

DISCRIMINATIVE TRAINING OF VBx DIARIZATION

Dominik Klement¹, Mireia Diez¹, Federico Landini¹, Lukáš Burget¹, Anna Silnova¹,
Marc Delcroix², Naohiro Tawara²

¹Brno University of Technology, Speech@FIT, ²NTT Corporation, Japan

ABSTRACT

Bayesian HMM clustering of x-vector sequences (VBx) has become a widely adopted diarization baseline model in publications and challenges. It uses an HMM to model speaker turns, a generatively trained probabilistic linear discriminant analysis (PLDA) for speaker distribution modeling, and Bayesian inference to estimate the assignment of x-vectors to speakers. This paper presents a new framework for updating the VBx parameters using discriminative training, which directly optimizes a predefined loss. We also propose a new loss that better correlates with the diarization error rate compared to binary cross-entropy — the default choice for diarization end-to-end systems. Proof-of-concept results across three datasets (AMI, CALLHOME, and DIHARD II) demonstrate the method’s capability of automatically finding hyperparameters, achieving comparable performance to those found by extensive grid search, which typically requires additional hyperparameter behavior knowledge. Moreover, we show that discriminative fine-tuning of PLDA can further improve the model’s performance. We release the source code with this publication.

Index Terms— speaker diarization, VBx, clustering, variational Bayes, discriminative training

1. INTRODUCTION

Speaker diarization systems split audio recordings into speaker-homogenous segments and assign them speaker labels. Currently, three main approaches are used: clustering-based [1, 2], end-to-end [3, 4], or hybrid (combination of both)[5, 6]. While the latter two are the state-of-the-art [4, 6], clustering-based VBx diarization is competitive in wide-band multi-speaker domains [6]. Besides, VBx has been widely adopted in the community as a baseline both in publications [4, 7] and in diarization challenges [8, 9, 10] as well as a reclustering stage in some systems [11, 12].

VBx is a clustering-based approach [2], which requires as input a sequence of x-vectors [13] extracted from a uniformly segmented audio, and some corresponding initial labels, which are usually obtained by agglomerative hierarchical clustering (AHC). VBx further refines these speaker labels. The approach is based on a Bayesian hidden Markov model (HMM), where states correspond to speakers, and transitions correspond to speaker turns. Speaker-specific distributions are derived from a two-covariance probabilistic linear discriminant analysis (PLDA) model [14], pre-trained on a large number of speaker-labeled x-vectors. The model uses variational Bayes (VB) inference [15] to obtain the speaker distributions and the refined speaker labels.

Being a clustering-based method, different approaches and objectives are used to train different parts of the model: the x-vector extractor is trained discriminatively with a speaker classification objective; the PLDA is trained generatively to maximize the likelihood of

its training set, and VBx hyperparameters are selected by grid search to optimize diarization performance. Alternatively, end-to-end systems present a single neural model trained discriminatively for diarization but they require large amounts of (speaker-labeled) conversational speech data, which are scarce. Thus, artificially simulated training sets are built. However, they hinder system performance and are a field of research on its own [16, 17]. On the contrary, while the independent training of models used in clustering-based methods is seen as a downside, it allows for the usage of large amounts of real (single) speaker data to build robust speaker-discriminant models.

In the context of speaker recognition, discriminative training has been previously applied to generative models. In [18], a PLDA-like functional was proposed for evaluating the speaker verification scores, and its parameters were discriminatively trained using an objective function that directly addressed the speaker verification task, providing significant performance gains. In [19], an i-vector classifier was discriminatively trained using maximum mutual information to directly optimize the multi-class calibration criterion, to remove the need of using later a calibration backend.

In this paper, we propose a discriminative training framework for VBx (DVBx). With DVBx, all stages and parameters of the diarization pipeline can be jointly fine-tuned in an end-to-end fashion to minimize the diarization error directly. Compared to standard EEND approaches, which are usually trained from scratch on simulated speaker conversations, VBx uses an x-vector extractor trained on a large number of real speaker utterances [2], and a PLDA model trained on a large number of such x-vectors. These components provide the model with a strong speaker discriminative power and serve as a well-founded starting point for DVBx fine-tuning.

