# A Study on Decoupled Probabilistic Linear Discriminant Analysis

Di Wang\*<sup>†</sup>, Lantian Li\*, Hongzhi Yu<sup>†</sup>, Dong Wang\*

\* Center for Speech and Language Technologies, BNRist, Tsinghua University, China

E-mail: {wangdi,lilt}@cslt.org; wangdong99@mails.tsinghua.edu.cn

<sup>†</sup> Key Laboratory of China's Ethnic Languages and Information Technology of Ministry of Education,

Northwest Minzu University, China

E-mail: yhz@xbmu.cn

Abstract-Probabilistic linear discriminant analysis (PLDA) has broad application in open-set verification tasks, such as speaker verification. A key concern for PLDA is that the model is too simple (linear Gaussian) to deal with complicated data; however, the simplicity by itself is a major advantage of PLDA, as it leads to desirable generalization. An interesting research therefore is how to improve modeling capacity of PLDA while retaining the simplicity. This paper presents a decoupling approach, which involves a global model that is simple and generalizable, and a local model that is complex and expressive. While the global model holds a bird view on the entire data, the local model represents the details of individual classes. We conduct a preliminary study towards this direction and investigate a simple decoupling model including both the global and local models. The new model, which we call decoupled PLDA, is tested on a speaker verification task. Experimental results show that it consistently outperforms the vanilla PLDA when the model is based on raw speaker vectors. However, when the speaker vectors are processed by length normalization, the advantage of decoupled PLDA will be largely lost, suggesting future research on non-linear local models.

#### I. INTRODUCTION

Probabilistic linear discriminant analysis (PLDA) [1], [2] has been widely used as the scoring model in open-set verification tasks, such as speaker verification. It represents the data with a linear Gaussian model, where the between-class distribution is a Gaussian and the within-class distributions of individual classes are homogeneous Gaussians, i.e., Gaussians with a shared covariance.

An obvious weakness of the PLDA model is the underlying linear Gaussian assumption, which is too simple to represent data with complex distributions. However, simply employing a more complex model is not always recommended, as the generalization capability to unseen classes may be degraded [3], [4]. In fact, simplicity by itself is a key advantage of PLDA. An interesting research topic therefore is how to improve the model capacity, while retaining its simplicity.

In this paper, we present a decoupling approach, which views the data from two perspectives: a high-level and global perspective that focuses on the distribution of the entire data, including those unseen classes; and a low-level and local perspective that scrutinizes the distribution of individual classes. The high-level perspective relies on a rough model  $p_g$ , for which generalization to unseen classes is the most important,

while the low-level perspective relies on a detailed model  $p_l$ , by which an accurate within-class probability can be obtained. We combine  $p_g$  and  $p_l$  using the normalized likelihood (NL) formulation of PLDA [5], [6], [7], where  $p_g$  is employed to perform enrollment and normalization, while  $p_l$  is employed to perform prediction. This leads to a decoupled model that represents the rough/global distribution and detailed/local in a single statistical framework.

A key advantage of the decoupling approach is that the local model  $p_l$  can be very complex and so can represent complicated data. Investigation on this capability will be left for future work. In this study, we focus on a simple case where  $p_l$  remains a Gaussian, but independent of the global model  $p_g$ . The resultant model will be called **decoupled PLDA**. We tested the model on a speaker verification task, and the results demonstrated that even with this simple case, the decoupling approach can offer interesting performance improvement if the model was based on raw speaker vectors, though this advantage was largely lost if the speaker vectors were length normalized [8].

## II. THEORY

## A. Revisit PLDA

PLDA represents data by a linear Gaussian model. To simplify the presentation, we shall assume that the covariance matrix of the between-class distribution is diagonal and the shared covariance matrix of the within-class distributions of individual classes is an identify matrix. This can easily obtained by simultaneous diagonalization with a linear transform **W**. Put it more formally:

$$p(\boldsymbol{\mu}) = \mathcal{N}(\boldsymbol{\mu}; \mathbf{0}, \operatorname{diag}(\boldsymbol{\epsilon})) \tag{1}$$

$$p(\boldsymbol{x}|\boldsymbol{\mu}) = \mathcal{N}(\boldsymbol{x};\boldsymbol{\mu},\mathbf{I}).$$
(2)

where  $\boldsymbol{x}$  has been transformed by  $\mathbf{W}$  and  $\operatorname{diag}(\boldsymbol{\epsilon})$  is a diagonal matrix constructed from  $\boldsymbol{\epsilon}$ .

