length (1.25 sec) and high overlap (1 sec), to allow for temporal refinement of speaker assignments. We find that only 1 or 2 iterations are needed for this refinement process, as the clustering does not typically change very much.

Since the DNN is trained with particular target segment durations in mind, we find a small additional performance improvement by fine-tuning DNNs for each duration.

### 3.3. Improved Duration Modeling

We have addressed two areas of the correlation model for reducing effective counts and increasing uncertainty. First, the relation between the effective number of samples in a cluster and

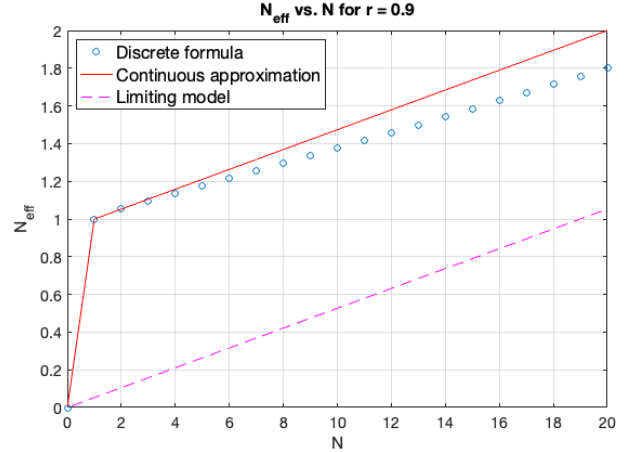


Figure 2: Comparison of functions for approximating effective number of samples.

the actual count is a somewhat complicated discrete formula:

$$N_{eff} = \frac{N}{1 + 2 \sum_{j=1}^{N-1} \frac{(N-j)}{N} r^j}. \quad (5)$$

In previous work, we have used the limiting approximation for large N of

$$N_{eff} = \left( \frac{1-r}{1+r} \right) N. \quad (6)$$

However, as shown in Fig. 2, this limiting function is not a close fit for small N with  $r = 0.9$ . Instead we introduce a more accurate continuous approximation:

$$N_{eff} = \min \left( N, \frac{(1-r)N + 2r}{1+r} \right). \quad (7)$$

Using this continuous approximation in place of our previous interpolation of the discrete version is much simpler to implement with no loss in performance.

Secondly, we have found that diverse tasks in diarization can have wildly varying audio lengths. In particular, long audio files result in a large number of segments in each speaker cluster, which reduces uncertainty and does not match our DNN training conditions. Therefore, we introduce a new parameter  $N_0$  to represent the target number of segments in a file, and reduce the observed counts across longer files by the scale factor  $N_0/N$ .

## 4. Experimental Results and Analysis

We train the DNN with LDC corpora Switchboard, Fisher, Mixer6, SRE2004-10, and VoxCeleb1 [26]. With augmentations, this results in a set of 5,175,668 utterances from 13,129 speakers. We use a 90/10 split between training and development sets, yielding a training set of 11,816 speakers. Segments vary between 1.5 and 2.5 seconds. Note that the DNN contains everything needed for diarization: both the speaker embedding (xvector) and internal PLDA parameters. No additional system training is needed.

### 4.1. Callhome

As our primary focus is telephone speech, we begin our evaluation of the two-pass LGP algorithm on the Callhome dataset.

Following standard practice, we report results for oracle speech activity marks with 250 ms forgiveness collar around speaker change points, and do not score segments with overlapping speakers. The results are shown in Table 3, where the ‘System’ column indicates the system being tested. Here, we set  $N_0 = N$  since the number of segments in each audio file matches the DNN training conditions. The results show that the proposed two-pass LGP system exceeds the best current published performance.

Table 3: *DER for various diarization algorithms for the Callhome dataset.*

System	DER
AHC	8.10
AHC+VB [1]	6.48
One-Pass LGP [14]	6.62
VBx [8]	4.42
Two-Pass LGP	3.92

In order to characterize the performance of each stage of the two-pass LGP algorithm, we also compute the DER after k-means initialization of the GMM clusters. The DER decreases between the clustering initialization with k-means and the first pass by 81.9%. The DER decreases between the first pass and the second pass by an additional 40.8%. The performance increase across each stage of processing indicates that while k-means is a good initializer of the GMM for clustering, by-itself, it does not achieve acceptable performance in assigning speakers to speech segments. The first pass of assignment dramatically helps overall performance by refining initial speaker assignments provided by k-means. As hypothesized, the second pass fine-tunes the speaker assignments successfully, and results in additional performance gain. With these results, we emphasize that the two-pass LGP algorithm used for the Callhome dataset is an integrated system and was not tuned on any development dataset.