While we envision DVBx as a method capable of state-of-the-art results, the main focus of this paper is to introduce the methodology and present fair comparisons with the well-established original VBx. It is out of the scope of this paper to compare to current state-of-the-art approaches, which would require the use of VAD and overlap detection, and which are not used in the current work. In later sections, we will outline how we plan to extend the approach in directions that will produce results on par with end-to-end models.

To prove the validity of the method, we present a set of experiments that automatically find the set of VBx hyperparameters. Their tuning is usually done by an extensive, time-consuming grid search. Moreover, one cannot search through all possible combinations and has to restrict the parameter search space to a finite predefined set of values, which may lead to finding a suboptimal solution. We show that by automatically optimizing the hyperparameters only, we can find a setting achieving comparable performance to the baseline VBx model with parameters found by grid search [2] on three real datasets: AMI, CALLHOME, and DIHARD II. Moreover, we show that the performance can improve even further by fine-tuning the PLDA parameters. All our code is publicly available <https://github.com/BUTSpeechFIT/DVBx>.

2. MODEL DESCRIPTION

In this section, we briefly describe the necessary parts of the VBx model closely related to our work and refer the reader to the original VBx publication [2] for more details and the full description of the VB inference, the derivation of the update formulae or the definition of the hyperparameters.

VBx uses a Bayesian HMM to model the sequence of d -dimensional \mathbf{x} -vectors $\hat{\mathbf{X}} = (\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2, \dots, \hat{\mathbf{x}}_T)^\top \in \mathbb{R}^{T \times d}$ representing T consecutive short speech segments in an input conversation. To solve the diarization task, we need to infer the sequence of HMM latent variables $\mathbf{Z} = (z_1, z_2, \dots, z_T)^\top \in \mathbb{N}^T$ representing the hard assignment of \mathbf{x} -vectors to the HMM states, where each HMM state corresponds to one speaker in the conversation. The state transition probabilities are modeled as

$$p(z_t = s | z_{t-1} = s') = (1 - P_t)\pi_s + \delta(s = s')P_t, \quad (1)$$

where δ refers to the Kronecker delta function, P_t is the probability of staying in the same state, and π_s is the probability of switching to state/speaker s . The HMM state distributions (i.e., distributions of the \mathbf{x} -vectors corresponding to each speaker) are Gaussian with specific mean and shared within-speaker covariance matrix Σ_w . VBx uses a Bayesian setting where the speaker means are latent variables with Gaussian prior with global mean \mathbf{m} and between-speaker covariance Σ_b . The parameters \mathbf{m} , Σ_b , and Σ_w are taken from a pre-trained PLDA model. For mathematical convenience, we first center the original \mathbf{x} -vectors around the global mean \mathbf{m} and transform them by LDA-like linear transformation \mathbf{E} to obtain \mathbf{x} -vectors $\mathbf{X} = (\hat{\mathbf{X}} - \mathbf{1}\mathbf{m})\mathbf{E}$ with zero global mean, diagonal between-speaker covariance matrix Φ and identity within-speaker covariance matrix. The matrices \mathbf{E} and Φ can be obtained by solving the generalized eigenvalue problem $\Sigma_b\mathbf{E} = \Sigma_w\mathbf{E}\Phi$. For the transformed \mathbf{x} -vectors, the HMM state distributions are

$$p(\mathbf{x}_t | z_t = s) = \mathcal{N}(\mathbf{x}_t; \mathbf{V}\mathbf{y}_s, \mathbf{I}), \quad (2)$$

where $\mathbf{V} = \Phi^{\frac{1}{2}}$ and \mathbf{y}_s is the state/speaker specific latent variable with standard normal prior.

For each input conversation, VBx uses VB inference [15] to iteratively update the approximate posterior distributions $q(\mathbf{y}_s)$, responsibilities $\gamma_{ts} = q(z_t = s)$ and the speaker priors π_s . After convergence, the responsibilities γ_{ts} represent the (probabilistic) assignment of \mathbf{x} -vectors to HMM states/speakers and therefore provide the solution to the diarization problem.