According to the hypothesis test theory [9], the following likelihood ratio (LR) has the highest power among other scores when determining if a test sample  $\boldsymbol{x}$  belongs to the class represented by the enrollment samples  $\{\boldsymbol{x}_1, ..., \boldsymbol{x}_n\}$ :

$$LR = \frac{p(\boldsymbol{x}, \boldsymbol{x}_1, \dots, \boldsymbol{x}_n)}{p(\boldsymbol{x})p(\boldsymbol{x}_1, \dots, \boldsymbol{x}_n)}.$$
(3)

It has been found that the decision based on PLDA LR is optimal in term of minimum Bayes risk (MBR) [7].

The parameters of the model involve  $\epsilon$  and W. They are usually optimized by maximizing likelihood (ML) with the following objective function:

$$\mathcal{L}(\boldsymbol{\epsilon}, \mathbf{W}) = \sum_{k=1}^{K} \ln p(\boldsymbol{x}_{1}^{k}, ..., \boldsymbol{x}_{n_{k}}^{k}), \qquad (4)$$

where  $\boldsymbol{x}_{i}^{k}$  denotes the *i*-th sample of the *k*-th class.

# B. Decoupling form

The LR form in (3) is equivalent to the following form [5], [6], [7]:

$$NL = \frac{\int p(\boldsymbol{x}|\boldsymbol{\mu}) p(\boldsymbol{\mu}|\boldsymbol{x}_1, \dots, \boldsymbol{x}_n) \mathrm{d}\boldsymbol{\mu}}{p(\boldsymbol{x})}.$$
 (5)

This was denoted by normalized likelihood (NL) in [7] in order to emphasize on the normalization role of the denominator. By this NL form, the LR score is decomposed into three components:

- The enrollment component p(μ|x<sub>1</sub>,...,x<sub>n</sub>) that produces the posterior of the class mean μ given the enrollment data x<sub>1</sub>,...,x<sub>n</sub>.
- The prediction component p(x|μ) that computes the probability that a test sample x belongs to a class represented by μ.
- The normalization component p(x) that computes the probability that x is produced by all potential classes.

Among the three components, the enrollment and normalization concern all potential classes and so are relevant to the global distribution, while the prediction concerns the distribution of a particular class and so is relevant to the local distribution. In the vanilla PLDA, all the three components are derived from the same underlying generative process, and so are *coupled* together. In cases where the generative model matches the data, the coupled model is optimal for scoring in terms of Bayes risk.

In practice, however, the data could be more complex than what the presumed generative model can represent. For example, in speaker verification, it was well known that both the between-class and within-class distributions of speaker vectors tend to be super-Gaussian [8], [10], [11], [12]. In this case, a linear Gaussian model such as PLDA is not sufficient. Nevertheless, designing a more complex generative model is not always ideal [3], [4], as it may jeopardize the model's generalization, another important concern for openset verification tasks.

The NL form provides a novel decoupling approach for the above dilemma. The key idea is that we never assume that the three scoring components are tied to a single generative process. Instead, we shall assume two generative processes, one for the global distribution, and the other for the local. Specifically, we hope all the inference related to unseen classes is based on a *global model* which should be simple, in order to ensure generalization, and the inference related to individual classes is based on a *local model* which could be complex, in order to ensure accuracy.

Among the three components in the NL form, the enrollment and normalization components are related to inference for unseen classes and so should be based on the global model, while the prediction component is for representing the within-class property and thus should be based on the local model. Intuitively, this decoupling approach can be understood as a multi-level representation for the data: a highlevel representation that describes the holistic property of the entire data, and a low-level representation that describes the details of a particular class.