#### 4.2. Wideband Datasets

Next, we evaluate the proposed two-pass LGP algorithm on the wideband datasets: DIHARD2 [27], AMI-Headset, and AMI-Beamformed [28], using the same system as described above for Callhome. Here, we set  $N_0 = 25$  to account for varying audio segment lengths as described in Section 3.3 for all evaluations conducted with the narrowband xvector system described in Section 2.2. Because the DIHARD2 and AMI datasets are wideband, we also evaluate the two-pass LGP algorithm with an xvector system trained on wideband data. The wideband xvector system is trained with Voxceleb1 [26] and Voxceleb2 [29]. With augmentations, this results in a set of 5,949,009 utterances from 6,825 speakers. We use a 90/10 split between training and development sets, yielding a training set of 6,143 speakers, with minibatches of size 128 with one segment per speaker. The remainder of training is exactly the same as what was reported for the narrowband xvector system. We set  $N_0 = 30$  for the wideband xvector system, to account for varying audio segments.

We compare both the two-pass narrowband LGP algorithm (2P-NB-LGP) and the two-pass wideband LGP algorithm (2P-WB-LGP) to the VBx algorithm, using the task dependent tuning factors specified by the authors of that algorithm [8]. For the DIHARD2 dataset, we evaluate using no collar and score overlapping speech segments. For AMI, we follow the evalua-

tion protocol specified by Section 4 in [8], and show results for evaluating with no collar and scoring overlapping speech segments.

Table 4 aggregates the results, and shows that both two-pass LGP systems are competitive across all tasks. Both the 2P-NB-LGP and 2P-WB-LGP exceed the best published DER performance for the evaluation split for both AMI-Beamformed (AMI-BF) and AMI-Mixed Headset (AMI-Mx) datasets. The 2P-WB-LGP system shows competitive performance in the evaluation split of the DIHARD2 dataset. Using the  $N_0$  robustness tuning factor, we revisit the performance of the 2P-NB-LGP algorithm on the Callhome dataset. A configuration of  $N_0 = 25$  reduces the DER performance on Callhome by 14.4%. We emphasize that while the NB and WB systems indicated by “2P-NB-LGP” and “2P-WB-LGP” in Table 4 are fundamentally different, each uses the same hyperparameters across the datasets tested.

Table 4: *Diarization performance on DIHARD2 and AMI datasets*

Task	System	Dev	Eval
DIHARD2	Landini [30]	17.90	18.21
	Lin [31]	21.36	18.84
	Lin [32]	18.76	18.44
	VBx	18.19	18.55
	2P-NB-LGP	19.21	20.83
	2P-WB-LGP	17.8	18.76
AMI-BF	VBx	17.66	20.84
	2P-NB-LGP	18.56	19.95
	2P-WB-LGP	18.94	19.84
AMI-Mx	VBx	16.33	18.99
	2P-NB-LGP	16.21	17.57
	2P-WB-LGP	16.74	17.79

## 5. Conclusion

This paper has presented the first published number below 4% DER for the Callhome benchmark task, without any task-dependent parameter estimation or tuning. This is achieved through significant improvements to the LGP diarization system, and particularly the two-pass approach to improve time resolution. Additionally, since the LGP DNN is directly trained to optimize PLDA performance over segment embeddings, no further parameter training is needed to perform diarization. The single narrowband system also provides competitive performance across a range of tasks (Callhome, DIHARD2, and AMI). In addition, a single wideband version of the system can essentially match the tuned performance of the VBx approach over both DIHARD2 and AMI.

## 6. References

- [1] G. Sell, D. Snyder, A. McCree, D. Garcia-Romero, J. Villalba, M. Maciejewski, V. Manohar, N. Dehak, D. Povey, S. Watanabe *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.
- [2] M. Diez, F. Landini, L. Burget, J. Rohdin, A. Silnova, K. Zmoliková, O. Novotný, K. Veselý, O. Glombek, O. Plchot *et al.*, “But system for DIHARD speech diarization challenge 2018,” in *Proc. Interspeech*, 2018, pp. 2798–2802.