The approximate posteriors $q(\mathbf{y}_s) = \mathcal{N}(\mathbf{y}_s; \boldsymbol{\alpha}_s, \mathbf{L}_s^{-1})$ are updated as:

$$\boldsymbol{\alpha}_s = \frac{F_A}{F_B} \mathbf{L}_s^{-1} \sum_t \gamma_{ts} \mathbf{V}^\top \mathbf{x}_t, \quad \mathbf{L}_s = \mathbf{I} + \frac{F_A}{F_B} \left(\sum_t \gamma_{ts} \right) \Phi, \quad (3)$$

where F_A and F_B are model hyperparameters (described below). The responsibilities γ_{ts} are updated using the standard (non-Bayesian) HMM forward-backward algorithm, except that the usual HMM state log-likelihood $\log p(\mathbf{x}_t | z_t)$ is replaced by a more complex term accounting for the uncertainty encoded in $q(\mathbf{y}_s)$:

$$\log \bar{p}(\mathbf{x}_t | z_t) \propto F_A \left[\boldsymbol{\alpha}_s^\top \mathbf{V} \mathbf{x}_t - \frac{1}{2} \text{Tr}[\Phi[\mathbf{L}_s^{-1} + \boldsymbol{\alpha}_s \boldsymbol{\alpha}_s^\top]] \right]. \quad (4)$$

Maximum likelihood type II updates are used for π_s (see equation (24) in [2]).

Scalars F_A, F_B are model hyperparameters that were introduced to control the VB inference. The acoustic scaling factor F_A

counteracts the independence assumption between observations. The speaker regularization coefficient F_B can be used to control the number of output speakers (a high F_B results in keeping fewer speakers).

The responsibilities γ_{ts} for the first VB iteration are typically initialized from labels obtained by preceding AHC. Let $\hat{\gamma}_t$ be a one-hot vector encoding the hard assignment of a speaker label to \mathbf{x} -vector \mathbf{x}_t as provided by AHC. Rather than using elements of such vectors directly as the initial responsibilities, we apply label smoothing:

$$\gamma_t = \text{softmax}(\tau \cdot \hat{\gamma}_t), \quad (5)$$

where a lower value of hyperparameter τ makes the initial assignment of \mathbf{x} -vectors to speakers more uncertain.

3. DISCRIMINATIVE TRAINING

Many diarization datasets have training and/or development data that can be used for tuning parameters. Conventional VBx uses an already trained (and fixed) \mathbf{x} -vector extractor and PLDA model, and such data are only used to tune its hyperparameters. In contrast, DVBx can utilize the same data to train all VBx parameters discriminatively. It is possible because VB update equations involve mathematical operations through which we can backpropagate loss gradients w.r.t model parameters, and the iterative VB inference can be unfolded, similarly to recurrent neural networks. In this section, we describe the details of the DVBx setup.

3.1. Losses

Similar to many end-to-end diarization approaches [3, 4], the model is optimized using permutation invariant training (PIT) scheme [20]:

$$\mathcal{L} = \frac{1}{TS} \min_{\phi \in \text{perm}(S)} \sum_{t=1}^T H(\gamma_t^\phi, \hat{\mathbf{l}}_t), \quad (6)$$

where $\text{perm}(S)$ is the set of all permutations of sequence $(1, 2, \dots, S)$, $\gamma_t^\phi = (\gamma_{t1}^\phi, \gamma_{t2}^\phi, \dots, \gamma_{tS}^\phi)^\top \in \langle 0, 1 \rangle^S$ represents the permuted responsibilities predicted by the VBx model for t -th frame, and $\hat{\mathbf{l}}_t = (\hat{l}_{t1}, \hat{l}_{t2}, \dots, \hat{l}_{tS})^\top \in \{0, 1\}^S$ are the ground truth labels indicating speaker activities of all S speaker in the t -th frame. For the loss function H , we first considered binary cross entropy (BCE) as normally used for end-to-end diarization system training:

$$H_B(\gamma_t^\phi, \hat{\mathbf{l}}_t) = \sum_{s=1}^S -\hat{l}_{ts} \log(\gamma_{ts}^\phi) - (1 - \hat{l}_{ts}) \log(1 - \gamma_{ts}^\phi). \quad (7)$$

However, VB inference is known to underestimate latent variable uncertainty, which leads to overly confident (close to either 0 or 1) responsibilities γ_{ts} . Consequently, the BCE loss becomes very high, as it heavily penalizes confident errors. Overconfidence can adversely affect the performance of DVBx, as the loss primarily drives parameter updates to mitigate overconfidence rather than improving diarization performance. To compensate for such behavior, we optionally calibrate the responsibilities by passing them through the same function as Eq. (5) right before calculating the loss. The corresponding parameter τ_{calib} is also trained to ‘‘please’’ the BCE loss.