# C. Decoupled PLDA

As the first study of the decoupling approach, we test a simple case in this paper: the global model is a linear Gaussian, and the local model is another independent Gaussian. This model, which we call decoupled PLDA, is very similar to the vanilla PLDA, with the only exception that the prediction component is never based on the within-class distribution learned by the global model, but a new Gaussian. Formally, assume the global model as follows:

$$p_g(\boldsymbol{\mu}) = \mathcal{N}(\boldsymbol{\mu}; \mathbf{0}, \operatorname{diag}(\boldsymbol{\epsilon}))$$
 (6)

$$p_g(\boldsymbol{x}|\boldsymbol{\mu}) = \mathcal{N}(\boldsymbol{x};\boldsymbol{\mu},\mathbf{I}), \tag{7}$$

and the local model supposes that after a linear transform **M**, the within-class distribution follows a Gaussian:

$$p_l(\mathbf{M}\boldsymbol{x}|\boldsymbol{\mu}) = \mathcal{N}(\mathbf{M}\boldsymbol{x};\boldsymbol{\mu},\mathbf{I}), \tag{8}$$

where  $\mathbf{M}$  is a transform matrix. The two models are integrated by the NL formulation, shown as follows:

$$NL = \frac{\int p_l(\mathbf{M}\boldsymbol{x}|\boldsymbol{\mu}) p_g(\boldsymbol{\mu}|\boldsymbol{x}_1,...,\boldsymbol{x}_n) \mathrm{d}\boldsymbol{\mu}}{p_a(\boldsymbol{x})}.$$
 (9)

The decoupled PLDA can be trained following the maximum likelihood (ML) criterion. The training involves two steps:

- **Global model training**: The same as training vanilla PLDA, following the ML criterion. More details can be found in the original paper [1], [2].
- Local model training: Optimize the linear transform M by maximizing the likelihood function as follows:

$$\mathcal{L}(\mathbf{M}) = \prod_{k}^{K} \prod_{i=1}^{n_{k}} \int p_{l}(\mathbf{M}\boldsymbol{x}_{i}^{k}|\boldsymbol{\mu}) p_{g}(\boldsymbol{\mu}|\boldsymbol{x}_{1}^{k},...,\boldsymbol{x}_{n_{k}}^{k}) \mathrm{d}\boldsymbol{\mu}$$

$$= \prod_{k}^{K} \prod_{i=1}^{n_{k}} \int \mathcal{N}(\mathbf{M}\boldsymbol{x}_{i}^{k};\boldsymbol{\mu},\mathbf{I}) p_{g}(\boldsymbol{\mu}|\boldsymbol{x}_{1}^{k},...,\boldsymbol{x}_{n_{k}}^{k}) \mathrm{d}\boldsymbol{\mu}$$

$$= \prod_{k}^{K} \prod_{i=1}^{n_{k}} \mathcal{N}(\mathbf{M}\boldsymbol{x}_{i}^{k};\frac{n_{k}\boldsymbol{\epsilon}}{n_{k}\boldsymbol{\epsilon}+1}\bar{\boldsymbol{x}}_{k},(1+\frac{\boldsymbol{\epsilon}}{n_{k}\boldsymbol{\epsilon}+1})\mathbf{I}),$$
(10)

where  $\bar{x}_k$  is the average of the samples in the k-th class.

# D. Discussion

1) Difference from PLDA: One may argue that since the decoupled PLDA and vanilla PLDA are trained using the same data and follow the same model assumption (both  $p_g$  and  $p_l$  are Gaussians), they should be identical when optimized, i.e.  $\mathbf{M} = \mathbf{I}$ . This intuition however is not correct, due to two main reasons. Firstly, since  $p_l$  and  $p_g$  are independent, the posterior derived from  $p_g$ , i.e.,  $p_g(\boldsymbol{\mu}|\mathbf{x}_1^k,...,\mathbf{x}_{n_k}^k)$ , cannot be regarded as the posterior of  $\boldsymbol{\mu}$  in the local model. Instead, it should be regarded as a *prior* for the distribution of  $\boldsymbol{\mu}$  for each data sample. With this prior, the local model is trained by improving the likelihood of the data on the local model. This is the key point of the decoupled model, where we assume two random processes, rather than a single one as in the vanilla PLDA.

Secondly, when training the local model, we treat all the samples are independent, no matter whether they belong to the same class or not. This is clear from the form of  $\mathcal{L}(\mathbf{M})$  presented in (10), where the integration over  $\boldsymbol{\mu}$  is conducted for each training sample. Note that if it is not the case, the objective would be as follows:

$$\mathcal{L}(\mathbf{M}) = \prod_{k}^{K} \int \left[ \prod_{i=1}^{n_{k}} p_{l}(\mathbf{M}\boldsymbol{x}_{i}^{k} | \boldsymbol{\mu}) \right] p_{g}(\boldsymbol{\mu} | \boldsymbol{x}_{1}^{k}, ..., \boldsymbol{x}_{n_{k}}^{k}) \mathrm{d}\boldsymbol{\mu}.$$
(11)

The reason of assuming independent samples is that it matches the scenario in the *test* phase, where all the test samples are independent. The objective for the local model, therefore, essentially encourages positive samples obtaining higher likelihood with the hypothesized class means (inferred by the global model) during test.