- [3] 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," *Proc. Interspeech 2018*, pp. 2793–2797, 2018.
- [4] Z. Zajic, M. Kunešová, J. Zelinka, and M. Hruš, "Zcu-ntis speaker diarization system for the DIHARD 2018 challenge," in *Proc. INTERSPEECH*, 2018, pp. 2788–2792.
- [5] I. Vinals, P. Gimeno, A. Ortega, A. Miguel, and E. Lleida, "Estimation of the number of speakers with variational bayesian plda in the DIHARD diarization challenge," in *Proc. INTERSPEECH*, 2018, pp. 2803–2807.
- [6] J. Patino, H. Delgado, and N. Evans, "The eurecom submission to the first DIHARD challenge," in *Proc. Interspeech*, vol. 2018, 2018, pp. 2813–2817.
- [7] N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, and M. Liberman, "First DIHARD challenge evaluation plan," 2018.
- [8] F. Landini, J. Profant, M. Diez, and L. Burget, "Bayesian hmm clustering of x-vector sequences (vbx) in speaker diarization: theory, implementation and analysis on standard tasks," *arXiv preprint arXiv:2012.14952*, 2020.
- [9] Y. Fujita, S. Watanabe, S. Horiguchi, Y. Xue, and K. Nagamatsu, "End-to-end neural diarization: Reformulating speaker diarization as simple multi-label classification," *arXiv preprint arXiv:2003.02966*, 2020.
- [10] Z. Huang, S. Watanabe, Y. Fujita, P. García, Y. Shao, D. Povey, and S. Khudanpur, "Speaker diarization with region proposal network," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6514–6518.
- [11] I. Medennikov, M. Korenevsky, T. Prisyach, Y. Khokhlov, M. Korenevskaya, I. Sorokin, T. Timofeeva, A. Mitrofanov, A. Andrusenko, I. Podluzhny *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.
- [12] K. Kinoshita, M. Delcroix, and N. Tawara, "Integrating end-to-end neural and clustering-based diarization: Getting the best of both worlds," *arXiv preprint arXiv:2010.13366*, 2020.
- [13] S. Horiguchi, Y. Fujita, S. Watanabe, Y. Xue, and K. Nagamatsu, "End-to-end speaker diarization for an unknown number of speakers with encoder-decoder based attractors," *arXiv preprint arXiv:2005.09921*, 2020.
- [14] A. McCree, G. Sell, and D. Garcia-Romero, "Speaker diarization using leave-one-out gaussian plda clustering of dnn embeddings," in *Interspeech*, 2019, pp. 381–385.
- [15] S. Ioffe, "Probabilistic linear discriminant analysis," in *European Conference on Computer Vision*. Springer, 2006, pp. 531–542.
- [16] S. J. D. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in *Proc. ICCV*, 2007, pp. 1–8.
- [17] J. Villalba, A. Ortega, A. Miguel, and E. Lleida, "Variational bayesian plda for speaker diarization in the mgb challenge," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 667–674.
- [18] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Wiley, 2001.
- [19] P. Kenny, "Bayesian analysis of speaker diarization with eigen-voice priors," *CRIM, Montreal, Technical Report*, 2008.
- [20] M. Diez, L. Burget, and P. Matejka, "Speaker diarization based on bayesian hmm with eigenvoice priors," in *Proceedings of Odyssey*, 2018, pp. 147–154.
- [21] K. Fukunaga, *Introduction to Statistical Pattern Recognition*. Academic Press, 1990.
- [22] A. McCree, G. Sell, and D. Garcia-Romero, "Extended variability modeling and unsupervised adaptation for plda speaker recognition," in *Interspeech*, 2017, pp. 1552–1556.
- [23] D. Snyder, D. Garcia-Romero, G. Sell, A. McCree, D. Povey, and S. Khudanpur, "Speaker recognition for multi-speaker conversations using x-vectors," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [24] 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.
- [25] D. Garcia-Romero, G. Sell, and A. McCree, "Magneto: X-vector magnitude estimation network plus offset for improved speaker recognition," in *Proc. Odyssey 2020 The Speaker and Language Recognition Workshop*, 2020, pp. 1–8.
- [26] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," *Proc. Interspeech 2017*, pp. 2616–2620, 2017.
- [27] N. Ryant, K. Church, C. Cieri, A. Cristia, J. Du, S. Ganapathy, and M. Liberman, "The second dihard diarization challenge: Dataset, task, and baselines," *arXiv preprint arXiv:1906.07839*, 2019.
- [28] J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal *et al.*, "The ami meeting corpus: A pre-announcement," in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [29] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *arXiv preprint arXiv:1806.05622*, 2018.
- [30] F. Landini, S. Wang, M. Diez, L. Burget, P. Matějka, K. Žmolíková, L. Mošner, A. Silnova, O. Plchot, O. Novotný *et al.*, "But system for the second dihard speech diarization challenge," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6529–6533.
- [31] Q. Lin, W. Cai, L. Yang, J. Wang, J. Zhang, and M. Li, "Dihard ii is still hard: Experimental results and discussions from the dku-lenovo team," *arXiv preprint arXiv:2002.12761*, 2020.
- [32] Q. Lin, Y. Hou, and M. Li, "Self-attentive similarity measurement strategies in speaker diarization," in *Proc. Interspeech*, vol. 2020, 2020, pp. 284–288.