Despite BCE being the standard loss for training diarization systems [21], it does not directly correlate with the diarization error rate (DER) (see Section 5). Therefore, we propose expected detection error loss (EDE):

$$H_E(\gamma_t^\phi, \hat{\mathbf{l}}_t) = \sum_{s=1}^S (1 - \gamma_{ts}^\phi) \hat{l}_{ts} + \gamma_{ts}^\phi (1 - \hat{l}_{ts}), \quad (8)$$

which calculates expected miss, if $\hat{l}_{ts} = 1$, and expected false alarm, if $\hat{l}_{ts} = 0$. The main difference between EDE and the DER metric is that EDE operates with probabilities rather than hard decisions, and it double-counts speaker confusion frames, as confusion can be viewed as both a miss and a false alarm from the perspective of different speakers. Note that EDE is not as sensitive to overconfident responsibilities as BCE; hence, it does not require calibration. However, if also calibrated by Eq. (5), the larger the positive calibration constant τ_{calib} , the closer the EDE is to (hard) detection error due to over-saturated responsibilities.

3.2. Backpropagation through VB Iterations

The straightforward way to backpropagate through VBx is to run VB inference for a predefined number of iterations, obtain the refined speaker labels after the last iteration, compute the loss, and backpropagate the gradients. However, this leads to vanishing gradients when using a high number of VB iterations.

As a solution, we compute the loss of Eq. (6) using the responsibilities after each VB iteration and average them to obtain the final loss value. It is mathematically equal to averaging the gradients across the VB iterations, which speeds up the model convergence due to easier gradient propagation.

3.3. Parameters

The following VBx parameters were considered for fine-tuning in this publication. In some cases, re-parametrizations were used to ensure they remain valid. F_A and F_B are trained directly. To ensure P_l is a valid probability, we optimize $\logit(P_l)$ and apply sigmoid afterward. To ensure positive τ , we train $\log(\tau)$ (the same for τ_{calib}). Regarding PLDA parameters, the transformation matrix \mathbf{E} is trained directly, and for the diagonal covariance matrix $\mathbf{\Phi}$, we train the log of the diagonal.

4. EXPERIMENTAL SETUP

4.1. Data

The proposed discriminative training does not require large amounts of training data, as only a few scalars are optimized, and the PLDA model is generatively pre-trained, so the discriminative training only fine-tunes its parameters.

To evaluate the method's performance, we use three datasets: DIHARDII [22], CALLHOME [23], and AMI Beamformed [24]. DIHARD II (DH) comes with a pre-defined split into two parts: development and testing. The development part is further split into two equal-sized speaker-disjoint (if possible) parts for training (8 hours) and validation (10 hours). For CALLHOME (CH), we use the well-adopted split¹ into two partitions: part1, part2, containing 249 and 250 files respectively. We further split part1 into two subsets: train (127 files - 4 hours) and validation (122 files - 3.5 hours), such that each part contains approximately the same number of files per each number of speakers (2-7). Lastly, we use the Full-corpus-ASR partition [2] (train set 65 hours, dev set 7.5 hours) in the case of AMI.

4.1.1. Labels

Like most standard clustering-based methods, VBx does not handle overlaps. Therefore, designing the ground truth (GT) labels of the overlap segments for discriminative training is not straightforward. First, we considered assigning each x -vector the GT label of

¹<https://github.com/BUTSpeechFIT/CALLHOME.sublists>

the dominant speaker in the segment, as this is what the VBx system can estimate in the optimal case. Next, we considered the GT labels containing the values $p \in \langle 0, 1 \rangle$ representing the proportion of each speaker's speech to the total amount of speech in the segment. In our preliminary experiments, the latter approach outperformed the former one in multiple training settings, so we decided to use it for the experiments presented in this paper.

4.2. Model

The model takes b (batch size) pre-extracted x -vector sequences, extracted by ResNet101 [25, 26], from b utterances and processes them in parallel. Each concurrent instance runs VBx on a single utterance (i.e., we do not require equal-length utterances in a single batch) and computes the gradients w.r.t. the trainable parameters. Subsequently, gradients from parallel instances are averaged, and the trainable parameters are updated accordingly.