2) Single-set training: One design choice for the decoupled PLDA is to split the training data into an enrollment set and a test set, and train the global and local models using the two sets respectively. This would simulate the enroll-test scenario in a more realistic way. In this study, we choose a *single-set* scheme where the global and local models are trained using the same dataset, partly due to the limited training data in our experiments.

3) Early stop for unideal normalization: A key problem of the decoupled PLDA is that with the increased likelihood  $p_l(\mathbf{x}|\boldsymbol{\mu})$  computed by the local model, the normalization term  $p_g(\mathbf{x})$  computed by the global model will be not theoretically

correct. In this preliminary study, an empirical solution was used to address this problem: we monitor the performance of the model on the *training set* when training the local model, and stop the training process when the performance reaches the best. The resultant model trades off the gain between the empowered local model and the loss caused by the unideal normalization. More theoretical solution will be left for future work.

## III. RELATED WORK

Numerous techniques have been proposed to alleviate the deficiency of vanilla PLDA in representing complicated data. For example, the heavy-tailed PLDA [10] uses a Student's t distribution to handle long-tail distributions. Mixture PLDA was proposed to handle the diverse statistics under different conditions [13], [14], [15]. Deep normalization was proposed to gaussianize the distribution of speaker vectors [11], which was further extended to a neural discriminant analysis (NDA) [4]. Although promising, few of the techniques can improve modeling capacity while maintaining simplicity.

# IV. EXPERIMENTS

We evaluate decoupled PLDA by a speaker verification task. Speaker verification (SV) targets for determining whether a speech segment is spoken by a claimed speaker or not [16], [17], [18]. Modern SV approaches firstly convert variable-length speech segments to fixed-dimensional vectors (a.k.a., speaker vectors), and then employ PLDA to score the speaker vectors.

Two types of speaker vectors are most popular: i-vector [19] and x-vector [20]. The former is derived based on a statistical model, while the latter is based on deep neural networks [21], [22], [23]. In spite of other explorations presented recently [24], [25], [26], [27], [28], [29], [30], the i-vector/x-vector plus PLDA architecture is widely recognized as the state-of-the-art, partly due to the standard recipe published in the Kaldi toolkit [31]. In this paper, we will evaluate decoupled PLDA (dePLDA for short) with both i-vector and x-vector, and compare its performance with that of vanilla PLDA.

# A. Data

**VoxCeleb** [32], [33]: An open-source speaker dataset collected from media sources by University of Oxford. This dataset contains 2,000+ hours of speech signals from 7,000+ speakers. This dataset was used to train the i-vector and x-vector models, as well as the PLDA model used in the test on the SITW dataset.

**SITW** [34]: A standard evaluation dataset consists of 299 speakers. The core-core trials built on the SITW.Dev and SITW.Eval sets were used for evaluation.

**AIShell-1** [35]: An open-source Chinese Mandarin speech dataset. The training set (used for PLDA and decoupled PLDA training) involves 360, 897 utterances from 340 speakers, and the test set involves 64, 495 utterances from 60 speakers.

**HI-MIA** [36]: An open-source text-dependent speaker recognition dataset. All the speech utterances contain the word

		i·	vector		х	-vector	
Model	Dim	SITW.Eval	AISHELL-1	HI-MIA   Dim	SITW.Eval	AISHELL-1	HI-MIA
PLDA	400	8.201	1.179	1.008 512	6.862	1.146	0.698
PLDA	128	7.381	1.372	1.318 128	5.604	1.368	0.930
dePLDA	400	7.627	1.075	<b>0.698</b> 512	5.820	0.792	0.465
dePLDA	128	6.998	1 278	1 240 128	4.950	1.028	0 775

TABLE IEER(%) RESULTS WITH PLDA AND DEPLDA.

'Hi MIA', recorded by a microphone 1 meter away from the speaker. The training set (used for PLDA and decoupled PLDA training) involves 15,186 utterances from 254 speakers, and the test set involves 5,048 utterances from 86 speakers.

## B. Experimental setup

The i-vector and x-vector models were implemented using the Kaldi toolkit [31], following the SITW recipe. Although the recipe does not produce the best performance known so far on the SITW dataset, it is the most popular and reproducible, and so suitable for concept demonstration in our preliminary study.