4.2.1. Training Setup

We experimented with two training strategies: a) training the hyperparameters (F_A, F_B, τ, P_l) and then PLDA with hyperparameters fixed, and b) joint training of all parameters. If all are jointly trained, as PLDA has orders of magnitude more parameters than the number of hyperparameters, PLDA fine-tuning quickly decreases the loss but prevents hyperparameters from reaching optimal values. Therefore, we use the two-stage training strategy.

At the beginning of the inference, F_A and F_B are set to their theoretically correct value of 1. Furthermore, τ_{calib} default value is also set to 1. Since there are no theoretically correct values for P_l and τ , we considered several initial values, all of which had almost no effect on the model performance. Therefore, we set their values to the ones selected for the original VBx [2].

VB inference is run for 10 iterations. We set $b = 8$ and train the model for 500 epochs with Adam optimizer [27].

In the VBx model, the trainable parameters have different dynamic ranges, and small changes in some of them have a strong impact on performance. This could be easily observed as gradients w.r.t. F_A were much higher than w.r.t. other hyperparameters. Therefore, our model required searching for suitable learning rates (LR) for the different parameters. LR was set to $5e-4$ for F_A and to $1e-2$ for the rest of the hyperparameters. For PLDA matrices with many more trainable parameters, the LR was set to $1e-3$.

Experiments showed that P_l converged to 0 independently of its initial value. Moreover, we observed that training with $P_l = 0$ throughout the whole training process did not significantly decrease the model performance. Therefore, we set $P_l = 0$ in all DVVBx experiments. This further simplifies the inference, as the HMM becomes a Gaussian mixture model (GMM) (see Eq. 1). Moreover, training GMM VBx on longer utterances requires significantly less memory and is several times faster than HMM VBx.

During the two-stage training, we first select the best model after the hyperparameters training according to the best validation DER. Then, we train the PLDA parameters with hyperparameters being fixed and provide results based on the best validation DER score.

4.3. Baseline

As a first baseline, we selected the original VBx model [2], as this work directly extends it. Its hyperparameters F_A, F_B , and P_l were estimated using grid search separately for each dataset to minimize DER on the development set. Label smoothing, however, was hand-tuned to $\tau = 7$ and remained constant for all datasets.




Fig. 1. BCE and EDE behavior throughout 500 epochs on the CH train set.

As the HMM is simplified to a GMM in DVbX, we also include GMM VBx as a baseline. To obtain optimal results for this baseline, we re-ran the grid search for F_A , F_B .

4.4. Evaluation

We use oracle VAD (during training and evaluation) and 40 VB iterations with stopping criterion (i.e., stop VB if the evidence lower bound [15] stops increasing). The models are evaluated using DER implemented in pyannote [28] with the standard collar=0.0 s for AMI and DH and collar=0.25 s for CH.

5. RESULTS

5.1. Hyperparameters Fine-tuning

Table 1 presents the main experimental results. It first shows the performance of the baselines and their corresponding hyperparameter values. Secondly, it shows the results of hyperparameter training experiments comparing the three losses: BCE, BCE with calibration (BCE+calib.), and EDE.

As mentioned earlier, BCE optimization and DER performance do not correlate well, which is illustrated in Fig. 1. BCE requires using a validation set for early stopping, as parameters would eventually diverge, significantly harming the performance. Moreover, substantially lower F_A , F_B values are achieved by BCE because it tries to compensate for the overconfident responsibilities. BCE+calib. solves this issue but performs worse than both BCE and EDE. The proposed EDE not only provides reasonable hyperparameter values but converges faster than BCE (see Fig. 1) and achieves comparable or better DER in all datasets. Hence, further comparisons are conducted on models trained with EDE.

Furthermore, Table 1 shows that DVbX with EDE outperforms the GMM VBx baseline on all datasets while achieving comparable performance with the original HMM VBx baseline, which proves that DVbX can be used to successfully tune hyperparameters. It also shows that the GMM simplification does not harm the model performance much.

5.2. PLDA Fine-tuning

PLDA fine-tuning experiments are shown in Table 2. It first presents the following systems: a) GMM VBx baseline, b) discriminatively trained hyperparameters, c) fine-tuned PLDA with fixed hyperparameters obtained in b). While the fine-tuning obtains comparable performance to the baseline in DH and CH, for AMI, it yields an absolute 2.61% DER improvement w.r.t. the GMM VBx baseline and 1.92% DER improvement w.r.t. DVbX hyperparameters training. We hypothesize that, in AMI, DVbX can leverage the larger development set to fix the mismatch between the pre-trained PLDA and the target domain. Lastly, for the sake of completeness, we also tested PLDA FT using the GMM VBx baseline parameters, where a similar behavior was observed.

Table 1. DERs (%) using three losses: BCE, BCE with calibration, EDE on datasets: AMI, CH, DH. The first two rows in each dataset section display the baseline results.

Data	System	τ	F_A	F_B	DER
DH	HMM VBx [2]	7.00	0.20	6.00	18.55
	GMM VBx	7.00	0.20	5.00	18.93
	DVBx - BCE	2.90	0.25	4.38	18.98
	DVBx - BCE+calib.	12.88	0.43	10.14	18.84
	DVBx - EDE	9.62	0.33	9.64	<u>18.76</u>
CH	HMM VBx [2]	7.00	0.40	17.00	13.53
	GMM VBx	7.00	0.30	13.00	13.63
	DVBx - BCE	0.97	0.08	1.39	13.53
	DVBx - BCE+calib.	1.93	0.51	11.16	14.52
	DVBx - EDE	12.40	0.26	9.47	13.48
AMI	HMM VBx [2]	7.00	0.40	64.00	20.84
	GMM VBx	7.00	0.50	63.00	21.49
	DVBx - BCE	12.35	0.12	8.89	21.06
	DVBx - BCE+calib.	15.10	0.21	13.90	21.72
	DVBx - EDE	3.48	0.25	25.31	<u>20.91</u>

Table 2. DERs (%) of PLDA fine-tuning (FT).

System	DH	CH	AMI
a) GMM VBx	18.93	13.63	21.49
b) DVbX trained F_A, F_B, τ	18.76	13.48	20.91
c) b) + PLDA FT	18.66	13.38	18.99
d) a) + PLDA FT	18.93	13.63	18.88

6. CONCLUSION

In this work, we introduced a new principled way of tuning the VBx model hyperparameters without the need for running an extensive grid search requiring an understanding of hyperparameter sensitivity. To successfully do so, we introduced a new loss, EDE, that better correlates with the DER metric than the standard BCE loss. We want to highlight that while we used EDE only in the context of DVbX, its usage can be extended to EEND system training. Furthermore, we showed that we can further improve the model performance by fine-tuning the PLDA parameters. The presented work is the first step towards building a robust DVbX, combining the capabilities of a principled model with discriminative training. In future work, we plan to backpropagate gradients to the x-vector extractor to fine-tune together the whole diarization pipeline. Moreover, we would like to couple this method with the recently proposed well-performing MS-VBx [6], allowing the model to handle overlaps.

7. ACKNOWLEDGEMENTS

The work was partly supported by Czech Ministry of Interior project No. VJ01010108 "ROZKAZ", Czech National Science Foundation (GACR) project NEUREM3 No. 19-26934X, and Horizon 2020 Marie Skłodowska-Curie grant ESPERANTO, No. 101007666. Computing on IT4I supercomputer was supported by the Czech Ministry of Education, Youth and Sports through the e-INFRA CZ (ID:90254).