The dimensionality of i-vectors and x-vectors was set to 400 and 512, respectively. The PLDA and dePLDA models were trained based on the produced i-vectors and x-vectors. PLDA was trained using the Kaldi toolkit, and decoupled PLDA was implemented by PyTorch<sup>1</sup>. The transform matrix M in decoupled PLDA was constrained to be diagonal and was initialized as an identity matrix. The Adam optimizer [37] was used to optimize M, and the best M is selected using SITW.Dev for the SITW test, and their own training sets for the AIShell-1 test and HI-MIA test. The selection is based on the performance in terms of the equal error rate (EER).

# C. Training process

The change of the EER results on SITW.Dev core-core trials with dePLDA trained on VoxCeleb are reported in Fig. 1. It can be seen that the performance was improved with the training proceeding, and reached the best result after 10 iterations. After that, more iterations lead to performance degradation. This is the expected pattern, and reflects the trade-off between the empowered local model and the mismatch between the local and global models.

# D. Main results

The EER results on SITW.Eval, AISHELL-1 and HI-MIA are reported in Table I. Two configurations were tested: (1) in the full-dimensional case, the PLDA/dePLDA models were trained with the raw speaker vectors; (2) in the low-dimensional case, the dimensionality of the speaker vectors was firstly reduced to 128 by LDA (as recommended in the Kaldi recipe), and then the PLDA/dePLDA models were trained using the low-dimensional speaker vectors. The results show that with both the configurations, dePLDA outperforms



Fig. 1. The trend of the EER results in the dePLDA training process. Note that at the beginning  $\mathbf{M}$  is initialized as  $\mathbf{I}$ , which is equivalent to the vanilla PLDA.

PLDA in a consistent way, demonstrating that the decoupling approach is rational and effective. Note that the dimension reduction does not improve the performance in the AISHELL-1 and HI-MIA tests, which we attribute to the unreliable LDA model with the limited training data, as well as the suboptimal dimensionality (128) that was tuned on SITW.

## E. Inconsistence with length normalization

Length normalization (LN) [8] is a standard preprocessing in speaker verification. It projects speaker vectors to a spherical surface, so that the (marginal) distribution tends to be more Gaussian. This simple operation has been demonstrated highly effective for both i-vector systems and x-vector systems [8]. Unfortunately, LN and dePLDA are inconsistent, as the local model shrinks the speaker vectors, which may counteract the effect of LN.

A possibility is to use a *partial* LN, i.e., conducting LN within the global model only. This is not theoretically correct as it may cause additional mismatch between the local model and the global model. Nevertheless, we found the partial LN may lead to performance gains in some cases.

The results with i-vectors and x-vectors are shown in Table II and Table III respectively. It can be seen that LN consistently improved the performance of the vanilla PLDA in all situations. The partial LN, however, lead to improved performance with dePLDA in some cases, but in other cases, the performance was degraded. Accompanied by LN, vanilla PLDA generally outperformed dePLDA, although in some situations dePLDA accompanied by partial LN showed better

<sup>&</sup>lt;sup>1</sup>https://gitlab.com/csltstu/enroll-test-mismatch

TABLE II EER(%) results with PLDA and dePLDA based on the i-vector model.

Model	Dim	LN	SITW.Eval	AISHELL-1	HI-MIA
PLDA PLDA PLDA PLDA	400 400 128 128	√ √	8.201 6.315 7.381 <b>5.878</b>	1.179 1.071 1.372 1.335	1.008 <b>0.698</b> 1.318 1.008
dePLDA dePLDA dePLDA dePLDA	400 400 128 128	√ √	7.627 6.014 6.998 5.987	1.075 <b>1.056</b> 1.278 1.349	<b>0.698</b> 0.775 1.240 0.775

TABLE III EER(%) results with PLDA and dePLDA based on the x-vector model.

Model	Dim	LN	SITW.Eval	AISHELL-1	HI-MIA
PLDA PLDA PLDA PLDA	512 512 128 128	√ √	6.862 4.511 5.604 <b>3.636</b>	1.146 0.806 1.368 1.019	0.698 0.698 0.930 0.620
dePLDA dePLDA dePLDA dePLDA	512 512 128 128	√ √	5.820 4.292 4.950 4.019	0.792 <b>0.740</b> 1.028 1.198	<b>0.465</b> 0.853 0.775 0.620

performance (e.g., results on SITW with 512-dim x-vectors or 400-dim i-vectors).