8. REFERENCES

- [1] Tae Jin Park, Kyu J Han, Manoj Kumar, and Shrikanth Narayanan, "Auto-tuning spectral clustering for speaker diarization using normalized maximum eigengap," *IEEE Signal Processing Letters*, vol. 27, pp. 381–385, 2019.
- [2] Federico Landini, Ján Profant, Mireia Diez, and Lukáš Burget, "Bayesian HMM clustering of x-vector sequences (VBx) in speaker diarization: Theory, implementation and analysis on standard tasks," *Computer Speech & Language*, vol. 71, pp. 101254, 2022.
- [3] Shota Horiguchi, Yusuke Fujita, Shinji Watanabe, Yawen Xue, and Kenji Nagamatsu, "End-to-End Speaker Diarization for an Unknown Number of Speakers with Encoder-Decoder Based Attractors," in *Proc. Interspeech 2020*, 2020, pp. 269–273.
- [4] Jiaming Wang, Zhihao Du, and Shiliang Zhang, "Told: a novel two-stage overlap-aware framework for speaker diarization," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [5] Hervé Bredin, "Pyannote. audio 2.1 speaker diarization pipeline: Principle, benchmark, and recipe," in *Proceedings for Interspeech*, 2023.
- [6] Marc Delcroix et al., "Multi-Stream Extension of Variational Bayesian HMM Clustering (MS-VBx) for Combined End-to-End and Vector Clustering-based Diarization," in *Proc. INTERSPEECH 2023*, 2023, pp. 3477–3481.
- [7] Desh Raj, Zili Huang, and Sanjeev Khudanpur, "Multi-class spectral clustering with overlaps for speaker diarization," in *2021 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 582–589.
- [8] Fan Yu, Shiliang Zhang, Yihui Fu, Lei Xie, Siqi Zheng, Zhihao Du, Weilong Huang, Pengcheng Guo, Zhijie Yan, Bin Ma, et al., "M2met: The icassp 2022 multi-channel multi-party meeting transcription challenge," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6167–6171.
- [9] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al., "Ego4d: Around the world in 3,000 hours of egocentric video," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18995–19012.
- [10] G. Cheng et al., "The conversational short-phrase speaker diarization (cssd) task: Dataset, evaluation metric and baselines," in *2022 13th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2022, pp. 488–492.
- [11] Shota Horiguchi, Paola Garcia, Yusuke Fujita, Shinji Watanabe, and Kenji Nagamatsu, "End-to-end speaker diarization as post-processing," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7188–7192.
- [12] Prachi Singh, Amrit Kaul, and Sriram Ganapathy, "Supervised hierarchical clustering using graph neural networks for speaker diarization," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [13] David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [14] Niko Brummer and Edward Villiers, "The speaker partitioning problem," *Proc. of Odyssey 2010*, 01 2010.
- [15] Christopher M. Bishop, *Pattern Recognition and Machine Learning*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [16] Natsuo Yamashita, Shota Horiguchi, and Takeshi Homma, "Improving the naturalness of simulated conversations for end-to-end neural diarization," in *Odyssey 2022*, 2022.
- [17] Federico Landini, Alicia Lozano-Diez, Mireia Diez, and Lukáš Burget, "From Simulated Mixtures to Simulated Conversations as Training Data for End-to-End Neural Diarization," in *Proc. Interspeech*, 2022, pp. 5095–5099.
- [18] Lukáš Burget, Oldřich Plchot, Sandro Cumani, Ondřej Glembek, Pavel Matějka, and Niko Brümmer, "Discriminatively trained probabilistic linear discriminant analysis for speaker verification," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2011, pp. 4832–4835.
- [19] Alan McCree, "Multiclass discriminative training of i-vector language recognition," *The Speaker and Language Recognition Workshop (Odyssey 2014)*, 2014.
- [20] Dong Yu, Morten Kolbæk, Zheng-Hua Tan, and Jesper Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 241–245.
- [21] Yusuke Fujita, Naoyuki Kanda, Shota Horiguchi, Kenji Nagamatsu, and Shinji Watanabe, "End-to-End Neural Speaker Diarization with Permutation-Free Objectives," in *Proc. Interspeech 2019*, 2019, pp. 4300–4304.
- [22] N. Ryant et al., "The Second DIHARD Diarization Challenge: Dataset, task, and baselines," in *Proceedings of Interspeech*, 2019.
- [23] A. F. Martin and M. A. Przybocki, "Speaker recognition in a multi-speaker environment," in *7th European Conference on Speech Communication and Technology, Eurospeech*, September 2001, vol. 7, num. 2, pp. 787–790.
- [24] Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al., "The AMI meeting corpus: A pre-announcement," in *International workshop on machine learning for multimodal interaction*. Springer, 2006, pp. 28–39.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [26] Hossein Zeinali, Shuai Wang, Anna Silnova, Pavel Matějka, and Oldřich Plchot, "But system description to voxceleb speaker recognition challenge 2019," in *Proceedings of The VoxCeleb Challenge Workshop 2019*, 2019, pp. 1–4.
- [27] Diederik P. Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR*, 2015.
- [28] Hervé Bredin, "pyannote.metrics: a toolkit for reproducible evaluation, diagnostic, and error analysis of speaker diarization systems," in *Proceedings of Interspeech 2017*, 2017.