The loss of the comparative advantage of dePLDA when LN is employed can be attributed to the improved Gaussianality of the length-normalized data. Notice that if the data is truly linear Gaussian and perfectly matches the PLDA assumption, the vanilla PLDA is optimal. The decoupling approach can contribute only if the data is complex and deviate from linear Gaussian. LN arguably regulates the speaker vectors and makes them more Gaussian, which may reduce the potential contribution of the decoupling approach.

## V. CONCLUSIONS

This paper reported a preliminary study towards decoupled scoring for open-set verification tasks, where a global model is designed to provide a bird view for the entire profile of the data, and a local model to provide a detailed view for individual classes. To verify its effectiveness, we tested a simple decoupled PLDA model (dePLDA), where both the global and local models are Gaussian. Compared to the vanilla PLDA, the new model uses a separate Gaussian to match the within-class distribution, and simulates sample independence in the test phase.

We verify the decoupled PLDA model with a speaker verification task on three datasets. The results show that working on raw speaker vectors, the decoupled PLDA consistently outperformed the vanilla PLDA, confirming that the decoupling approach is effective. However, when length normalization is employed, the comparative advantage of dePLDA was largely lost. This is probably attributed to the improved Gaussianality of the speaker vectors after length normalization, which diminishes the potential of the decoupling approach. Nevertheless, our experiments indeed demonstrated that the decoupling approach is rational and effective, and the performance might be significantly improved if the local model is more expressive, especially when it is non-Gaussian.

#### ACKNOWLEDGMENT

This work was supported by the National Natural Science Foundation of China (NSFC) under Grants No.61633013 and No.62171250, and also the Huawei Innovation Research Program under Cross-Device Speaker Recognition Project Contract No.YBN2019125091.

# REFERENCES

- S. Ioffe, "Probabilistic linear discriminant analysis," in European Conference on Computer Vision. Springer, 2006, pp. 531–542.
- [2] S. J. Prince and J. H. Elder, "Probabilistic linear discriminant analysis for inferences about identity," in 2007 IEEE 11th International Conference on Computer Vision. IEEE, 2007, pp. 1–8.
- [3] J. Villalba, N. Brümmer, and N. Dehak, "Tied variational autoencoder backends for i-vector speaker recognition." in *INTERSPEECH*, 2017, pp. 1004–1008.
- [4] L. Li, D. Wang, and T. F. Zheng, "Neural discriminant analysis for deep speaker embedding," arXiv preprint arXiv:2005.11905, 2020.
- [5] B. J. Borgström and A. McCree, "Discriminatively trained bayesian speaker comparison of i-vectors," in 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, 2013, pp. 7659–7662.
- [6] A. McCree, G. Sell, and D. Garcia-Romero, "Extended variability modeling and unsupervised adaptation for PLDA speaker recognition." in *INTERSPEECH*, 2017, pp. 1552–1556.
- [7] D. Wang, "Remarks on optimal scores for speaker recognition," arXiv preprint arXiv:2010.04862, 2020.
- [8] D. Garcia-Romero and C. Y. Espy-Wilson, "Analysis of i-vector length normalization in speaker recognition systems," in *INTERSPEECH*, 2011.
- [9] J. Neyman and E. S. Pearson, "Ix. on the problem of the most efficient tests of statistical hypotheses," *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 231, no. 694-706, pp. 289–337, 1933.
- [10] P. Kenny, "Bayesian speaker verification with heavy-tailed priors." in Proceedings of Odyssey: The Speaker and Language Recognition Workshop, 2010, p. 14.
- [11] Y. Cai, L. Li, A. Abel, X. Zhu, and D. Wang, "Deep normalization for speaker vectors," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 733–744, 2020.
- [12] D. Wang, "A simulation study on optimal scores for speaker recognition," EURASIP Journal on Audio, Speech, and Music Processing, vol. 2020, no. 1, pp. 1–23, 2020.
- [13] M.-W. Mak, X. Pang, and J.-T. Chien, "Mixture of PLDA for noise robust i-vector speaker verification," *IEEE/ACM Transactions on Audio*, *Speech, and Language Processing*, vol. 24, no. 1, pp. 130–142, 2015.
- [14] J. Xie, S. Sia, P. Garcia, D. Povey, and S. Khudanpur, "Mixture of speaker-type PLDAs for children's speech diarization," arXiv preprint arXiv:2008.13213, 2020.
- [15] M. Senoussaoui, P. Kenny, N. Brümmer, E. d. Villiers, and P. Dumouchel, "Mixture of PLDA models in i-vector space for genderindependent speaker recognition," in *INTERSPEECH*, 2011.
- [16] J. P. Campbell, "Speaker recognition: A tutorial," *Proceedings of the IEEE*, vol. 85, no. 9, pp. 1437–1462, 1997.
- [17] D. A. Reynolds, "An overview of automatic speaker recognition technology," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 4. IEEE, 2002, pp. IV–4072.
- [18] J. H. Hansen and T. Hasan, "Speaker recognition by machines and humans: A tutorial review," *IEEE Signal processing magazine*, vol. 32, no. 6, pp. 74–99, 2015.
- [19] N. Dehak, P. J. Kenny, R. Dehak, P. Dumouchel, and P. Ouellet, "Frontend factor analysis for speaker verification," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 4, pp. 788–798, 2011.

- [20] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing* (ICASSP). IEEE, 2018, pp. 5329–5333.
- [21] L. Deng and D. Yu, "Deep learning: methods and applications," Foundations and trends in signal processing, vol. 7, no. 3–4, pp. 197–387, 2014.
- [22] E. Variani, X. Lei, E. McDermott, I. L. Moreno, and J. Gonzalez-Dominguez, "Deep neural networks for small footprint text-dependent speaker verification," in *IEEE International Conference on Acoustics*, *Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4052–4056.
- [23] L. Li, Y. Chen, Y. Shi, Z. Tang, and D. Wang, "Deep speaker feature learning for text-independent speaker verification," in *INTERSPEECH*, 2017, pp. 1542–1546.
- [24] G. Heigold, I. Moreno, S. Bengio, and N. Shazeer, "End-to-end textdependent speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5115–5119.
- [25] S.-X. Zhang, Z. Chen, Y. Zhao, J. Li, and Y. Gong, "End-to-end attention based text-dependent speaker verification," in *Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 171–178.
- [26] F. R. Rahman Chowdhury, Q. Wang, I. L. Moreno, and L. Wan, "Attention-based models for text-dependent speaker verification," in *IEEE International Conference on Acoustics, Speech and Signal Pro*cessing (ICASSP), 2018, pp. 5359–5363.
- [27] W. Cai, J. Chen, and M. Li, "Exploring the encoding layer and loss function in end-to-end speaker and language recognition system," in *Proceedings of Odyssey: The Speaker and Language Recognition* Workshop, 2018, pp. 74–81.
- [28] J. weon Jung, H.-S. Heo, J. ho Kim, H. jin Shim, and H.-J. Yu, "RawNet: Advanced end-to-end deep neural network using raw waveforms for textindependent speaker verification," in *INTERSPEECH*, 2019, pp. 1268– 1272.
- [29] S. Wang, J. Rohdin, L. Burget, O. Plchot, Y. Qian, K. Yu, and J. Cernocky, "On the usage of phonetic information for text-independent speaker embedding extraction," in *INTERSPEECH*, 2019, pp. 1148– 1152.
- [30] J. Wang, K.-C. Wang, M. T. Law, F. Rudzicz, and M. Brudno, "Centroid-based deep metric learning for speaker recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3652–3656.
- [31] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz et al., "The Kaldi speech recognition toolkit," in *IEEE workshop on automatic speech* recognition and understanding, 2011.
- [32] A. Nagrani, J. S. Chung, and A. Zisserman, "VoxCeleb: a large-scale speaker identification dataset," in *INTERSPEECH*, 2017.
- [33] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *INTERSPEECH*, 2018, pp. 1086–1090.
- [34] M. McLaren, L. Ferrer, D. Castan, and A. Lawson, "The speakers in the wild (SITW) speaker recognition database." in *INTERSPEECH*, 2016, pp. 818–822.
- [35] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHELL-1: An opensource mandarin speech corpus and a speech recognition baseline," in *Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*. IEEE, 2017, pp. 1–5.
- [36] X. Qin, H. Bu, and M. Li, "HI-MIA: A far-field text-dependent speaker verification database and the baselines," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7609–7613.
- [37] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," arXiv preprint arXiv:1412.6980, 2014.